

# Xenomai

3.1

Generated by Doxygen 1.8.14



# Contents

<b>1</b>	<b>API service tags</b>	<b>1</b>
<b>2</b>	<b>Deprecated List</b>	<b>3</b>
<b>3</b>	<b>Module Index</b>	<b>5</b>
3.1	Modules . . . . .	5
<b>4</b>	<b>Data Structure Index</b>	<b>9</b>
4.1	Data Structures . . . . .	9
<b>5</b>	<b>File Index</b>	<b>13</b>
5.1	File List . . . . .	13
<b>6</b>	<b>Module Documentation</b>	<b>23</b>
6.1	Channels and ranges . . . . .	23
6.1.1	Detailed Description . . . . .	25
6.2	Big dual kernel lock . . . . .	27
6.2.1	Detailed Description . . . . .	27
6.2.2	Macro Definition Documentation . . . . .	27
6.2.2.1	cobalt_atomic_enter . . . . .	27
6.2.2.2	cobalt_atomic_leave . . . . .	28
6.2.2.3	RTDM_EXECUTE_ATOMICALY . . . . .	29
6.3	Spinlock with preemption deactivation . . . . .	30
6.3.1	Detailed Description . . . . .	31
6.3.2	Macro Definition Documentation . . . . .	31
6.3.2.1	rtdm_lock_get_irqsave . . . . .	31

---

6.3.2.2	<code>rtm_lock_irqrestore</code>	31
6.3.2.3	<code>rtm_lock_irqsave</code>	31
6.3.3	Function Documentation	32
6.3.3.1	<code>rtm_lock_get()</code>	32
6.3.3.2	<code>rtm_lock_init()</code>	32
6.3.3.3	<code>rtm_lock_put()</code>	33
6.3.3.4	<code>rtm_lock_put_irqrestore()</code>	33
6.4	User-space driver core	34
6.4.1	Detailed Description	35
6.4.2	Macro Definition Documentation	35
6.4.2.1	<code>UDD_IRQ_CUSTOM</code>	36
6.4.2.2	<code>UDD_IRQ_NONE</code>	36
6.4.2.3	<code>UDD_MEM_LOGICAL</code>	36
6.4.2.4	<code>UDD_MEM_NONE</code>	36
6.4.2.5	<code>UDD_MEM_PHYS</code>	36
6.4.2.6	<code>UDD_MEM_VIRTUAL</code>	37
6.4.2.7	<code>UDD_RTIOC_IRQDIS</code>	37
6.4.2.8	<code>UDD_RTIOC_IRQEN</code>	37
6.4.2.9	<code>UDD_RTIOC_IRQSIG</code>	37
6.4.3	Function Documentation	38
6.4.3.1	<code>udd_disable_irq()</code>	38
6.4.3.2	<code>udd_enable_irq()</code>	38
6.4.3.3	<code>udd_get_device()</code>	39
6.4.3.4	<code>udd_notify_event()</code>	40
6.4.3.5	<code>udd_register_device()</code>	40
6.4.3.6	<code>udd_unregister_device()</code>	41
6.5	Thread runtime statistics	42
6.5.1	Detailed Description	42
6.6	Thread state flags	43
6.6.1	Detailed Description	44

---

---

6.6.2	Macro Definition Documentation	44
6.6.2.1	XNHELD	44
6.6.2.2	XNMIGRATE	44
6.6.2.3	XNPEND	44
6.6.2.4	XNREADY	45
6.6.2.5	XNSUSP	45
6.6.2.6	XNTRAPLB	45
6.7	Thread information flags	46
6.7.1	Detailed Description	46
6.8	CAN Devices	47
6.8.1	Detailed Description	54
6.8.2	Macro Definition Documentation	57
6.8.2.1	CAN_CTRLMODE_3_SAMPLES	58
6.8.2.2	CAN_CTRLMODE_LISTENONLY	58
6.8.2.3	CAN_CTRLMODE_LOOPBACK	58
6.8.2.4	CAN_ERR_LOSTARB_UNSPEC	58
6.8.2.5	CAN_RAW_ERR_FILTER	58
6.8.2.6	CAN_RAW_FILTER	59
6.8.2.7	CAN_RAW_LOOPBACK	60
6.8.2.8	CAN_RAW_RECV_OWN_MSGS	61
6.8.2.9	RTCAN_RTIOC_RCV_TIMEOUT	61
6.8.2.10	RTCAN_RTIOC_SND_TIMEOUT	61
6.8.2.11	RTCAN_RTIOC_TAKE_TIMESTAMP	62
6.8.2.12	SIOCGCANBAUDRATE	63
6.8.2.13	SIOCGCANCTRLMODE	63
6.8.2.14	SIOCGCANCUSTOMBITTIME	64
6.8.2.15	SIOCGCANSTATE	64
6.8.2.16	SIOCGIFINDEX	65
6.8.2.17	SIOCSCANBAUDRATE	66
6.8.2.18	SIOCSCANCTRLMODE	66

---

---

6.8.2.19	SIOCSCANCUSTOMBITTIME	67
6.8.2.20	SIOCSCANMODE	68
6.8.2.21	SOL_CAN_RAW	69
6.8.3	Typedef Documentation	69
6.8.3.1	can_filter_t	69
6.8.3.2	can_frame_t	70
6.8.4	Enumeration Type Documentation	70
6.8.4.1	CAN_BITTIME_TYPE	70
6.8.4.2	CAN_MODE	70
6.8.4.3	CAN_STATE	71
6.9	RTDM	72
6.9.1	Detailed Description	73
6.9.2	Macro Definition Documentation	73
6.9.2.1	RTDM_TIMEOUT_INFINITE	73
6.9.2.2	RTDM_TIMEOUT_NONE	73
6.9.3	Typedef Documentation	73
6.9.3.1	nanosecs_abs_t	74
6.9.3.2	nanosecs_rel_t	74
6.10	RTDM User API	75
6.10.1	Detailed Description	75
6.11	Serial Devices	76
6.12	Testing Devices	78
6.13	Real-time IPC	79
6.13.1	Detailed Description	81
6.13.2	Macro Definition Documentation	81
6.13.2.1	BUFP_BUFSZ	81
6.13.2.2	BUFP_LABEL	82
6.13.2.3	IDDP_LABEL	83
6.13.2.4	IDDP_POOLSZ	84
6.13.2.5	SO_RCVTIMEO	85

---

---

6.13.2.6	SO_SNDTIMEO	85
6.13.2.7	XDDP_BUFSZ	85
6.13.2.8	XDDP_EVTDOWN	86
6.13.2.9	XDDP_EVTIN	86
6.13.2.10	XDDP_EVTNOBUF	87
6.13.2.11	XDDP_EVTOUT	87
6.13.2.12	XDDP_LABEL	87
6.13.2.13	XDDP_MONITOR	88
6.13.2.14	XDDP_POOLSZ	89
6.13.3	Enumeration Type Documentation	89
6.13.3.1	anonymous enum	89
6.13.4	Function Documentation	90
6.13.4.1	bind__AF_RTIPC()	90
6.13.4.2	close__AF_RTIPC()	92
6.13.4.3	connect__AF_RTIPC()	92
6.13.4.4	getpeername__AF_RTIPC()	93
6.13.4.5	getsockname__AF_RTIPC()	94
6.13.4.6	getsockopt__AF_RTIPC()	94
6.13.4.7	recvmsg__AF_RTIPC()	94
6.13.4.8	sendmsg__AF_RTIPC()	95
6.13.4.9	setsockopt__AF_RTIPC()	96
6.13.4.10	socket__AF_RTIPC()	97
6.14	Asynchronous Procedure Calls	98
6.14.1	Detailed Description	98
6.14.2	Function Documentation	98
6.14.2.1	xnapc_alloc()	99
6.14.2.2	xnapc_free()	99
6.14.2.3	xnapc_schedule()	100
6.15	In-kernel arithmetics	101
6.15.1	Detailed Description	101

---

---

6.15.2 Function Documentation . . . . .	101
6.15.2.1 xnarch_generic_full_divmod64() . . . . .	101
6.16 Buffer descriptor . . . . .	102
6.16.1 Detailed Description . . . . .	103
6.16.2 Function Documentation . . . . .	104
6.16.2.1 xnbufd_copy_from_kmem() . . . . .	104
6.16.2.2 xnbufd_copy_to_kmem() . . . . .	105
6.16.2.3 xnbufd_invalidate() . . . . .	106
6.16.2.4 xnbufd_map_kread() . . . . .	107
6.16.2.5 xnbufd_map_kwrite() . . . . .	107
6.16.2.6 xnbufd_map_uread() . . . . .	108
6.16.2.7 xnbufd_map_uwrite() . . . . .	108
6.16.2.8 xnbufd_reset() . . . . .	109
6.16.2.9 xnbufd_unmap_kread() . . . . .	109
6.16.2.10xnbufd_unmap_kwrite() . . . . .	110
6.16.2.11xnbufd_unmap_uread() . . . . .	110
6.16.2.12xnbufd_unmap_uwrite() . . . . .	111
6.17 Clock services . . . . .	112
6.17.1 Detailed Description . . . . .	112
6.17.2 Function Documentation . . . . .	112
6.17.2.1 xnclock_adjust() . . . . .	112
6.17.2.2 xnclock_deregister() . . . . .	113
6.17.2.3 xnclock_register() . . . . .	113
6.17.2.4 xnclock_tick() . . . . .	114
6.18 Debugging services . . . . .	115
6.18.1 Detailed Description . . . . .	115
6.19 Dynamic memory allocation services . . . . .	116
6.19.1 Detailed Description . . . . .	116
6.19.2 Function Documentation . . . . .	116
6.19.2.1 xnheap_alloc() . . . . .	116



---

6.19.2.2	xnheap_destroy()	117
6.19.2.3	xnheap_free()	117
6.19.2.4	xnheap_set_name()	118
6.20	Cobalt	119
6.20.1	Detailed Description	119
6.21	Cobalt kernel	120
6.21.1	Detailed Description	121
6.21.1.1	Dual kernel service tags	121
6.22	Interrupt management	123
6.22.1	Detailed Description	123
6.22.2	Function Documentation	123
6.22.2.1	xnintr_affinity()	124
6.22.2.2	xnintr_attach()	125
6.22.2.3	xnintr_destroy()	126
6.22.2.4	xnintr_detach()	126
6.22.2.5	xnintr_disable()	127
6.22.2.6	xnintr_enable()	127
6.22.2.7	xnintr_init()	127
6.23	Locking services	130
6.23.1	Detailed Description	130
6.23.2	Macro Definition Documentation	130
6.23.2.1	splexit	130
6.23.2.2	splhigh	131
6.23.2.3	spltest	131
6.24	Lightweight key-to-object mapping service	132
6.24.1	Detailed Description	132
6.24.2	Function Documentation	133
6.24.2.1	xnmap_create()	133
6.24.2.2	xnmap_delete()	133
6.24.2.3	xnmap_enter()	134

---

6.24.2.4	xnmap_fetch()	135
6.24.2.5	xnmap_fetch_nocheck()	135
6.24.2.6	xnmap_remove()	136
6.25	Registry services	137
6.25.1	Detailed Description	137
6.25.2	Function Documentation	137
6.25.2.1	xnregistry_bind()	138
6.25.2.2	xnregistry_enter()	139
6.25.2.3	xnregistry_lookup()	140
6.25.2.4	xnregistry_remove()	140
6.25.2.5	xnregistry_unlink()	141
6.26	Driver programming interface	142
6.26.1	Detailed Description	142
6.27	Driver to driver services	143
6.27.1	Detailed Description	144
6.27.2	Function Documentation	144
6.27.2.1	rtdm_accept()	144
6.27.2.2	rtdm_bind()	145
6.27.2.3	rtdm_close()	146
6.27.2.4	rtdm_connect()	147
6.27.2.5	rtdm_getpeername()	148
6.27.2.6	rtdm_getsockname()	149
6.27.2.7	rtdm_getsockopt()	149
6.27.2.8	rtdm_ioctl()	150
6.27.2.9	rtdm_listen()	151
6.27.2.10	rtdm_open()	152
6.27.2.11	rtdm_read()	153
6.27.2.12	rtdm_recv()	153
6.27.2.13	rtdm_recvfrom()	154
6.27.2.14	rtdm_recvmsg()	155

---

6.27.2.15	<code>rtm_send()</code>	156
6.27.2.16	<code>rtm_sendmsg()</code>	156
6.27.2.17	<code>rtm_sendto()</code>	157
6.27.2.18	<code>rtm_setsockopt()</code>	158
6.27.2.19	<code>rtm_shutdown()</code>	159
6.27.2.20	<code>rtm_socket()</code>	160
6.27.2.21	<code>rtm_write()</code>	160
6.28	Device Profiles	162
6.28.1	Detailed Description	164
6.28.2	Macro Definition Documentation	164
6.28.2.1	<code>RTIOC_DEVICE_INFO</code>	164
6.28.2.2	<code>RTIOC_PURGE</code>	164
6.29	Device Registration Services	165
6.29.1	Detailed Description	166
6.29.2	Macro Definition Documentation	166
6.29.2.1	<code>RTDM_DEVICE_TYPE_MASK</code>	166
6.29.2.2	<code>RTDM_EXCLUSIVE</code>	167
6.29.2.3	<code>RTDM_FIXED_MINOR</code>	167
6.29.2.4	<code>RTDM_MAX_MINOR</code>	167
6.29.2.5	<code>RTDM_NAMED_DEVICE</code>	167
6.29.2.6	<code>RTDM_PROTOCOL_DEVICE</code>	167
6.29.3	Function Documentation	167
6.29.3.1	<code>rtm_close_handler()</code>	167
6.29.3.2	<code>rtm_dev_register()</code>	168
6.29.3.3	<code>rtm_dev_unregister()</code>	169
6.29.3.4	<code>rtm_drv_set_sysclass()</code>	169
6.29.3.5	<code>rtm_get_unmapped_area_handler()</code>	170
6.29.3.6	<code>rtm_ioctl_handler()</code>	171
6.29.3.7	<code>rtm_mmap_handler()</code>	171
6.29.3.8	<code>rtm_open_handler()</code>	172

---

6.29.3.9	<code>rtdm_read_handler()</code>	172
6.29.3.10	<code>rtdm_recvmsg_handler()</code>	173
6.29.3.11	<code>rtdm_select_handler()</code>	173
6.29.3.12	<code>rtdm_sendmsg_handler()</code>	174
6.29.3.13	<code>rtdm_socket_handler()</code>	174
6.29.3.14	<code>rtdm_write_handler()</code>	175
6.30	Clock Services	176
6.30.1	Detailed Description	176
6.30.2	Function Documentation	176
6.30.2.1	<code>rtdm_clock_read()</code>	176
6.30.2.2	<code>rtdm_clock_read_monotonic()</code>	177
6.31	Task Services	178
6.31.1	Detailed Description	179
6.31.2	Typedef Documentation	179
6.31.2.1	<code>rtdm_task_proc_t</code>	179
6.31.3	Function Documentation	180
6.31.3.1	<code>rtdm_task_busy_sleep()</code>	180
6.31.3.2	<code>rtdm_task_busy_wait()</code>	180
6.31.3.3	<code>rtdm_task_current()</code>	181
6.31.3.4	<code>rtdm_task_destroy()</code>	181
6.31.3.5	<code>rtdm_task_init()</code>	182
6.31.3.6	<code>rtdm_task_join()</code>	183
6.31.3.7	<code>rtdm_task_set_period()</code>	183
6.31.3.8	<code>rtdm_task_set_priority()</code>	184
6.31.3.9	<code>rtdm_task_should_stop()</code>	184
6.31.3.10	<code>rtdm_task_sleep()</code>	184
6.31.3.11	<code>rtdm_task_sleep_abs()</code>	185
6.31.3.12	<code>rtdm_task_sleep_until()</code>	186
6.31.3.13	<code>rtdm_task_unblock()</code>	186
6.31.3.14	<code>rtdm_task_wait_period()</code>	186

---

6.31.3.15	<code>rtdm_wait_complete()</code>	187
6.31.3.16	<code>rtdm_wait_is_completed()</code>	187
6.31.3.17	<code>rtdm_wait_prepare()</code>	188
6.32	Timer Services	189
6.32.1	Detailed Description	190
6.32.2	Typedef Documentation	190
6.32.2.1	<code>rtdm_timer_handler_t</code>	190
6.32.3	Enumeration Type Documentation	190
6.32.3.1	<code>rtdm_timer_mode</code>	190
6.32.4	Function Documentation	190
6.32.4.1	<code>rtdm_timer_destroy()</code>	190
6.32.4.2	<code>rtdm_timer_init()</code>	191
6.32.4.3	<code>rtdm_timer_start()</code>	191
6.32.4.4	<code>rtdm_timer_start_in_handler()</code>	192
6.32.4.5	<code>rtdm_timer_stop()</code>	193
6.32.4.6	<code>rtdm_timer_stop_in_handler()</code>	193
6.33	Synchronisation Services	194
6.33.1	Detailed Description	196
6.33.2	Enumeration Type Documentation	196
6.33.2.1	<code>rtdm_selecttype</code>	196
6.33.3	Function Documentation	196
6.33.3.1	<code>rtdm_for_each_waiter()</code>	196
6.33.3.2	<code>rtdm_for_each_waiter_safe()</code>	197
6.33.3.3	<code>rtdm_timedwait()</code>	197
6.33.3.4	<code>rtdm_timedwait_condition()</code>	198
6.33.3.5	<code>rtdm_timedwait_condition_locked()</code>	199
6.33.3.6	<code>rtdm_timedwait_locked()</code>	200
6.33.3.7	<code>rtdm_toseq_init()</code>	201
6.33.3.8	<code>rtdm_wait()</code>	201
6.33.3.9	<code>rtdm_wait_condition()</code>	202

---

6.33.3.10	<code>rtm_wait_condition_locked()</code>	203
6.33.3.11	<code>rtm_wait_locked()</code>	203
6.33.3.12	<code>rtm_waitqueue_broadcast()</code>	204
6.33.3.13	<code>rtm_waitqueue_destroy()</code>	204
6.33.3.14	<code>rtm_waitqueue_flush()</code>	205
6.33.3.15	<code>rtm_waitqueue_init()</code>	205
6.33.3.16	<code>rtm_waitqueue_lock()</code>	206
6.33.3.17	<code>rtm_waitqueue_signal()</code>	206
6.33.3.18	<code>rtm_waitqueue_unlock()</code>	207
6.33.3.19	<code>rtm_waitqueue_wakeup()</code>	207
6.34	Event Services	208
6.34.1	Detailed Description	208
6.34.2	Function Documentation	208
6.34.2.1	<code>rtm_event_clear()</code>	209
6.34.2.2	<code>rtm_event_destroy()</code>	210
6.34.2.3	<code>rtm_event_init()</code>	210
6.34.2.4	<code>rtm_event_pulse()</code>	211
6.34.2.5	<code>rtm_event_select()</code>	211
6.34.2.6	<code>rtm_event_signal()</code>	212
6.34.2.7	<code>rtm_event_timedwait()</code>	212
6.34.2.8	<code>rtm_event_wait()</code>	213
6.35	Semaphore Services	214
6.35.1	Detailed Description	214
6.35.2	Function Documentation	214
6.35.2.1	<code>rtm_sem_destroy()</code>	214
6.35.2.2	<code>rtm_sem_down()</code>	215
6.35.2.3	<code>rtm_sem_init()</code>	215
6.35.2.4	<code>rtm_sem_select()</code>	216
6.35.2.5	<code>rtm_sem_timeddown()</code>	217
6.35.2.6	<code>rtm_sem_up()</code>	217

6.36	Mutex services	219
6.36.1	Detailed Description	219
6.36.2	Function Documentation	219
6.36.2.1	rtdm_mutex_destroy()	219
6.36.2.2	rtdm_mutex_init()	220
6.36.2.3	rtdm_mutex_lock()	220
6.36.2.4	rtdm_mutex_timedlock()	221
6.36.2.5	rtdm_mutex_unlock()	222
6.37	Interrupt Management Services	223
6.37.1	Detailed Description	224
6.37.2	Macro Definition Documentation	224
6.37.2.1	rtdm_irq_get_arg	224
6.37.3	Typedef Documentation	224
6.37.3.1	rtdm_irq_handler_t	224
6.37.4	Function Documentation	225
6.37.4.1	rtdm_irq_disable()	225
6.37.4.2	rtdm_irq_enable()	225
6.37.4.3	rtdm_irq_free()	226
6.37.4.4	rtdm_irq_request()	227
6.38	Non-Real-Time Signalling Services	228
6.38.1	Detailed Description	228
6.38.2	Typedef Documentation	228
6.38.2.1	rtdm_nrtsig_handler_t	228
6.38.3	Function Documentation	229
6.38.3.1	rtdm_nrtsig_destroy()	229
6.38.3.2	rtdm_nrtsig_init()	229
6.38.3.3	rtdm_nrtsig_pend()	230
6.38.3.4	rtdm_schedule_nrt_work()	230
6.39	Utility Services	231
6.39.1	Detailed Description	232

---

6.39.2 Function Documentation . . . . .	232
6.39.2.1 rtdm_available() . . . . .	232
6.39.2.2 rtdm_copy_from_user() . . . . .	233
6.39.2.3 rtdm_copy_to_user() . . . . .	233
6.39.2.4 rtdm_free() . . . . .	234
6.39.2.5 rtdm_in_rt_context() . . . . .	234
6.39.2.6 rtdm_iomap_to_user() . . . . .	235
6.39.2.7 rtdm_malloc() . . . . .	236
6.39.2.8 rtdm_mmap_iomem() . . . . .	236
6.39.2.9 rtdm_mmap_kmem() . . . . .	237
6.39.2.10 rtdm_mmap_to_user() . . . . .	237
6.39.2.11 rtdm_mmap_vmem() . . . . .	238
6.39.2.12 rtdm_munmap() . . . . .	239
6.39.2.13 rtdm_printk() . . . . .	240
6.39.2.14 rtdm_printk_ratelimited() . . . . .	240
6.39.2.15 rtdm_read_user_ok() . . . . .	241
6.39.2.16 rtdm_rt_capable() . . . . .	241
6.39.2.17 rtdm_rw_user_ok() . . . . .	242
6.39.2.18 rtdm_safe_copy_from_user() . . . . .	242
6.39.2.19 rtdm_safe_copy_to_user() . . . . .	243
6.39.2.20 rtdm_strncpy_from_user() . . . . .	244
6.40 SCHED_QUOTA scheduling policy . . . . .	245
6.40.1 Detailed Description . . . . .	245
6.41 Thread scheduling control . . . . .	246
6.41.1 Detailed Description . . . . .	246
6.41.2 Function Documentation . . . . .	246
6.41.2.1 xnsched_rotate() . . . . .	246
6.41.2.2 xnsched_run() . . . . .	247
6.42 Synchronous I/O multiplexing . . . . .	248
6.42.1 Detailed Description . . . . .	248



---

6.42.2 Function Documentation . . . . .	249
6.42.2.1 xselect() . . . . .	249
6.42.2.2 xselect_bind() . . . . .	250
6.42.2.3 xselect_destroy() . . . . .	250
6.42.2.4 xselect_init() . . . . .	251
6.42.2.5 xselect_signal() . . . . .	251
6.42.2.6 xselector_destroy() . . . . .	252
6.42.2.7 xselector_init() . . . . .	252
6.43 Thread synchronization services . . . . .	253
6.43.1 Detailed Description . . . . .	254
6.43.2 Function Documentation . . . . .	254
6.43.2.1 xnsynch_acquire() . . . . .	254
6.43.2.2 xnsynch_destroy() . . . . .	255
6.43.2.3 xnsynch_flush() . . . . .	255
6.43.2.4 xnsynch_init() . . . . .	256
6.43.2.5 xnsynch_init_protect() . . . . .	257
6.43.2.6 xnsynch_peek_pendq() . . . . .	258
6.43.2.7 xnsynch_release() . . . . .	258
6.43.2.8 xnsynch_sleep_on() . . . . .	259
6.43.2.9 xnsynch_try_acquire() . . . . .	260
6.43.2.10xnsynch_wakeup_one_sleeper() . . . . .	260
6.43.2.11xnsynch_wakeup_this_sleeper() . . . . .	261
6.44 Thread services . . . . .	262
6.44.1 Detailed Description . . . . .	263
6.44.2 Function Documentation . . . . .	263
6.44.2.1 xthread_cancel() . . . . .	263
6.44.2.2 xthread_current() . . . . .	264
6.44.2.3 xthread_from_task() . . . . .	265
6.44.2.4 xthread_harden() . . . . .	265
6.44.2.5 xthread_init() . . . . .	265

6.44.2.6	xnthread_join()	267
6.44.2.7	xnthread_map()	267
6.44.2.8	xnthread_relax()	268
6.44.2.9	xnthread_resume()	269
6.44.2.10	xnthread_set_mode()	270
6.44.2.11	xnthread_set_periodic()	271
6.44.2.12	xnthread_set_schedparam()	272
6.44.2.13	xnthread_set_slice()	273
6.44.2.14	xnthread_start()	273
6.44.2.15	xnthread_suspend()	274
6.44.2.16	xnthread_test_cancel()	275
6.44.2.17	xnthread_unblock()	276
6.44.2.18	xnthread_wait_period()	276
6.45	Timer services	278
6.45.1	Detailed Description	279
6.45.2	Function Documentation	279
6.45.2.1	__xntimer_migrate()	279
6.45.2.2	program_htick_shot()	280
6.45.2.3	switch_htick_mode()	280
6.45.2.4	xntimer_destroy()	281
6.45.2.5	xntimer_get_date()	281
6.45.2.6	xntimer_get_overruns()	282
6.45.2.7	xntimer_get_timeout()	282
6.45.2.8	xntimer_grab_hardware()	283
6.45.2.9	xntimer_init()	284
6.45.2.10	xntimer_interval()	284
6.45.2.11	xntimer_release_hardware()	285
6.45.2.12	xntimer_start()	285
6.45.2.13	xntimer_stop()	286
6.46	Virtual file services	287

---

6.46.1 Detailed Description . . . . .	288
6.46.2 Function Documentation . . . . .	289
6.46.2.1 xnvfile_destroy() . . . . .	289
6.46.2.2 xnvfile_get_blob() . . . . .	289
6.46.2.3 xnvfile_get_integer() . . . . .	290
6.46.2.4 xnvfile_get_string() . . . . .	290
6.46.2.5 xnvfile_init_dir() . . . . .	291
6.46.2.6 xnvfile_init_link() . . . . .	292
6.46.2.7 xnvfile_init_regular() . . . . .	292
6.46.2.8 xnvfile_init_snapshot() . . . . .	293
6.46.3 Variable Documentation . . . . .	294
6.46.3.1 cobalt_vfroot [1/2] . . . . .	294
6.46.3.2 cobalt_vfroot [2/2] . . . . .	294
6.47 Analogy framework . . . . .	295
6.47.1 Detailed Description . . . . .	295
6.48 Driver API . . . . .	296
6.48.1 Detailed Description . . . . .	296
6.49 Driver management services . . . . .	297
6.49.1 Detailed Description . . . . .	297
6.49.2 Function Documentation . . . . .	297
6.49.2.1 a4l_register_drv() . . . . .	297
6.49.2.2 a4l_unregister_drv() . . . . .	298
6.50 Subdevice management services . . . . .	299
6.50.1 Detailed Description . . . . .	300
6.50.2 Function Documentation . . . . .	301
6.50.2.1 a4l_add_subd() . . . . .	301
6.50.2.2 a4l_alloc_subd() . . . . .	301
6.50.2.3 a4l_get_subd() . . . . .	302
6.51 Buffer management services . . . . .	303
6.51.1 Detailed Description . . . . .	304

6.51.2 Function Documentation . . . . .	304
6.51.2.1 a4l_buf_commit_absget() . . . . .	304
6.51.2.2 a4l_buf_commit_absput() . . . . .	305
6.51.2.3 a4l_buf_commit_get() . . . . .	305
6.51.2.4 a4l_buf_commit_put() . . . . .	306
6.51.2.5 a4l_buf_count() . . . . .	307
6.51.2.6 a4l_buf_evt() . . . . .	308
6.51.2.7 a4l_buf_get() . . . . .	308
6.51.2.8 a4l_buf_prepare_absget() . . . . .	309
6.51.2.9 a4l_buf_prepare_absput() . . . . .	310
6.51.2.10a4l_buf_prepare_get() . . . . .	310
6.51.2.11a4l_buf_prepare_put() . . . . .	311
6.51.2.12a4l_buf_put() . . . . .	311
6.51.2.13a4l_get_chan() . . . . .	312
6.51.2.14a4l_get_cmd() . . . . .	312
6.52 Interrupt management services . . . . .	313
6.52.1 Detailed Description . . . . .	313
6.52.2 Function Documentation . . . . .	313
6.52.2.1 a4l_free_irq() . . . . .	313
6.52.2.2 a4l_get_irq() . . . . .	314
6.52.2.3 a4l_request_irq() . . . . .	314
6.53 Misc services . . . . .	316
6.53.1 Detailed Description . . . . .	316
6.53.2 Function Documentation . . . . .	316
6.53.2.1 a4l_get_time() . . . . .	316
6.54 Clocks and timers . . . . .	317
6.54.1 Detailed Description . . . . .	318
6.54.2 Function Documentation . . . . .	318
6.54.2.1 clock_getres() . . . . .	318
6.54.2.2 clock_gettime() . . . . .	319

---

6.54.2.3	<code>clock_nanosleep()</code>	320
6.54.2.4	<code>clock_settime()</code>	321
6.54.2.5	<code>nanosleep()</code>	322
6.54.2.6	<code>timer_create()</code>	322
6.54.2.7	<code>timer_delete()</code>	324
6.54.2.8	<code>timer_getoverrun()</code>	324
6.54.2.9	<code>timer_gettime()</code>	325
6.54.2.10	<code>timer_settime()</code>	326
6.55	Condition variables	327
6.55.1	Detailed Description	328
6.55.2	Function Documentation	328
6.55.2.1	<code>pthread_cond_broadcast()</code>	328
6.55.2.2	<code>pthread_cond_destroy()</code>	329
6.55.2.3	<code>pthread_cond_init()</code>	329
6.55.2.4	<code>pthread_cond_signal()</code>	330
6.55.2.5	<code>pthread_cond_timedwait()</code>	331
6.55.2.6	<code>pthread_cond_wait()</code>	332
6.55.2.7	<code>pthread_condattr_destroy()</code>	333
6.55.2.8	<code>pthread_condattr_getclock()</code>	333
6.55.2.9	<code>pthread_condattr_getpshared()</code>	334
6.55.2.10	<code>pthread_condattr_init()</code>	335
6.55.2.11	<code>pthread_condattr_setclock()</code>	336
6.55.2.12	<code>pthread_condattr_setpshared()</code>	336
6.56	POSIX interface	338
6.56.1	Detailed Description	339
6.57	Message queues	340
6.57.1	Detailed Description	341
6.57.2	Function Documentation	341
6.57.2.1	<code>mq_close()</code>	341
6.57.2.2	<code>mq_getattr()</code>	342

---

6.57.2.3	<a href="#">mq_notify()</a>	342
6.57.2.4	<a href="#">mq_open()</a>	343
6.57.2.5	<a href="#">mq_receive()</a>	345
6.57.2.6	<a href="#">mq_send()</a>	346
6.57.2.7	<a href="#">mq_setattr()</a>	347
6.57.2.8	<a href="#">mq_timedreceive()</a>	347
6.57.2.9	<a href="#">mq_timedsend()</a>	348
6.57.2.10	<a href="#">mq_unlink()</a>	349
6.58	<a href="#">Mutual exclusion</a>	351
6.58.1	<a href="#">Detailed Description</a>	352
6.58.2	<a href="#">Function Documentation</a>	352
6.58.2.1	<a href="#">__attribute__()</a>	353
6.58.2.2	<a href="#">pthread_mutex_destroy()</a>	353
6.58.2.3	<a href="#">pthread_mutex_getprioceiling()</a>	354
6.58.2.4	<a href="#">pthread_mutex_init()</a>	354
6.58.2.5	<a href="#">pthread_mutex_lock()</a>	355
6.58.2.6	<a href="#">pthread_mutex_setprioceiling()</a>	356
6.58.2.7	<a href="#">pthread_mutex_timedlock()</a>	357
6.58.2.8	<a href="#">pthread_mutex_trylock()</a>	358
6.58.2.9	<a href="#">pthread_mutex_unlock()</a>	359
6.58.2.10	<a href="#">pthread_mutexattr_destroy()</a>	359
6.58.2.11	<a href="#">pthread_mutexattr_getprotocol()</a>	360
6.58.2.12	<a href="#">pthread_mutexattr_getpshared()</a>	361
6.58.2.13	<a href="#">pthread_mutexattr_gettype()</a>	361
6.58.2.14	<a href="#">pthread_mutexattr_init()</a>	362
6.58.2.15	<a href="#">pthread_mutexattr_setprotocol()</a>	363
6.58.2.16	<a href="#">pthread_mutexattr_setpshared()</a>	364
6.58.2.17	<a href="#">pthread_mutexattr_settype()</a>	365
6.59	<a href="#">Process scheduling</a>	366
6.59.1	<a href="#">Detailed Description</a>	367

---

---

6.59.2 Function Documentation . . . . .	367
6.59.2.1 sched_get_priority_max() . . . . .	367
6.59.2.2 sched_get_priority_max_ex() . . . . .	368
6.59.2.3 sched_get_priority_min() . . . . .	368
6.59.2.4 sched_get_priority_min_ex() . . . . .	369
6.59.2.5 sched_getconfig_np() . . . . .	370
6.59.2.6 sched_getscheduler() . . . . .	371
6.59.2.7 sched_getscheduler_ex() . . . . .	372
6.59.2.8 sched_setconfig_np() . . . . .	372
6.59.2.9 sched_setscheduler() . . . . .	374
6.59.2.10 sched_setscheduler_ex() . . . . .	375
6.59.2.11 sched_yield() . . . . .	376
6.60 Semaphores . . . . .	378
6.60.1 Detailed Description . . . . .	378
6.60.2 Function Documentation . . . . .	379
6.60.2.1 sem_close() . . . . .	379
6.60.2.2 sem_destroy() . . . . .	379
6.60.2.3 sem_init() . . . . .	380
6.60.2.4 sem_post() . . . . .	381
6.60.2.5 sem_timedwait() . . . . .	382
6.60.2.6 sem_trywait() . . . . .	383
6.60.2.7 sem_unlink() . . . . .	384
6.60.2.8 sem_wait() . . . . .	385
6.61 Thread management . . . . .	386
6.61.1 Detailed Description . . . . .	386
6.61.2 Function Documentation . . . . .	386
6.61.2.1 pthread_create() . . . . .	387
6.61.2.2 pthread_join() . . . . .	388
6.61.2.3 pthread_kill() . . . . .	389
6.61.2.4 pthread_setmode_np() . . . . .	389

---

6.61.2.5 pthread_setname_np()	391
6.62 Scheduling management	392
6.62.1 Detailed Description	392
6.62.2 Function Documentation	392
6.62.2.1 pthread_getschedparam()	393
6.62.2.2 pthread_getschedparam_ex()	393
6.62.2.3 pthread_setschedparam()	394
6.62.2.4 pthread_setschedparam_ex()	395
6.62.2.5 pthread_yield()	396
6.63 Smokey API	398
6.64 Asynchronous acquisition API	402
6.64.1 Detailed Description	404
6.64.2 Function Documentation	404
6.64.2.1 a4l_get_bufsize()	404
6.64.2.2 a4l_mark_bufrw()	405
6.64.2.3 a4l_mmap()	406
6.64.2.4 a4l_poll()	406
6.64.2.5 a4l_set_bufsize()	407
6.64.2.6 a4l_snd_cancel()	408
6.64.2.7 a4l_snd_command()	408
6.65 Asynchronous acquisition API	410
6.65.1 Detailed Description	410
6.65.2 Function Documentation	410
6.65.2.1 a4l_async_read()	410
6.65.2.2 a4l_async_write()	411
6.66 Software calibration API	412
6.66.1 Detailed Description	412
6.66.2 Function Documentation	412
6.66.2.1 a4l_dcaltoraw()	412
6.66.2.2 a4l_get_softcal_converter()	413



---

6.66.2.3	a4l_rawtodcal()	413
6.66.2.4	a4l_read_calibration_file()	414
6.67	Descriptor Syscall API	415
6.67.1	Detailed Description	415
6.67.2	Function Documentation	415
6.67.2.1	a4l_sys_desc()	416
6.68	Descriptor API	417
6.68.1	Detailed Description	417
6.68.2	Function Documentation	417
6.68.2.1	a4l_close()	417
6.68.2.2	a4l_fill_desc()	418
6.68.2.3	a4l_get_chinfo()	418
6.68.2.4	a4l_get_rnginfo()	419
6.68.2.5	a4l_get_subdinfo()	420
6.68.2.6	a4l_open()	420
6.69	Math API	421
6.69.1	Detailed Description	421
6.69.2	Function Documentation	421
6.69.2.1	a4l_math_mean()	421
6.69.2.2	a4l_math_polyfit()	422
6.69.2.3	a4l_math_stddev()	422
6.69.2.4	a4l_math_stddev_of_mean()	423
6.70	Range / conversion API	424
6.70.1	Detailed Description	424
6.70.2	Function Documentation	424
6.70.2.1	a4l_dtoraw()	424
6.70.2.2	a4l_find_range()	425
6.70.2.3	a4l_ftoraw()	426
6.70.2.4	a4l_rawtod()	426
6.70.2.5	a4l_rawtof()	427

---

6.70.2.6	a4l_rawtoul()	427
6.70.2.7	a4l_sizeof_chan()	428
6.70.2.8	a4l_sizeof_subd()	428
6.70.2.9	a4l_ultoraw()	429
6.71	Level 1 API	430
6.71.1	Detailed Description	430
6.72	Synchronous acquisition API	431
6.72.1	Detailed Description	433
6.72.2	Function Documentation	433
6.72.2.1	a4l_snd_insn()	433
6.72.2.2	a4l_snd_insnlist()	434
6.73	Level 2 API	435
6.73.1	Detailed Description	435
6.74	Synchronous acquisition API	436
6.74.1	Detailed Description	436
6.74.2	Function Documentation	436
6.74.2.1	a4l_config_subd()	436
6.74.2.2	a4l_sync_dio()	437
6.74.2.3	a4l_sync_read()	438
6.74.2.4	a4l_sync_write()	438
6.75	Analogy user API	440
6.75.1	Detailed Description	440
6.76	Level 0 API	441
6.76.1	Detailed Description	441
6.77	Basic Syscall API	442
6.77.1	Detailed Description	442
6.77.2	Function Documentation	442
6.77.2.1	a4l_sys_close()	442
6.77.2.2	a4l_sys_open()	443
6.77.2.3	a4l_sys_read()	443

---

6.77.2.4	<code>a4l_sys_write()</code> . . . . .	444
6.78	Attach / detach Syscall API . . . . .	445
6.78.1	Detailed Description . . . . .	445
6.78.2	Function Documentation . . . . .	445
6.78.2.1	<code>a4l_sys_attach()</code> . . . . .	445
6.78.2.2	<code>a4l_sys_bufcfg()</code> . . . . .	446
6.78.2.3	<code>a4l_sys_detach()</code> . . . . .	446
6.79	Alarm services . . . . .	448
6.79.1	Detailed Description . . . . .	448
6.79.2	Function Documentation . . . . .	449
6.79.2.1	<code>rt_alarm_create()</code> . . . . .	449
6.79.2.2	<code>rt_alarm_delete()</code> . . . . .	450
6.79.2.3	<code>rt_alarm_inquire()</code> . . . . .	450
6.79.2.4	<code>rt_alarm_start()</code> . . . . .	451
6.79.2.5	<code>rt_alarm_stop()</code> . . . . .	452
6.80	Buffer services . . . . .	453
6.80.1	Detailed Description . . . . .	454
6.80.2	Macro Definition Documentation . . . . .	454
6.80.2.1	<code>B_PRIO</code> . . . . .	454
6.80.3	Function Documentation . . . . .	454
6.80.3.1	<code>rt_buffer_bind()</code> . . . . .	455
6.80.3.2	<code>rt_buffer_clear()</code> . . . . .	455
6.80.3.3	<code>rt_buffer_create()</code> . . . . .	456
6.80.3.4	<code>rt_buffer_delete()</code> . . . . .	457
6.80.3.5	<code>rt_buffer_inquire()</code> . . . . .	458
6.80.3.6	<code>rt_buffer_read()</code> . . . . .	458
6.80.3.7	<code>rt_buffer_read_timed()</code> . . . . .	459
6.80.3.8	<code>rt_buffer_read_until()</code> . . . . .	460
6.80.3.9	<code>rt_buffer_unbind()</code> . . . . .	461
6.80.3.10	<code>rt_buffer_write()</code> . . . . .	461

---

6.80.3.11	rt_buffer_write_timed()	462
6.80.3.12	rt_buffer_write_until()	463
6.81	Condition variable services	464
6.81.1	Detailed Description	465
6.81.2	Function Documentation	465
6.81.2.1	rt_cond_bind()	465
6.81.2.2	rt_cond_broadcast()	466
6.81.2.3	rt_cond_create()	466
6.81.2.4	rt_cond_delete()	467
6.81.2.5	rt_cond_inquire()	468
6.81.2.6	rt_cond_signal()	469
6.81.2.7	rt_cond_unbind()	469
6.81.2.8	rt_cond_wait()	470
6.81.2.9	rt_cond_wait_timed()	470
6.81.2.10	rt_cond_wait_until()	471
6.82	Event flag group services	472
6.82.1	Detailed Description	473
6.82.2	Macro Definition Documentation	473
6.82.2.1	EV_ANY	473
6.82.2.2	EV_PRIO	473
6.82.3	Function Documentation	474
6.82.3.1	rt_event_bind()	474
6.82.3.2	rt_event_clear()	475
6.82.3.3	rt_event_create()	475
6.82.3.4	rt_event_delete()	476
6.82.3.5	rt_event_inquire()	477
6.82.3.6	rt_event_signal()	477
6.82.3.7	rt_event_unbind()	478
6.82.3.8	rt_event_wait()	478
6.82.3.9	rt_event_wait_timed()	479

---

6.82.3.10	<code>rt_event_wait_until()</code> . . . . .	480
6.83	Heap management services . . . . .	482
6.83.1	Detailed Description . . . . .	483
6.83.2	Macro Definition Documentation . . . . .	483
6.83.2.1	<code>H_PRIO</code> . . . . .	483
6.83.3	Function Documentation . . . . .	483
6.83.3.1	<code>rt_heap_alloc()</code> . . . . .	483
6.83.3.2	<code>rt_heap_alloc_timed()</code> . . . . .	484
6.83.3.3	<code>rt_heap_alloc_until()</code> . . . . .	486
6.83.3.4	<code>rt_heap_bind()</code> . . . . .	486
6.83.3.5	<code>rt_heap_create()</code> . . . . .	487
6.83.3.6	<code>rt_heap_delete()</code> . . . . .	488
6.83.3.7	<code>rt_heap_free()</code> . . . . .	489
6.83.3.8	<code>rt_heap_inquire()</code> . . . . .	489
6.83.3.9	<code>rt_heap_unbind()</code> . . . . .	490
6.84	Alchemy API . . . . .	491
6.84.1	Detailed Description . . . . .	492
6.85	Mutex services . . . . .	493
6.85.1	Detailed Description . . . . .	494
6.85.2	Function Documentation . . . . .	494
6.85.2.1	<code>rt_mutex_acquire()</code> . . . . .	494
6.85.2.2	<code>rt_mutex_acquire_timed()</code> . . . . .	494
6.85.2.3	<code>rt_mutex_acquire_until()</code> . . . . .	495
6.85.2.4	<code>rt_mutex_bind()</code> . . . . .	496
6.85.2.5	<code>rt_mutex_create()</code> . . . . .	497
6.85.2.6	<code>rt_mutex_delete()</code> . . . . .	498
6.85.2.7	<code>rt_mutex_inquire()</code> . . . . .	498
6.85.2.8	<code>rt_mutex_release()</code> . . . . .	499
6.85.2.9	<code>rt_mutex_unbind()</code> . . . . .	499
6.86	Message pipe services . . . . .	501

---

6.86.1 Detailed Description . . . . .	502
6.86.2 Macro Definition Documentation . . . . .	502
6.86.2.1 P_MINOR_AUTO . . . . .	502
6.86.2.2 P_URGENT . . . . .	502
6.86.3 Function Documentation . . . . .	502
6.86.3.1 rt_pipe_bind() . . . . .	502
6.86.3.2 rt_pipe_create() . . . . .	503
6.86.3.3 rt_pipe_delete() . . . . .	504
6.86.3.4 rt_pipe_read() . . . . .	505
6.86.3.5 rt_pipe_read_timed() . . . . .	506
6.86.3.6 rt_pipe_read_until() . . . . .	507
6.86.3.7 rt_pipe_stream() . . . . .	507
6.86.3.8 rt_pipe_unbind() . . . . .	508
6.86.3.9 rt_pipe_write() . . . . .	508
6.87 Message queue services . . . . .	510
6.87.1 Detailed Description . . . . .	511
6.87.2 Macro Definition Documentation . . . . .	511
6.87.2.1 Q_PRIO . . . . .	511
6.87.3 Function Documentation . . . . .	511
6.87.3.1 rt_queue_alloc() . . . . .	512
6.87.3.2 rt_queue_bind() . . . . .	512
6.87.3.3 rt_queue_create() . . . . .	513
6.87.3.4 rt_queue_delete() . . . . .	514
6.87.3.5 rt_queue_flush() . . . . .	515
6.87.3.6 rt_queue_free() . . . . .	515
6.87.3.7 rt_queue_inquire() . . . . .	516
6.87.3.8 rt_queue_read() . . . . .	517
6.87.3.9 rt_queue_read_timed() . . . . .	517
6.87.3.10rt_queue_read_until() . . . . .	518
6.87.3.11rt_queue_receive() . . . . .	519

---

6.87.3.12	<code>rt_queue_receive_timed()</code> . . . . .	519
6.87.3.13	<code>rt_queue_receive_until()</code> . . . . .	521
6.87.3.14	<code>rt_queue_send()</code> . . . . .	521
6.87.3.15	<code>rt_queue_unbind()</code> . . . . .	522
6.88	Semaphore services . . . . .	523
6.88.1	Detailed Description . . . . .	524
6.88.2	Macro Definition Documentation . . . . .	524
6.88.2.1	<code>S_PRIO</code> . . . . .	524
6.88.3	Function Documentation . . . . .	524
6.88.3.1	<code>rt_sem_bind()</code> . . . . .	524
6.88.3.2	<code>rt_sem_broadcast()</code> . . . . .	525
6.88.3.3	<code>rt_sem_create()</code> . . . . .	526
6.88.3.4	<code>rt_sem_delete()</code> . . . . .	527
6.88.3.5	<code>rt_sem_inquire()</code> . . . . .	527
6.88.3.6	<code>rt_sem_p()</code> . . . . .	528
6.88.3.7	<code>rt_sem_p_timed()</code> . . . . .	528
6.88.3.8	<code>rt_sem_p_until()</code> . . . . .	529
6.88.3.9	<code>rt_sem_unbind()</code> . . . . .	530
6.88.3.10	<code>rt_sem_v()</code> . . . . .	530
6.89	Task management services . . . . .	532
6.89.1	Detailed Description . . . . .	534
6.89.2	Macro Definition Documentation . . . . .	534
6.89.2.1	<code>T_LOCK</code> . . . . .	534
6.89.2.2	<code>T_LOPRIO</code> . . . . .	534
6.89.2.3	<code>T_WARNSW</code> . . . . .	535
6.89.3	Function Documentation . . . . .	535
6.89.3.1	<code>rt_task_bind()</code> . . . . .	535
6.89.3.2	<code>rt_task_create()</code> . . . . .	536
6.89.3.3	<code>rt_task_delete()</code> . . . . .	537
6.89.3.4	<code>rt_task_inquire()</code> . . . . .	538

---

6.89.3.5	<code>rt_task_join()</code>	539
6.89.3.6	<code>rt_task_receive()</code>	540
6.89.3.7	<code>rt_task_receive_timed()</code>	540
6.89.3.8	<code>rt_task_receive_until()</code>	542
6.89.3.9	<code>rt_task_reply()</code>	542
6.89.3.10	<code>rt_task_resume()</code>	543
6.89.3.11	<code>rt_task_same()</code>	544
6.89.3.12	<code>rt_task_self()</code>	544
6.89.3.13	<code>rt_task_send()</code>	545
6.89.3.14	<code>rt_task_send_timed()</code>	545
6.89.3.15	<code>rt_task_send_until()</code>	547
6.89.3.16	<code>rt_task_set_affinity()</code>	548
6.89.3.17	<code>rt_task_set_mode()</code>	548
6.89.3.18	<code>rt_task_set_periodic()</code>	550
6.89.3.19	<code>rt_task_set_priority()</code>	551
6.89.3.20	<code>rt_task_shadow()</code>	551
6.89.3.21	<code>rt_task_sleep()</code>	553
6.89.3.22	<code>rt_task_sleep_until()</code>	554
6.89.3.23	<code>rt_task_slice()</code>	555
6.89.3.24	<code>rt_task_spawn()</code>	556
6.89.3.25	<code>rt_task_start()</code>	556
6.89.3.26	<code>rt_task_suspend()</code>	557
6.89.3.27	<code>rt_task_unbind()</code>	558
6.89.3.28	<code>rt_task_unblock()</code>	558
6.89.3.29	<code>rt_task_wait_period()</code>	559
6.89.3.30	<code>rt_task_yield()</code>	560
6.90	Timer management services	561
6.90.1	Detailed Description	561
6.90.2	Typedef Documentation	562
6.90.2.1	<code>RT_TIMER_INFO</code>	562

---



---

6.90.3 Function Documentation . . . . .	562
6.90.3.1 rt_timer_inquire() . . . . .	562
6.90.3.2 rt_timer_ns2ticks() . . . . .	562
6.90.3.3 rt_timer_read() . . . . .	563
6.90.3.4 rt_timer_spin() . . . . .	563
6.90.3.5 rt_timer_ticks2ns() . . . . .	564
6.91 VxWorks® emulator . . . . .	565
6.92 pSOS® emulator . . . . .	566
6.93 Transition Kit . . . . .	567
6.93.1 Detailed Description . . . . .	567
6.93.2 Function Documentation . . . . .	567
6.93.2.1 COMPAT__rt_alarm_create() . . . . .	568
6.93.2.2 COMPAT__rt_event_clear() . . . . .	568
6.93.2.3 COMPAT__rt_event_create() . . . . .	569
6.93.2.4 COMPAT__rt_event_signal() . . . . .	570
6.93.2.5 COMPAT__rt_pipe_create() . . . . .	571
6.93.2.6 COMPAT__rt_task_create() . . . . .	572
6.93.2.7 COMPAT__rt_task_set_periodic() . . . . .	574
6.93.2.8 pthread_make_periodic_np() . . . . .	575
6.93.2.9 pthread_wait_np() . . . . .	576
6.93.2.10rt_alarm_wait() . . . . .	577

<b>7 Data Structure Documentation</b>	<b>579</b>
7.1 a4l_channel Struct Reference	579
7.1.1 Detailed Description	579
7.1.2 Field Documentation	579
7.1.2.1 flags	579
7.1.2.2 nb_bits	579
7.2 a4l_channels_desc Struct Reference	580
7.2.1 Detailed Description	580
7.2.2 Field Documentation	580
7.2.2.1 chans	580
7.2.2.2 length	581
7.2.2.3 mode	581
7.3 a4l_cmd_desc Struct Reference	581
7.3.1 Detailed Description	582
7.3.2 Field Documentation	582
7.3.2.1 data_len	582
7.3.2.2 idx_subd	582
7.4 a4l_descriptor Struct Reference	582
7.4.1 Detailed Description	583
7.4.2 Field Documentation	583
7.4.2.1 board_name	583
7.4.2.2 driver_name	583
7.4.2.3 fd	583
7.4.2.4 idx_read_subd	583
7.4.2.5 idx_write_subd	584
7.4.2.6 magic	584
7.4.2.7 nb_subd	584
7.4.2.8 sbdata	584
7.4.2.9 sbsize	584
7.5 a4l_driver Struct Reference	584

---

7.5.1 Detailed Description . . . . .	585
7.6 a4l_instruction Struct Reference . . . . .	585
7.6.1 Detailed Description . . . . .	586
7.6.2 Field Documentation . . . . .	586
7.6.2.1 idx_subd . . . . .	586
7.7 a4l_instruction_list Struct Reference . . . . .	586
7.7.1 Detailed Description . . . . .	587
7.8 a4l_range Struct Reference . . . . .	587
7.8.1 Detailed Description . . . . .	587
7.8.2 Field Documentation . . . . .	587
7.8.2.1 flags . . . . .	587
7.8.2.2 max . . . . .	587
7.8.2.3 min . . . . .	588
7.9 a4l_subdevice Struct Reference . . . . .	588
7.9.1 Detailed Description . . . . .	589
7.10 atomic_t Struct Reference . . . . .	589
7.10.1 Detailed Description . . . . .	590
7.11 can_bittime Struct Reference . . . . .	590
7.11.1 Detailed Description . . . . .	591
7.12 can_bittime_btr Struct Reference . . . . .	591
7.12.1 Detailed Description . . . . .	591
7.13 can_bittime_std Struct Reference . . . . .	591
7.13.1 Detailed Description . . . . .	592
7.14 can_filter Struct Reference . . . . .	592
7.14.1 Detailed Description . . . . .	592
7.14.2 Field Documentation . . . . .	592
7.14.2.1 can_id . . . . .	593
7.14.2.2 can_mask . . . . .	593
7.15 can_frame Struct Reference . . . . .	593
7.15.1 Detailed Description . . . . .	594

---

7.15.2 Field Documentation . . . . .	594
7.15.2.1 can_id . . . . .	594
7.16 can_ifreq Struct Reference . . . . .	594
7.16.1 Detailed Description . . . . .	595
7.17 macb_dma_desc Struct Reference . . . . .	595
7.17.1 Detailed Description . . . . .	595
7.18 macb_tx_skb Struct Reference . . . . .	595
7.18.1 Detailed Description . . . . .	596
7.19 RT_ALARM_INFO Struct Reference . . . . .	596
7.19.1 Detailed Description . . . . .	596
7.20 RT_BUFFER_INFO Struct Reference . . . . .	596
7.20.1 Detailed Description . . . . .	597
7.21 RT_COND_INFO Struct Reference . . . . .	597
7.21.1 Detailed Description . . . . .	597
7.22 RT_EVENT_INFO Struct Reference . . . . .	597
7.22.1 Detailed Description . . . . .	598
7.23 RT_HEAP_INFO Struct Reference . . . . .	598
7.23.1 Detailed Description . . . . .	598
7.23.2 Field Documentation . . . . .	598
7.23.2.1 heapsize . . . . .	598
7.23.2.2 usablemem . . . . .	599
7.23.2.3 usedmem . . . . .	599
7.24 RT_MUTEX_INFO Struct Reference . . . . .	599
7.24.1 Detailed Description . . . . .	599
7.24.2 Field Documentation . . . . .	599
7.24.2.1 owner . . . . .	600
7.25 RT_QUEUE_INFO Struct Reference . . . . .	600
7.25.1 Detailed Description . . . . .	600
7.26 RT_SEM_INFO Struct Reference . . . . .	600
7.26.1 Detailed Description . . . . .	601

---

7.27 RT_TASK_INFO Struct Reference . . . . .	601
7.27.1 Detailed Description . . . . .	601
7.28 rt_timer_info Struct Reference . . . . .	602
7.28.1 Detailed Description . . . . .	602
7.28.2 Field Documentation . . . . .	602
7.28.2.1 date . . . . .	602
7.29 rtdm_dev_context Struct Reference . . . . .	603
7.29.1 Detailed Description . . . . .	603
7.29.2 Field Documentation . . . . .	603
7.29.2.1 device . . . . .	604
7.30 rtdm_device Struct Reference . . . . .	604
7.30.1 Detailed Description . . . . .	605
7.30.2 Field Documentation . . . . .	605
7.30.2.1 "@15 . . . . .	605
7.30.2.2 driver . . . . .	605
7.30.2.3 label . . . . .	605
7.30.2.4 minor . . . . .	605
7.31 rtdm_device_info Struct Reference . . . . .	606
7.31.1 Detailed Description . . . . .	606
7.32 rtdm_driver Struct Reference . . . . .	606
7.32.1 Detailed Description . . . . .	607
7.32.2 Field Documentation . . . . .	607
7.32.2.1 base_minor . . . . .	607
7.32.2.2 context_size . . . . .	608
7.32.2.3 device_count . . . . .	608
7.32.2.4 device_flags . . . . .	608
7.32.2.5 profile_info . . . . .	608
7.33 rtdm_fd_ops Struct Reference . . . . .	608
7.33.1 Detailed Description . . . . .	609
7.33.2 Field Documentation . . . . .	609

7.33.2.1	close	610
7.33.2.2	get_unmapped_area	610
7.33.2.3	ioctl_nrt	610
7.33.2.4	ioctl_rt	610
7.33.2.5	mmap	610
7.33.2.6	open	610
7.33.2.7	read_nrt	611
7.33.2.8	read_rt	611
7.33.2.9	recvmsg_nrt	611
7.33.2.10	recvmsg_rt	611
7.33.2.11	select	611
7.33.2.12	sendmsg_nrt	611
7.33.2.13	sendmsg_rt	612
7.33.2.14	socket	612
7.33.2.15	write_nrt	612
7.33.2.16	write_rt	612
7.34	rtm_gpio_readout Struct Reference	612
7.34.1	Detailed Description	612
7.35	rtm_profile_info Struct Reference	613
7.35.1	Detailed Description	613
7.36	rtm_sm_ops Struct Reference	613
7.36.1	Detailed Description	613
7.37	rtm_spi_config Struct Reference	614
7.37.1	Detailed Description	614
7.38	rtipc_port_label Struct Reference	614
7.38.1	Detailed Description	614
7.38.2	Field Documentation	614
7.38.2.1	label	615
7.39	rtser_config Struct Reference	615
7.39.1	Detailed Description	616

7.40 rtser_event Struct Reference . . . . .	616
7.40.1 Detailed Description . . . . .	616
7.41 rtser_status Struct Reference . . . . .	616
7.41.1 Detailed Description . . . . .	617
7.42 sockaddr_can Struct Reference . . . . .	617
7.42.1 Detailed Description . . . . .	617
7.42.2 Field Documentation . . . . .	617
7.42.2.1 can_ifindex . . . . .	617
7.43 sockaddr_ipc Struct Reference . . . . .	618
7.43.1 Detailed Description . . . . .	618
7.43.2 Field Documentation . . . . .	618
7.43.2.1 sipc_port . . . . .	618
7.44 udd_device Struct Reference . . . . .	619
7.44.1 Detailed Description . . . . .	620
7.44.2 Field Documentation . . . . .	620
7.44.2.1 close . . . . .	620
7.44.2.2 device_flags . . . . .	620
7.44.2.3 device_subclass . . . . .	621
7.44.2.4 interrupt . . . . .	621
7.44.2.5 ioctl . . . . .	621
7.44.2.6 irq . . . . .	622
7.44.2.7 mem_regions . . . . .	622
7.44.2.8 mmap . . . . .	622
7.44.2.9 open . . . . .	622
7.45 udd_memregion Struct Reference . . . . .	623
7.45.1 Detailed Description . . . . .	623
7.45.2 Field Documentation . . . . .	624
7.45.2.1 addr . . . . .	624
7.45.2.2 len . . . . .	624
7.45.2.3 type . . . . .	624

7.46 udd_device::udd_reserved Struct Reference . . . . .	624
7.46.1 Detailed Description . . . . .	625
7.47 udd_signotify Struct Reference . . . . .	625
7.47.1 Detailed Description . . . . .	625
7.47.2 Field Documentation . . . . .	625
7.47.2.1 pid . . . . .	625
7.47.2.2 sig . . . . .	626
7.48 xnsched Struct Reference . . . . .	626
7.48.1 Detailed Description . . . . .	626
7.48.2 Field Documentation . . . . .	626
7.48.2.1 cpu . . . . .	626
7.48.2.2 curr . . . . .	627
7.48.2.3 htimer . . . . .	627
7.48.2.4 inesting . . . . .	627
7.48.2.5 lflags . . . . .	627
7.48.2.6 resched . . . . .	627
7.48.2.7 rrbtimer . . . . .	627
7.48.2.8 rt . . . . .	627
7.48.2.9 status . . . . .	628
7.49 xnfile_lock_ops Struct Reference . . . . .	628
7.49.1 Detailed Description . . . . .	628
7.49.2 Field Documentation . . . . .	628
7.49.2.1 get . . . . .	628
7.49.2.2 put . . . . .	629
7.50 xnfile_regular_iterator Struct Reference . . . . .	629
7.50.1 Detailed Description . . . . .	629
7.50.2 Field Documentation . . . . .	629
7.50.2.1 pos . . . . .	630
7.50.2.2 private . . . . .	630
7.50.2.3 seq . . . . .	630



---

7.50.2.4 vfile . . . . .	630
7.51 xnvfile_regular_ops Struct Reference . . . . .	630
7.51.1 Detailed Description . . . . .	631
7.51.2 Field Documentation . . . . .	631
7.51.2.1 begin . . . . .	631
7.51.2.2 end . . . . .	631
7.51.2.3 next . . . . .	632
7.51.2.4 rewind . . . . .	632
7.51.2.5 show . . . . .	633
7.51.2.6 store . . . . .	633
7.52 xnvfile_rev_tag Struct Reference . . . . .	634
7.52.1 Detailed Description . . . . .	634
7.52.2 Field Documentation . . . . .	634
7.52.2.1 rev . . . . .	634
7.53 xnvfile_snapshot Struct Reference . . . . .	635
7.53.1 Detailed Description . . . . .	635
7.54 xnvfile_snapshot_iterator Struct Reference . . . . .	635
7.54.1 Detailed Description . . . . .	636
7.54.2 Field Documentation . . . . .	636
7.54.2.1 databuf . . . . .	637
7.54.2.2 endfn . . . . .	637
7.54.2.3 nrdata . . . . .	637
7.54.2.4 private . . . . .	637
7.54.2.5 seq . . . . .	637
7.54.2.6 vfile . . . . .	637
7.55 xnvfile_snapshot_ops Struct Reference . . . . .	638
7.55.1 Detailed Description . . . . .	638
7.55.2 Field Documentation . . . . .	638
7.55.2.1 begin . . . . .	638
7.55.2.2 end . . . . .	639
7.55.2.3 next . . . . .	639
7.55.2.4 rewind . . . . .	640
7.55.2.5 show . . . . .	640
7.55.2.6 store . . . . .	641

<b>8 File Documentation</b>	<b>643</b>
8.1 include/cobalt/kernel/rtdm/analogy/channel_range.h File Reference	643
8.1.1 Detailed Description	645
8.2 include/cobalt/kernel/rtdm/analogy/driver.h File Reference	646
8.2.1 Detailed Description	646
8.3 include/cobalt/kernel/rtdm/driver.h File Reference	647
8.3.1 Detailed Description	652
8.3.2 Macro Definition Documentation	652
8.3.2.1 RTDM_PROFILE_INFO	652
8.3.3 Function Documentation	653
8.3.3.1 rtdm_fd_device()	653
8.3.3.2 rtdm_fd_is_user()	653
8.3.3.3 rtdm_fd_to_private()	654
8.3.3.4 rtdm_private_to_fd()	654
8.4 include/cobalt/kernel/rtdm/analogy/subdevice.h File Reference	654
8.4.1 Detailed Description	656
8.5 include/rtdm/can.h File Reference	656
8.5.1 Detailed Description	657
8.6 include/rtdm/uapi/can.h File Reference	657
8.6.1 Detailed Description	663
8.7 include/cobalt/kernel/rtdm/fd.h File Reference	664
8.7.1 Detailed Description	665
8.7.2 Function Documentation	665
8.7.2.1 rtdm_fd_get()	665
8.7.2.2 rtdm_fd_lock()	666
8.7.2.3 rtdm_fd_put()	666
8.7.2.4 rtdm_fd_select()	667
8.7.2.5 rtdm_fd_unlock()	668
8.8 include/rtdm/uapi/gpiopwm.h File Reference	669
8.8.1 Detailed Description	670

8.9	<a href="#">include/rtdm/uapi/ipc.h File Reference</a>	670
8.9.1	<a href="#">Detailed Description</a>	673
8.10	<a href="#">include/rtdm/rtdm.h File Reference</a>	674
8.10.1	<a href="#">Detailed Description</a>	674
8.11	<a href="#">include/rtdm/uapi/rtdm.h File Reference</a>	674
8.11.1	<a href="#">Detailed Description</a>	676
8.12	<a href="#">include/rtdm/serial.h File Reference</a>	676
8.12.1	<a href="#">Detailed Description</a>	677
8.13	<a href="#">include/rtdm/uapi/serial.h File Reference</a>	677
8.13.1	<a href="#">Detailed Description</a>	682
8.13.2	<a href="#">Macro Definition Documentation</a>	682
8.13.2.1	<a href="#">RTSER_RTIOC_BREAK_CTL</a>	682
8.13.2.2	<a href="#">RTSER_RTIOC_GET_CONFIG</a>	683
8.13.2.3	<a href="#">RTSER_RTIOC_GET_CONTROL</a>	683
8.13.2.4	<a href="#">RTSER_RTIOC_GET_STATUS</a>	683
8.13.2.5	<a href="#">RTSER_RTIOC_SET_CONFIG</a>	684
8.13.2.6	<a href="#">RTSER_RTIOC_SET_CONTROL</a>	685
8.13.2.7	<a href="#">RTSER_RTIOC_WAIT_EVENT</a>	685
8.14	<a href="#">include/rtdm/testing.h File Reference</a>	686
8.14.1	<a href="#">Detailed Description</a>	686
8.15	<a href="#">include/rtdm/uapi/testing.h File Reference</a>	687
8.15.1	<a href="#">Detailed Description</a>	688
8.16	<a href="#">include/cobalt/kernel/rtdm/udd.h File Reference</a>	689
8.16.1	<a href="#">Detailed Description</a>	690
8.17	<a href="#">include/rtdm/uapi/udd.h File Reference</a>	690
8.17.1	<a href="#">Detailed Description</a>	691
8.18	<a href="#">include/rtdm/analogy.h File Reference</a>	691
8.18.1	<a href="#">Detailed Description</a>	692
8.19	<a href="#">include/rtdm/uapi/analogy.h File Reference</a>	693
8.19.1	<a href="#">Detailed Description</a>	697

8.19.2 Macro Definition Documentation . . . . .	697
8.19.2.1 A4L_RNG_FACTOR . . . . .	698
8.20 lib/analogy/calibration.h File Reference . . . . .	698
8.20.1 Detailed Description . . . . .	698
8.21 lib/analogy/internal.h File Reference . . . . .	699
8.21.1 Detailed Description . . . . .	699
8.22 lib/analogy/async.c File Reference . . . . .	700
8.22.1 Detailed Description . . . . .	700
8.23 lib/analogy/calibration.c File Reference . . . . .	701
8.23.1 Detailed Description . . . . .	702
8.24 lib/analogy/descriptor.c File Reference . . . . .	702
8.24.1 Detailed Description . . . . .	703
8.25 lib/analogy/range.c File Reference . . . . .	703
8.25.1 Detailed Description . . . . .	704
8.26 lib/analogy/root_leaf.h File Reference . . . . .	704
8.26.1 Detailed Description . . . . .	705
8.27 lib/analogy/sync.c File Reference . . . . .	705
8.27.1 Detailed Description . . . . .	706
8.28 lib/analogy/sys.c File Reference . . . . .	706
8.28.1 Detailed Description . . . . .	707
<b>9 Example Documentation . . . . .</b>	<b>709</b>
9.1 bufp-label.c . . . . .	709
9.2 bufp-readwrite.c . . . . .	711
9.3 can-rtt.c . . . . .	714
9.4 cross-link.c . . . . .	719
9.5 iddp-label.c . . . . .	722
9.6 iddp-sendrecv.c . . . . .	725
9.7 rtcانconfig.c . . . . .	728
9.8 rtcانrecv.c . . . . .	731
9.9 rtcانsend.c . . . . .	734
9.10 xddp-echo.c . . . . .	738
9.11 xddp-label.c . . . . .	741
9.12 xddp-stream.c . . . . .	745
<b>Index . . . . .</b>	<b>749</b>

# Chapter 1

## API service tags

All services from the Cobalt/POSIX library, or which belong to APIs based on the Copperplate library may be restricted to particular calling contexts, or entail specific side-effects.

In dual kernel mode, the Cobalt API underlies all other application-oriented APIs, providing POSIX real-time services over the Cobalt real-time core. Therefore, the information below applies to all application-oriented APIs available with Xenomai, such as the Cobalt/POSIX library, the Alchemy API, and to all RTOS emulators as well. To describe this information, each service documented by this section bears a set of tags when applicable.

The table below matches the tags used throughout the documentation with the description of their meaning for the caller.

### Attention

By Xenomai thread, we mean any thread created by a Xenomai API service, including real-time Cobalt/POSIX threads in dual kernel mode. By regular/plain POSIX thread, we mean any thread directly created by the standard *glibc-based* POSIX service over Mercury or Cobalt (i.e. NPTL/linuxthreads `__STD(pthread_create())`), excluding such threads which have been promoted to the real-time domain afterwards (aka "shadowed") over Cobalt.

### Context tags

Tag	Context on entry
xthread-only	Must be called from a Xenomai thread
xhandler-only	Must be called from a Xenomai handler. See note.
xcontext	May be called from any Xenomai context (thread or handler).
pthread-only	Must be called from a regular POSIX thread
thread-unrestricted	May be called from a Xenomai or regular POSIX thread indifferently
xthread-nowait	May be called from a Xenomai thread unrestricted, or from a regular thread as a non-blocking service only. See note.
unrestricted	May be called from any context previously described

## Note

A Xenomai handler is used for callback-based notifications from Copperplate-based APIs, such as timeouts. This context is *NOT* mapped to a regular Linux signal handler, it is actually underlaid by a special thread context, so that async-unsafe POSIX services may be invoked internally by the API implementation when running on behalf of such handler. Therefore, calling Xenomai API services from asynchronous regular signal handlers is fundamentally unsafe.

Over Cobalt, the main thread is a particular case, which starts as a regular POSIX thread, then is automatically switched to a Cobalt thread as part of the initialization process, before the main() routine is invoked, unless automatic bootstrap was disabled (see [http://xenomai.org/2015/05/application-setup-and-init/#Application\\_entry\\_CC](http://xenomai.org/2015/05/application-setup-and-init/#Application_entry_CC)).

**Possible side-effects when running the application over the Cobalt core (i.e. dual kernel configuration)**

Tag	Description
switch-primary	the caller may switch to primary mode
switch-secondary	the caller may switch to secondary mode

## Note

As a rule of thumb, any service which might block the caller, causes a switch to primary mode if invoked from secondary mode. This rule might not apply in case the service can complete fully from user-space without any syscall entailed, due to a particular optimization (e.g. fast acquisition of semaphore resources directly from user-space in the non-contended case). Therefore, the switch-{primary, secondary} tags denote either services which *will* always switch the caller to the mode mentioned, or *might* have to do so, depending on the context. The absence of such tag indicates that such services can complete in either modes and as such will entail no switch.

# Chapter 2

## Deprecated List

Global **COMPAT\_\_rt\_alarm\_create** (RT\_ALARM \*alarm, const char \*name)

This is a compatibility service from the Transition Kit.

Global **COMPAT\_\_rt\_event\_clear** (RT\_EVENT \*event, unsigned long mask, unsigned long \*mask\_r)

This is a compatibility service from the Transition Kit.

Global **COMPAT\_\_rt\_event\_create** (RT\_EVENT \*event, const char \*name, unsigned long ivalue, int mode)

This is a compatibility service from the Transition Kit.

Global **COMPAT\_\_rt\_event\_signal** (RT\_EVENT \*event, unsigned long mask)

This is a compatibility service from the Transition Kit.

Global **COMPAT\_\_rt\_task\_create** (RT\_TASK \*task, const char \*name, int stksize, int prio, int mode)

This is a compatibility service from the Transition Kit.

Global **COMPAT\_\_rt\_task\_set\_periodic** (RT\_TASK \*task, RTIME idate, RTIME period)

This is a compatibility service from the Transition Kit.

### Module IOCTLS

Passing struct ifreq as a request descriptor for CAN IOCTLS is still accepted for backward compatibility, however it is recommended to switch to struct [can\\_ifreq](#) at the first opportunity.

### Module IOCTLS

Passing struct ifreq as a request descriptor for CAN IOCTLS is still accepted for backward compatibility, however it is recommended to switch to struct [can\\_ifreq](#) at the first opportunity.

Global **pthread\_make\_periodic\_np** (pthread\_t thread, struct timespec \*starttp, struct timespec \*periodtp)

This service is a non-portable extension of the Xenomai 2.x POSIX interface, not available with Xenomai 3.x. Instead, Cobalt-based applications should set up a periodic timer using the [timer\\_create\(\)](#), [timer\\_settime\(\)](#) call pair, then wait for release points via [sigwaitinfo\(\)](#). Overruns can be detected by looking at the `signfo.si_overrun` field. Alternatively, applications may obtain a file descriptor referring to a Cobalt timer via the [timerfd\(\)](#) call, and [read\(\)](#) from it to wait for timeouts.

Global **pthread\_wait\_np** (unsigned long \*overruns\_r)

This service is a non-portable extension of the Xenomai 2.x POSIX interface, not available with Xenomai 3.x. Instead, Cobalt-based applications should set up a periodic timer using the [timer\\_create\(\)](#), [timer\\_settime\(\)](#) call pair, then wait for release points via [sigwaitinfo\(\)](#). Overruns can be detected by looking at the `signfo.si_overrun` field. Alternatively, applications may obtain a file descriptor referring to a Cobalt timer via the [timerfd\(\)](#) call, and [read\(\)](#) from it to wait for timeouts.

**Global `rt_alarm_wait` (RT\_ALARM \*alarm)**

This is a compatibility service from the Transition Kit.

**Global `RTDM_EXECUTE_ATOMICALY` (code\_block)**

This construct will be phased out in Xenomai 3.0. Please use `rtdm_waitqueue` services instead.

**Global `rtdm_task_sleep_until` (nanosecs\_abs\_t wakeup\_time)**

Use `rtdm_task_sleep_abs` instead!



# Chapter 3

## Module Index

### 3.1 Modules

Here is a list of all modules:

RTDM . . . . .	72
RTDM User API . . . . .	75
Driver programming interface . . . . .	142
Driver to driver services . . . . .	143
Device Registration Services . . . . .	165
Clock Services . . . . .	176
Task Services . . . . .	178
Timer Services . . . . .	189
Synchronisation Services . . . . .	194
Big dual kernel lock . . . . .	27
Spinlock with preemption deactivation . . . . .	30
Event Services . . . . .	208
Semaphore Services . . . . .	214
Mutex services . . . . .	219
Interrupt Management Services . . . . .	223
Non-Real-Time Signalling Services . . . . .	228
Utility Services . . . . .	231
Device Profiles . . . . .	162
User-space driver core . . . . .	34
CAN Devices . . . . .	47
Serial Devices . . . . .	76
Testing Devices . . . . .	78
Real-time IPC . . . . .	79
Cobalt . . . . .	119
Cobalt kernel . . . . .	120
Asynchronous Procedure Calls . . . . .	98
In-kernel arithmetics . . . . .	101
Buffer descriptor . . . . .	102
Clock services . . . . .	112
Debugging services . . . . .	115
Dynamic memory allocation services . . . . .	116
Interrupt management . . . . .	123
Locking services . . . . .	130
Lightweight key-to-object mapping service . . . . .	132

Registry services . . . . .	137
Thread scheduling control . . . . .	246
SCHED_QUOTA scheduling policy . . . . .	245
Synchronous I/O multiplexing . . . . .	248
Thread synchronization services . . . . .	253
Thread services . . . . .	262
Thread runtime statistics . . . . .	42
Thread state flags . . . . .	43
Thread information flags . . . . .	46
Timer services . . . . .	278
Virtual file services . . . . .	287
Analogy framework . . . . .	295
Driver API . . . . .	296
Channels and ranges . . . . .	23
Driver management services . . . . .	297
Subdevice management services . . . . .	299
Buffer management services . . . . .	303
Interrupt management services . . . . .	313
Misc services . . . . .	316
Analogy user API . . . . .	440
Level 1 API . . . . .	430
Asynchronous acquisition API . . . . .	402
Descriptor API . . . . .	417
Synchronous acquisition API . . . . .	431
Level 2 API . . . . .	435
Asynchronous acquisition API . . . . .	410
Software calibration API . . . . .	412
Math API . . . . .	421
Range / conversion API . . . . .	424
Synchronous acquisition API . . . . .	436
Level 0 API . . . . .	441
Descriptor Syscall API . . . . .	415
Basic Syscall API . . . . .	442
Attach / detach Syscall API . . . . .	445
POSIX interface . . . . .	338
Clocks and timers . . . . .	317
Condition variables . . . . .	327
Message queues . . . . .	340
Mutual exclusion . . . . .	351
Process scheduling . . . . .	366
Semaphores . . . . .	378
Thread management . . . . .	386
Scheduling management . . . . .	392
Smokey API . . . . .	398
Alchemy API . . . . .	491
Alarm services . . . . .	448
Buffer services . . . . .	453
Condition variable services . . . . .	464
Event flag group services . . . . .	472
Heap management services . . . . .	482
Mutex services . . . . .	493
Message pipe services . . . . .	501
Message queue services . . . . .	510
Semaphore services . . . . .	523
Task management services . . . . .	532
Timer management services . . . . .	561
VxWorks® emulator . . . . .	565

---

pSOS® emulator . . . . .	566
Transition Kit . . . . .	567



# Chapter 4

## Data Structure Index

### 4.1 Data Structures

Here are the data structures with brief descriptions:

<a href="#">a4l_channel</a>	Structure describing some channel's characteristics . . . . .	579
<a href="#">a4l_channels_desc</a>	Structure describing a channels set . . . . .	580
<a href="#">a4l_cmd_desc</a>	Structure describing the asynchronous instruction . . . . .	581
<a href="#">a4l_descriptor</a>	Structure containing device-information useful to users . . . . .	582
<a href="#">a4l_driver</a>	Structure containing driver declaration data . . . . .	584
<a href="#">a4l_instruction</a>	Structure describing the synchronous instruction . . . . .	585
<a href="#">a4l_instruction_list</a>	Structure describing the list of synchronous instructions . . . . .	586
<a href="#">a4l_range</a>	Structure describing a (unique) range . . . . .	587
<a href="#">a4l_subdevice</a>	Structure describing the subdevice . . . . .	588
<a href="#">atomic_t</a>	Copyright © 2011 Gilles Chanteperrix <a href="mailto:gilles.chanteperrix@xenomai.org">gilles.chanteperrix@xenomai.org</a> . . . . .	589
<a href="#">can_bittime</a>	Custom CAN bit-time definition . . . . .	590
<a href="#">can_bittime_btr</a>	Hardware-specific BTR bit-times . . . . .	591
<a href="#">can_bittime_std</a>	Standard bit-time parameters according to Bosch . . . . .	591
<a href="#">can_filter</a>	Filter for reception of CAN messages . . . . .	592
<a href="#">can_frame</a>	Raw CAN frame . . . . .	593
<a href="#">can_ifreq</a>	CAN interface request descriptor . . . . .	594
<a href="#">macb_dma_desc</a>	Hardware DMA descriptor . . . . .	595
<a href="#">macb_tx_skb</a>	Data about an skb which is being transmitted . . . . .	595

<a href="#">RT_ALARM_INFO</a>	
Alarm status descriptor	596
<a href="#">RT_BUFFER_INFO</a>	
Buffer status descriptor	596
<a href="#">RT_COND_INFO</a>	
Condition variable status descriptor	597
<a href="#">RT_EVENT_INFO</a>	
Event status descriptor	597
<a href="#">RT_HEAP_INFO</a>	
Heap status descriptor	598
<a href="#">RT_MUTEX_INFO</a>	
Mutex status descriptor	599
<a href="#">RT_QUEUE_INFO</a>	
Queue status descriptor	600
<a href="#">RT_SEM_INFO</a>	
Semaphore status descriptor	600
<a href="#">RT_TASK_INFO</a>	
Task status descriptor	601
<a href="#">rt_timer_info</a>	
Timer status descriptor	602
<a href="#">rtdm_dev_context</a>	
Device context	603
<a href="#">rtdm_device</a>	
RTDM device	604
<a href="#">rtdm_device_info</a>	
Device information	606
<a href="#">rtdm_driver</a>	
RTDM driver	606
<a href="#">rtdm_fd_ops</a>	
RTDM file operation descriptor	608
<a href="#">rtdm_gpio_readout</a>	612
<a href="#">rtdm_profile_info</a>	
RTDM profile information	613
<a href="#">rtdm_sm_ops</a>	
RTDM state management handlers	613
<a href="#">rtdm_spi_config</a>	614
<a href="#">rtipc_port_label</a>	
Port label information structure	614
<a href="#">rtser_config</a>	
Serial device configuration	615
<a href="#">rtser_event</a>	
Additional information about serial device events	616
<a href="#">rtser_status</a>	
Serial device status	616
<a href="#">sockaddr_can</a>	
Socket address structure for the CAN address family	617
<a href="#">sockaddr_ipc</a>	
Socket address structure for the RTIPC address family	618
<a href="#">udd_device</a>	619
<a href="#">udd_memregion</a>	623
<a href="#">udd_device::udd_reserved</a>	
Reserved to the UDD core	624
<a href="#">udd_signotify</a>	
UDD event notification descriptor	625
<a href="#">xnsched</a>	
Scheduling information structure	626
<a href="#">xnvfile_lock_ops</a>	
Vfile locking operations	628

---

<a href="#">xnvfile_regular_iterator</a>	
Regular vfile iterator . . . . .	629
<a href="#">xnvfile_regular_ops</a>	
Regular vfile operation descriptor . . . . .	630
<a href="#">xnvfile_rev_tag</a>	
Snapshot revision tag . . . . .	634
<a href="#">xnvfile_snapshot</a>	
Snapshot vfile descriptor . . . . .	635
<a href="#">xnvfile_snapshot_iterator</a>	
Snapshot-driven vfile iterator . . . . .	635
<a href="#">xnvfile_snapshot_ops</a>	
Snapshot vfile operation descriptor . . . . .	638





# Chapter 5

## File Index

### 5.1 File List

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

include/alchemy/ <b>alarm.h</b>	??
include/alchemy/ <b>buffer.h</b>	??
include/alchemy/ <b>compat.h</b>	??
include/alchemy/ <b>cond.h</b>	??
include/alchemy/ <b>event.h</b>	??
include/alchemy/ <b>heap.h</b>	??
include/alchemy/ <b>mutex.h</b>	??
include/alchemy/ <b>pipe.h</b>	??
include/alchemy/ <b>queue.h</b>	??
include/alchemy/ <b>sem.h</b>	??
include/alchemy/ <b>task.h</b>	??
include/alchemy/ <b>timer.h</b>	??
include/boilerplate/ <b>ancillaries.h</b>	??
include/boilerplate/ <b>atomic.h</b>	??
include/boilerplate/ <b>avl-inner.h</b>	??
include/boilerplate/ <b>avl.h</b>	??
include/boilerplate/ <b>compiler.h</b>	??
include/boilerplate/ <b>debug.h</b>	??
include/boilerplate/ <b>hash.h</b>	??
include/boilerplate/ <b>heapmem.h</b>	??
include/boilerplate/ <b>libc.h</b>	??
include/boilerplate/ <b>list.h</b>	??
include/boilerplate/ <b>lock.h</b>	??
include/boilerplate/ <b>namegen.h</b>	??
include/boilerplate/ <b>obstack.h</b>	??
include/boilerplate/ <b>private-list.h</b>	??
include/boilerplate/ <b>scope.h</b>	??
include/boilerplate/ <b>setup.h</b>	??
include/boilerplate/ <b>shared-list.h</b>	??
include/boilerplate/ <b>shavl.h</b>	??
include/boilerplate/ <b>time.h</b>	??
include/boilerplate/ <b>tunables.h</b>	??
include/cobalt/ <b>arith.h</b>	??
include/cobalt/ <b>fcntl.h</b>	??
include/cobalt/ <b>mqueue.h</b>	??

include/cobalt/ <b>pthread.h</b>	??
include/cobalt/ <b>sched.h</b>	??
include/cobalt/ <b>semaphore.h</b>	??
include/cobalt/ <b>signal.h</b>	??
include/cobalt/ <b>stdio.h</b>	??
include/cobalt/ <b>stdlib.h</b>	??
include/cobalt/ <b>syslog.h</b>	??
include/cobalt/ <b>ticks.h</b>	??
include/cobalt/ <b>time.h</b>	??
include/cobalt/ <b>trace.h</b>	??
include/cobalt/ <b>tunables.h</b>	??
include/cobalt/ <b>unistd.h</b>	??
include/cobalt/ <b>wrappers.h</b>	??
include/cobalt/boilerplate/ <b>limits.h</b>	??
include/cobalt/boilerplate/ <b>sched.h</b>	??
include/cobalt/boilerplate/ <b>signal.h</b>	??
include/cobalt/boilerplate/ <b>trace.h</b>	??
include/cobalt/boilerplate/ <b>wrappers.h</b>	??
include/cobalt/kernel/ <b>ancillaries.h</b>	??
include/cobalt/kernel/ <b>apc.h</b>	??
include/cobalt/kernel/ <b>arith.h</b>	??
include/cobalt/kernel/ <b>assert.h</b>	??
include/cobalt/kernel/ <b>bufd.h</b>	??
include/cobalt/kernel/ <b>clock.h</b>	??
include/cobalt/kernel/ <b>compat.h</b>	??
include/cobalt/kernel/ <b>heap.h</b>	??
include/cobalt/kernel/ <b>init.h</b>	??
include/cobalt/kernel/ <b>intr.h</b>	??
include/cobalt/kernel/ <b>list.h</b>	??
include/cobalt/kernel/ <b>lock.h</b>	??
include/cobalt/kernel/ <b>map.h</b>	??
include/cobalt/kernel/ <b>pipe.h</b>	??
include/cobalt/kernel/ <b>ppd.h</b>	??
include/cobalt/kernel/ <b>registry.h</b>	??
include/cobalt/kernel/ <b>sched-idle.h</b>	??
include/cobalt/kernel/ <b>sched-quota.h</b>	??
include/cobalt/kernel/ <b>sched-rt.h</b>	??
include/cobalt/kernel/ <b>sched-sporadic.h</b>	??
include/cobalt/kernel/ <b>sched-tp.h</b>	??
include/cobalt/kernel/ <b>sched-weak.h</b>	??
include/cobalt/kernel/ <b>sched.h</b>	??
include/cobalt/kernel/ <b>schedparam.h</b>	??
include/cobalt/kernel/ <b>schedqueue.h</b>	??
include/cobalt/kernel/ <b>select.h</b>	??
include/cobalt/kernel/ <b>stat.h</b>	??
include/cobalt/kernel/ <b>synch.h</b>	??
include/cobalt/kernel/ <b>thread.h</b>	??
include/cobalt/kernel/ <b>timer.h</b>	??
include/cobalt/kernel/ <b>trace.h</b>	??
include/cobalt/kernel/ <b>tree.h</b>	??
include/cobalt/kernel/ <b>vdso.h</b>	??
include/cobalt/kernel/ <b>vfile.h</b>	??
include/cobalt/kernel/rtdm/ <b>autotune.h</b>	??
include/cobalt/kernel/rtdm/ <b>can.h</b>	??
include/cobalt/kernel/rtdm/ <b>cobalt.h</b>	??
include/cobalt/kernel/rtdm/ <b>compat.h</b>	??
include/cobalt/kernel/rtdm/ <b>driver.h</b>	??
Real-Time Driver Model for Xenomai, driver API header	647

include/cobalt/kernel/rtdm/fd.h	664
include/cobalt/kernel/rtdm/gpio.h	??
include/cobalt/kernel/rtdm/gpiopwm.h	??
include/cobalt/kernel/rtdm/ipc.h	??
include/cobalt/kernel/rtdm/net.h	??
include/cobalt/kernel/rtdm/rtdm.h	??
include/cobalt/kernel/rtdm/serial.h	??
include/cobalt/kernel/rtdm/testing.h	??
include/cobalt/kernel/rtdm/udd.h	
Copyright (C) 2014 Philippe Gerum rpm@xenomai.org	689
include/cobalt/kernel/rtdm/analogy/buffer.h	??
include/cobalt/kernel/rtdm/analogy/channel_range.h	
Analogy for Linux, channel, range related features	643
include/cobalt/kernel/rtdm/analogy/command.h	??
include/cobalt/kernel/rtdm/analogy/context.h	??
include/cobalt/kernel/rtdm/analogy/device.h	??
include/cobalt/kernel/rtdm/analogy/driver.h	
Analogy for Linux, driver facilities	646
include/cobalt/kernel/rtdm/analogy/instruction.h	??
include/cobalt/kernel/rtdm/analogy/rtdm_helpers.h	??
include/cobalt/kernel/rtdm/analogy/subdevice.h	
Analogy for Linux, subdevice related features	654
include/cobalt/kernel/rtdm/analogy/transfer.h	??
include/cobalt/sys/cobalt.h	??
include/cobalt/sys/ioctl.h	??
include/cobalt/sys/mman.h	??
include/cobalt/sys/select.h	??
include/cobalt/sys/socket.h	??
include/cobalt/sys/time.h	??
include/cobalt/sys/timerfd.h	??
include/cobalt/uapi/cond.h	??
include/cobalt/uapi/corectl.h	??
include/cobalt/uapi/event.h	??
include/cobalt/uapi/monitor.h	??
include/cobalt/uapi/mutex.h	??
include/cobalt/uapi/sched.h	??
include/cobalt/uapi/sem.h	??
include/cobalt/uapi/signal.h	??
include/cobalt/uapi/syscall.h	??
include/cobalt/uapi/thread.h	??
include/cobalt/uapi/time.h	??
include/cobalt/uapi/asm-generic/arith.h	??
include/cobalt/uapi/asm-generic/features.h	??
include/cobalt/uapi/asm-generic/syscall.h	??
include/cobalt/uapi/kernel/heap.h	??
include/cobalt/uapi/kernel/limits.h	??
include/cobalt/uapi/kernel/pipe.h	??
include/cobalt/uapi/kernel/synch.h	??
include/cobalt/uapi/kernel/thread.h	??
include/cobalt/uapi/kernel/trace.h	??
include/cobalt/uapi/kernel/types.h	??
include/cobalt/uapi/kernel/urw.h	??
include/cobalt/uapi/kernel/vdso.h	??
include/copperplate/clockobj.h	??
include/copperplate/cluster.h	??
include/copperplate/debug.h	??
include/copperplate/eventobj.h	??
include/copperplate/heapobj.h	??

include/copperplate/ <b>reference.h</b>	??
include/copperplate/ <b>registry-obstack.h</b>	??
include/copperplate/ <b>registry.h</b>	??
include/copperplate/ <b>semobj.h</b>	??
include/copperplate/ <b>syncobj.h</b>	??
include/copperplate/ <b>threadobj.h</b>	??
include/copperplate/ <b>timerobj.h</b>	??
include/copperplate/ <b>traceobj.h</b>	??
include/copperplate/ <b>tunables.h</b>	??
include/copperplate/ <b>wrappers.h</b>	??
include/mercury/ <b>pthread.h</b>	??
include/mercury/boilerplate/ <b>limits.h</b>	??
include/mercury/boilerplate/ <b>sched.h</b>	??
include/mercury/boilerplate/ <b>signal.h</b>	??
include/mercury/boilerplate/ <b>trace.h</b>	??
include/mercury/boilerplate/ <b>wrappers.h</b>	??
include/psos/ <b>psos.h</b>	??
include/psos/ <b>tunables.h</b>	??
include/rtdm/ <b>analogy.h</b>	
Analogy for Linux, library facilities	691
include/rtdm/ <b>autotune.h</b>	??
include/rtdm/ <b>can.h</b>	656
include/rtdm/ <b>gpio.h</b>	??
include/rtdm/ <b>gpiopwm.h</b>	??
include/rtdm/ <b>ipc.h</b>	??
include/rtdm/ <b>net.h</b>	??
include/rtdm/ <b>rtdm.h</b>	674
include/rtdm/ <b>serial.h</b>	
Real-Time Driver Model for Xenomai, serial device profile header	676
include/rtdm/ <b>spi.h</b>	??
include/rtdm/ <b>testing.h</b>	
Real-Time Driver Model for Xenomai, testing device profile header	686
include/rtdm/ <b>udd.h</b>	??
include/rtdm/uapi/ <b>analogy.h</b>	
Analogy for Linux, UAPI bits	693
include/rtdm/uapi/ <b>autotune.h</b>	??
include/rtdm/uapi/ <b>can.h</b>	
Real-Time Driver Model for RT-Socket-CAN, CAN device profile header	657
include/rtdm/uapi/ <b>gpio.h</b>	??
include/rtdm/uapi/ <b>gpiopwm.h</b>	
Real-Time Driver Model for Xenomai, pwm header	669
include/rtdm/uapi/ <b>ipc.h</b>	
This file is part of the Xenomai project	670
include/rtdm/uapi/ <b>net.h</b>	??
include/rtdm/uapi/ <b>rtdm.h</b>	
Real-Time Driver Model for Xenomai, user API header	674
include/rtdm/uapi/ <b>serial.h</b>	
Real-Time Driver Model for Xenomai, serial device profile header	677
include/rtdm/uapi/ <b>spi.h</b>	??
include/rtdm/uapi/ <b>testing.h</b>	
Real-Time Driver Model for Xenomai, testing device profile header	687
include/rtdm/uapi/ <b>udd.h</b>	
This file is part of the Xenomai project	690
include/smokey/ <b>smokey.h</b>	??
include/trank/ <b>rtdk.h</b>	??
include/trank/ <b>trank.h</b>	??
include/trank/native/ <b>alarm.h</b>	??
include/trank/native/ <b>buffer.h</b>	??

include/trunk/native/ <b>cond.h</b>	??
include/trunk/native/ <b>event.h</b>	??
include/trunk/native/ <b>heap.h</b>	??
include/trunk/native/ <b>misc.h</b>	??
include/trunk/native/ <b>mutex.h</b>	??
include/trunk/native/ <b>pipe.h</b>	??
include/trunk/native/ <b>queue.h</b>	??
include/trunk/native/ <b>sem.h</b>	??
include/trunk/native/ <b>task.h</b>	??
include/trunk/native/ <b>timer.h</b>	??
include/trunk/native/ <b>types.h</b>	??
include/trunk/posix/ <b>pthread.h</b>	??
include/trunk/rtdm/ <b>rtcan.h</b>	??
include/trunk/rtdm/ <b>rtdm.h</b>	??
include/trunk/rtdm/ <b>rtipc.h</b>	??
include/trunk/rtdm/ <b>rtserial.h</b>	??
include/trunk/rtdm/ <b>rttesting.h</b>	??
include/vxworks/ <b>errnoLib.h</b>	??
include/vxworks/ <b>intLib.h</b>	??
include/vxworks/ <b>kernLib.h</b>	??
include/vxworks/ <b>lstLib.h</b>	??
include/vxworks/ <b>memPartLib.h</b>	??
include/vxworks/ <b>msgQLib.h</b>	??
include/vxworks/ <b>rngLib.h</b>	??
include/vxworks/ <b>semLib.h</b>	??
include/vxworks/ <b>sysLib.h</b>	??
include/vxworks/ <b>taskHookLib.h</b>	??
include/vxworks/ <b>taskInfo.h</b>	??
include/vxworks/ <b>taskLib.h</b>	??
include/vxworks/ <b>tickLib.h</b>	??
include/vxworks/ <b>types.h</b>	??
include/vxworks/ <b>wdLib.h</b>	??
include/xenomai/ <b>init.h</b>	??
include/xenomai/ <b>tunables.h</b>	??
include/xenomai/ <b>version.h</b>	??
kernel/cobalt/ <b>debug.h</b>	??
kernel/cobalt/ <b>procfs.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/ <b>calibration.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/ <b>features.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/ <b>fpctest.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/ <b>machine.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/ <b>syscall.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/ <b>syscall32.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/ <b>thread.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/ <b>wrappers.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/uapi/ <b>arith.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/uapi/ <b>features.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/uapi/ <b>fpctest.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/uapi/ <b>syscall.h</b>	??
kernel/cobalt/arch/arm/include/asm/xenomai/uapi/ <b>tsc.h</b>	??
kernel/cobalt/arch/arm64/include/asm/xenomai/ <b>calibration.h</b>	??
kernel/cobalt/arch/arm64/include/asm/xenomai/ <b>features.h</b>	??
kernel/cobalt/arch/arm64/include/asm/xenomai/ <b>fpctest.h</b>	??
kernel/cobalt/arch/arm64/include/asm/xenomai/ <b>machine.h</b>	??
kernel/cobalt/arch/arm64/include/asm/xenomai/ <b>syscall.h</b>	??
kernel/cobalt/arch/arm64/include/asm/xenomai/ <b>syscall32.h</b>	??
kernel/cobalt/arch/arm64/include/asm/xenomai/ <b>thread.h</b>	??
kernel/cobalt/arch/arm64/include/asm/xenomai/ <b>wrappers.h</b>	??

kernel/cobalt/arch/arm64/include/asm/xenomai/uapi/arith.h	??
kernel/cobalt/arch/arm64/include/asm/xenomai/uapi/features.h	??
kernel/cobalt/arch/arm64/include/asm/xenomai/uapi/fptest.h	??
kernel/cobalt/arch/arm64/include/asm/xenomai/uapi/syscall.h	??
kernel/cobalt/arch/arm64/include/asm/xenomai/uapi/tsc.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/calibration.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/features.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/fptest.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/machine.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/syscall.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/syscall32.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/thread.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/wrappers.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/uapi/arith.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/uapi/features.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/uapi/fptest.h	??
kernel/cobalt/arch/powerpc/include/asm/xenomai/uapi/syscall.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/c1e.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/calibration.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/features.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/fptest.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/machine.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/smi.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/syscall.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/syscall32-table.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/syscall32.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/thread.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/wrappers.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/uapi/arith.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/uapi/features.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/uapi/fptest.h	??
kernel/cobalt/arch/x86/include/asm/xenomai/uapi/syscall.h	??
kernel/cobalt/include/asm-generic/xenomai/machine.h	??
kernel/cobalt/include/asm-generic/xenomai/pci_ids.h	??
kernel/cobalt/include/asm-generic/xenomai/syscall.h	??
kernel/cobalt/include/asm-generic/xenomai/syscall32.h	??
kernel/cobalt/include/asm-generic/xenomai/thread.h	??
kernel/cobalt/include/asm-generic/xenomai/wrappers.h	??
kernel/cobalt/include/ipipe/thread_info.h	??
kernel/cobalt/include/linux/xenomai/wrappers.h	??
kernel/cobalt/posix/clock.h	??
kernel/cobalt/posix/cond.h	??
kernel/cobalt/posix/corectl.h	??
kernel/cobalt/posix/event.h	??
kernel/cobalt/posix/extension.h	??
kernel/cobalt/posix/internal.h	??
kernel/cobalt/posix/io.h	??
kernel/cobalt/posix/memory.h	??
kernel/cobalt/posix/monitor.h	??
kernel/cobalt/posix/mqueue.h	??
kernel/cobalt/posix/mutex.h	??
kernel/cobalt/posix/process.h	??
kernel/cobalt/posix/sched.h	??
kernel/cobalt/posix/sem.h	??
kernel/cobalt/posix/signal.h	??
kernel/cobalt/posix/syscall.h	??
kernel/cobalt/posix/syscall32.h	??
kernel/cobalt/posix/thread.h	??

kernel/cobalt/posix/timer.h	??
kernel/cobalt/posix/timerfd.h	??
kernel/cobalt/rtdm/internal.h	??
kernel/cobalt/trace/cobalt-core.h	??
kernel/cobalt/trace/cobalt-posix.h	??
kernel/cobalt/trace/cobalt-rtdm.h	??
kernel/drivers/analogy/proc.h	??
kernel/drivers/analogy/intel/8255.h	??
kernel/drivers/analogy/national_instruments/mite.h	??
kernel/drivers/analogy/national_instruments/ni_mio.h	??
kernel/drivers/analogy/national_instruments/ni_stc.h	??
kernel/drivers/analogy/national_instruments/ni_tio.h	??
kernel/drivers/can/rtcan_dev.h	??
kernel/drivers/can/rtcan_internal.h	??
kernel/drivers/can/rtcan_list.h	??
kernel/drivers/can/rtcan_raw.h	??
kernel/drivers/can/rtcan_socket.h	??
kernel/drivers/can/rtcan_version.h	??
kernel/drivers/can/mscan/rtcan_mscan.h	??
kernel/drivers/can/mscan/rtcan_mscan_regs.h	??
kernel/drivers/can/sja1000/rtcan_sja1000.h	??
kernel/drivers/can/sja1000/rtcan_sja1000_regs.h	??
kernel/drivers/ipc/internal.h	??
kernel/drivers/net/drivers/rt_at91_ether.h	??
kernel/drivers/net/drivers/rt_eth1394.h	??
kernel/drivers/net/drivers/rt_fec.h	??
kernel/drivers/net/drivers/rt_macb.h	??
kernel/drivers/net/drivers/rt_smc91111.h	??
kernel/drivers/net/drivers/e1000/e1000.h	??
kernel/drivers/net/drivers/e1000/e1000_hw.h	??
kernel/drivers/net/drivers/e1000/e1000_osdep.h	??
kernel/drivers/net/drivers/e1000/kcompat.h	??
kernel/drivers/net/drivers/e1000e/defines.h	??
kernel/drivers/net/drivers/e1000e/e1000.h	??
kernel/drivers/net/drivers/e1000e/hw.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_80003es2lan.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_82541.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_82543.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_82571.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_api.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_defines.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_hw.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_ich8lan.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_mac.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_manage.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_nvme.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_osdep.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_phy.h	??
kernel/drivers/net/drivers/experimental/e1000/e1000_regs.h	??
kernel/drivers/net/drivers/experimental/e1000/kcompat.h	??
kernel/drivers/net/drivers/experimental/rt2500/rt2500pci.h	??
kernel/drivers/net/drivers/experimental/rt2500/rt2x00.h	??
kernel/drivers/net/drivers/igb/e1000_82575.h	??
kernel/drivers/net/drivers/igb/e1000_defines.h	??
kernel/drivers/net/drivers/igb/e1000_hw.h	??
kernel/drivers/net/drivers/igb/e1000_i210.h	??
kernel/drivers/net/drivers/igb/e1000_mac.h	??

kernel/drivers/net/drivers/igb/e1000_mbx.h	??
kernel/drivers/net/drivers/igb/e1000_nvm.h	??
kernel/drivers/net/drivers/igb/e1000_phy.h	??
kernel/drivers/net/drivers/igb/e1000_regs.h	??
kernel/drivers/net/drivers/igb/igb.h	??
kernel/drivers/net/drivers/mpc52xx_fec/rt_mpc52xx_fec.h	??
kernel/drivers/net/drivers/tulip/tulip.h	??
kernel/drivers/net/stack/include/ipv4_chrdev.h	??
kernel/drivers/net/stack/include/nomac_chrdev.h	??
kernel/drivers/net/stack/include/rtcfg_chrdev.h	??
kernel/drivers/net/stack/include/rtdev.h	??
kernel/drivers/net/stack/include/rtdev_mgr.h	??
kernel/drivers/net/stack/include/rtnet.h	??
kernel/drivers/net/stack/include/rtnet_chrdev.h	??
kernel/drivers/net/stack/include/rtnet_internal.h	??
kernel/drivers/net/stack/include/rtnet_iovec.h	??
kernel/drivers/net/stack/include/rtnet_port.h	??
kernel/drivers/net/stack/include/rtnet_rtpc.h	??
kernel/drivers/net/stack/include/rtnet_socket.h	??
kernel/drivers/net/stack/include/rtskb.h	??
kernel/drivers/net/stack/include/rtskb_fifo.h	??
kernel/drivers/net/stack/include/rtwlan.h	??
kernel/drivers/net/stack/include/rtwlan_io.h	??
kernel/drivers/net/stack/include/stack_mgr.h	??
kernel/drivers/net/stack/include/tdma_chrdev.h	??
kernel/drivers/net/stack/include/ethernet/eth.h	??
kernel/drivers/net/stack/include/ipv4/af_inet.h	??
kernel/drivers/net/stack/include/ipv4/arp.h	??
kernel/drivers/net/stack/include/ipv4/icmp.h	??
kernel/drivers/net/stack/include/ipv4/ip_fragment.h	??
kernel/drivers/net/stack/include/ipv4/ip_input.h	??
kernel/drivers/net/stack/include/ipv4/ip_output.h	??
kernel/drivers/net/stack/include/ipv4/ip_sock.h	??
kernel/drivers/net/stack/include/ipv4/protocol.h	??
kernel/drivers/net/stack/include/ipv4/route.h	??
kernel/drivers/net/stack/include/ipv4/tcp.h	??
kernel/drivers/net/stack/include/ipv4/udp.h	??
kernel/drivers/net/stack/include/rtcfg/rtcfg.h	??
kernel/drivers/net/stack/include/rtcfg/rtcfg_client_event.h	??
kernel/drivers/net/stack/include/rtcfg/rtcfg_conn_event.h	??
kernel/drivers/net/stack/include/rtcfg/rtcfg_event.h	??
kernel/drivers/net/stack/include/rtcfg/rtcfg_file.h	??
kernel/drivers/net/stack/include/rtcfg/rtcfg_frame.h	??
kernel/drivers/net/stack/include/rtcfg/rtcfg_ioctl.h	??
kernel/drivers/net/stack/include/rtcfg/rtcfg_proc.h	??
kernel/drivers/net/stack/include/rtcfg/rtcfg_timer.h	??
kernel/drivers/net/stack/include/rtnet/rtnet_disc.h	??
kernel/drivers/net/stack/include/rtnet/rtnet_proc.h	??
kernel/drivers/net/stack/include/rtnet/rtnet_proto.h	??
kernel/drivers/net/stack/include/rtnet/rtnet_vnic.h	??
kernel/drivers/net/stack/include/rtnet/nomac/nomac.h	??
kernel/drivers/net/stack/include/rtnet/nomac/nomac_dev.h	??
kernel/drivers/net/stack/include/rtnet/nomac/nomac_ioctl.h	??
kernel/drivers/net/stack/include/rtnet/nomac/nomac_proto.h	??
kernel/drivers/net/stack/include/rtnet/tdma/tdma.h	??
kernel/drivers/net/stack/include/rtnet/tdma/tdma_dev.h	??
kernel/drivers/net/stack/include/rtnet/tdma/tdma_ioctl.h	??
kernel/drivers/net/stack/include/rtnet/tdma/tdma_proto.h	??



kernel/drivers/net/stack/include/rtnetlink/tdma/tdma_worker.h	??
kernel/drivers/net/stack/ipv4/tcp/timerwheel.h	??
kernel/drivers/serial/16550A_io.h	??
kernel/drivers/serial/16550A_pci.h	??
kernel/drivers/serial/16550A_pnp.h	??
kernel/drivers/spi/spi-device.h	??
kernel/drivers/spi/spi-master.h	??
lib/alchemy/alarm.h	??
lib/alchemy/buffer.h	??
lib/alchemy/cond.h	??
lib/alchemy/event.h	??
lib/alchemy/heap.h	??
lib/alchemy/internal.h	??
lib/alchemy/mutex.h	??
lib/alchemy/pipe.h	??
lib/alchemy/queue.h	??
lib/alchemy/reference.h	??
lib/alchemy/sem.h	??
lib/alchemy/task.h	??
lib/alchemy/timer.h	??
lib/analogy/async.c	
Analogy for Linux, command, transfer, etc	700
lib/analogy/calibration.c	
Analogy for Linux, device, subdevice, etc	701
lib/analogy/calibration.h	
Analogy for Linux, internal calibration declarations	698
lib/analogy/descriptor.c	
Analogy for Linux, descriptor related features	702
lib/analogy/internal.h	
Analogy for Linux, internal declarations	699
lib/analogy/range.c	
Analogy for Linux, range related features	703
lib/analogy/root_leaf.h	
Analogy for Linux, root / leaf system	704
lib/analogy/sync.c	
Analogy for Linux, instruction related features	705
lib/analogy/sys.c	
Analogy for Linux, descriptor related features	706
lib/cobalt/current.h	??
lib/cobalt/internal.h	??
lib/cobalt/umm.h	??
lib/cobalt/arch/arm/include/asm/xenomai/features.h	??
lib/cobalt/arch/arm/include/asm/xenomai/syscall.h	??
lib/cobalt/arch/arm/include/asm/xenomai/tsc.h	??
lib/cobalt/arch/arm64/include/asm/xenomai/features.h	??
lib/cobalt/arch/arm64/include/asm/xenomai/syscall.h	??
lib/cobalt/arch/arm64/include/asm/xenomai/tsc.h	??
lib/cobalt/arch/powerpc/include/asm/xenomai/features.h	??
lib/cobalt/arch/powerpc/include/asm/xenomai/syscall.h	??
lib/cobalt/arch/powerpc/include/asm/xenomai/tsc.h	??
lib/cobalt/arch/x86/include/asm/xenomai/features.h	??
lib/cobalt/arch/x86/include/asm/xenomai/syscall.h	??
lib/cobalt/arch/x86/include/asm/xenomai/tsc.h	??
lib/copperplate/internal.h	??
lib/copperplate/regd/sysregfs.h	??
lib/psos/internal.h	??
lib/psos/pt.h	??
lib/psos/queue.h	??

---

lib/psos/ <b>reference.h</b>	??
lib/psos/ <b>rn.h</b>	??
lib/psos/ <b>sem.h</b>	??
lib/psos/ <b>task.h</b>	??
lib/psos/ <b>tm.h</b>	??
lib/trank/ <b>internal.h</b>	??
lib/vxworks/ <b>memPartLib.h</b>	??
lib/vxworks/ <b>msgQLib.h</b>	??
lib/vxworks/ <b>reference.h</b>	??
lib/vxworks/ <b>rngLib.h</b>	??
lib/vxworks/ <b>semLib.h</b>	??
lib/vxworks/ <b>taskHookLib.h</b>	??
lib/vxworks/ <b>taskLib.h</b>	??
lib/vxworks/ <b>tickLib.h</b>	??
lib/vxworks/ <b>wdLib.h</b>	??


# Chapter 6

## Module Documentation

### 6.1 Channels and ranges

#### Channels

Collaboration diagram for Channels and ranges:



#### Data Structures

- struct [a4l\\_channel](#)  
*Structure describing some channel's characteristics.*
- struct [a4l\\_channels\\_desc](#)  
*Structure describing a channels set.*
- struct [a4l\\_range](#)  
*Structure describing a (unique) range.*

## Macros

- #define [A4L\\_CHAN\\_GLOBAL](#) 0x10  
*Internal use flag (must not be used by driver developer)*
- #define [A4L\\_RNG\\_GLOBAL](#) 0x8  
*Internal use flag (must not be used by driver developer)*
- #define [RANGE](#)(x, y)  
*Macro to declare a (unique) range with no unit defined.*
- #define [RANGE\\_V](#)(x, y)  
*Macro to declare a (unique) range in Volt.*
- #define [RANGE\\_mA](#)(x, y)  
*Macro to declare a (unique) range in milliAmpere.*
- #define [RANGE\\_ext](#)(x, y)  
*Macro to declare a (unique) range in some external reference.*
- #define [A4L\\_RNG\\_GLOBAL\\_RNGDESC](#) 0  
*Constant to define a ranges descriptor as global (inter-channel)*
- #define [A4L\\_RNG\\_PERCHAN\\_RNGDESC](#) 1  
*Constant to define a ranges descriptor as specific for a channel.*
- #define [RNG\\_GLOBAL](#)(x)  
*Macro to declare a ranges global descriptor in one line.*

## Channel reference

### Flags to define the channel's reference

- #define [A4L\\_CHAN\\_AREF\\_GROUND](#) 0x1  
*Ground reference.*
- #define [A4L\\_CHAN\\_AREF\\_COMMON](#) 0x2  
*Common reference.*
- #define [A4L\\_CHAN\\_AREF\\_DIFF](#) 0x4  
*Differential reference.*
- #define [A4L\\_CHAN\\_AREF\\_OTHER](#) 0x8  
*Misc reference.*

## Channels declaration mode

### Constant to define whether the channels in a descriptor are identical

- #define [A4L\\_CHAN\\_GLOBAL\\_CHANDESC](#) 0  
*Global declaration, the set contains channels with similar characteristics.*
- #define [A4L\\_CHAN\\_PERCHAN\\_CHANDESC](#) 1  
*Per channel declaration, the descriptor gathers differents channels.*

### 6.1.1 Detailed Description

#### Channels

According to the Analogy nomenclature, the channel is the elementary acquisition entity. One channel is supposed to acquire one data at a time. A channel can be:

- an analog input or an analog output;
- a digital input or a digital output;

Channels are defined by their type and by some other characteristics like:

- their resolutions for analog channels (which usually ranges from 8 to 32 bits);
- their references;

Such parameters must be declared for each channel composing a subdevice. The structure `a4l_channel` (struct `a4l_channel`) is used to define one channel.

Another structure named `a4l_channels_desc` (struct `a4l_channels_desc`) gathers all channels for a specific subdevice. This latter structure also stores :

- the channels count;
- the channels declaration mode (`A4L_CHAN_GLOBAL_CHANDESC` or `A4L_CHAN_PERCHAN_CHANDESC`): if all the channels composing a subdevice are identical, there is no need to declare the parameters for each channel; the global declaration mode eases the structure composition.

Usually the channels descriptor looks like this:

```
struct a4l_channels_desc example_chan = {
    mode: A4L_CHAN_GLOBAL_CHANDESC, -> Global declaration
                                     mode is set
    length: 8, -> 8 channels
    chans: {
        {A4L_CHAN_AREF_GROUND, 16}, -> Each channel is 16 bits
                                     wide with the ground as
                                     reference
    },
};
```

#### Ranges

So as to perform conversion from logical values acquired by the device to physical units, some range structure(s) must be declared on the driver side.

Such structures contain:

- the physical unit type (Volt, Ampere, none);
- the minimal and maximal values;

These range structures must be associated with the channels at subdevice registration time as a channel can work with many ranges. At configuration time (thanks to an Analogy command), one range will be selected for each enabled channel.

Consequently, for each channel, the developer must declare all the possible ranges in a structure called `struct a4l_rngtab`. Here is an example:

```
struct a4l_rngtab example_tab = {
    length: 2,
    rngs: {
        RANGE_V(-5,5),
        RANGE_V(-10,10),
    },
};
```

For each subdevice, a specific structure is designed to gather all the ranges tabs of all the channels. In this structure, called `struct a4l_rngdesc`, three fields must be filled:

- the declaration mode (`A4L_RNG_GLOBAL_RNGDESC` or `A4L_RNG_PERCHAN_RNGDESC`);
- the number of ranges tab;
- the tab of ranges tabs pointers;

Most of the time, the channels which belong to the same subdevice use the same set of ranges. So, there is no need to declare the same ranges for each channel. A macro is defined to prevent redundant declarations: `RNG_GLOBAL()`.

Here is an example:

```
struct a4l_rngdesc example_rng = RNG_GLOBAL(example_tab);
```

## 6.2 Big dual kernel lock

Collaboration diagram for Big dual kernel lock:



### Macros

- `#define cobalt_atomic_enter(__context)`  
*Enter atomic section (dual kernel only)*
- `#define cobalt_atomic_leave(__context)`  
*Leave atomic section (dual kernel only)*
- `#define RTDM_EXECUTE_ATOMICALY(code_block)`  
*Execute code block atomically (DEPRECATED)*

### 6.2.1 Detailed Description

### 6.2.2 Macro Definition Documentation

#### 6.2.2.1 cobalt\_atomic\_enter

```
#define cobalt_atomic_enter(  
    __context )
```

#### Value:

```
do {  
    xnlock_get_irqsave(&nklock, (__context));  
    xnsched_lock();  
} while (0)
```

Enter atomic section (dual kernel only)

This call opens a fully atomic section, serializing execution with respect to all interrupt handlers (including for real-time IRQs) and Xenomai threads running on all CPUs.

## Parameters

<code>__context</code>	name of local variable to store the context in. This variable updated by the real-time core will hold the information required to leave the atomic section properly.
------------------------	--

## Note

Atomic sections may be nested. The caller is allowed to sleep on a blocking Xenomai service from primary mode within an atomic section delimited by `cobalt_atomic_enter/cobalt_atomic_leave` calls. On the contrary, sleeping on a regular Linux kernel service while holding such lock is NOT valid.

Since the strongest lock is acquired by this service, it can be used to synchronize real-time and non-real-time contexts.

## Warning

This service is not portable to the Mercury core, and should be restricted to Cobalt-specific use cases, mainly for the purpose of porting existing dual-kernel drivers which still depend on the obsolete `RTDM_EXECUTE_ATOMICALY()` construct.

6.2.2.2 `cobalt_atomic_leave`

```
#define cobalt_atomic_leave(  
    __context )
```

**Value:**

```
do {  
    xnsched_unlock();  
    xnlock_put_irqrestore(&nklock, (__context));  
} while (0)
```

## Leave atomic section (dual kernel only)

This call closes an atomic section previously opened by a call to `cobalt_atomic_enter()`, restoring the preemption and interrupt state which prevailed prior to entering the exited section.

## Parameters

<code>__context</code>	name of local variable which stored the context.
------------------------	--

## Warning

This service is not portable to the Mercury core, and should be restricted to Cobalt-specific use cases.



## 6.2.2.3 RTDM\_EXECUTE\_ATOMICALY

```
#define RTDM_EXECUTE_ATOMICALY(  
    code_block )
```

**Value:**

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

## Execute code block atomically (DEPRECATED)

Generally, it is illegal to suspend the current task by calling `rtm_task_sleep()`, `rtm_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

<code>code_block</code>	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.

## Tags

[unrestricted](#)


**Deprecated** This construct will be phased out in Xenomai 3.0. Please use `rtm_waitqueue` services instead.

## See also

[cobalt\\_atomic\\_enter\(\)](#).

## 6.3 Spinlock with preemption deactivation

Collaboration diagram for Spinlock with preemption deactivation:



### Macros

- #define [RTDM\\_LOCK\\_UNLOCKED](#)(\_\_name) IPIPE\_SPIN\_LOCK\_UNLOCKED  
*Static lock initialisation.*
- #define [rtdm\\_lock\\_get\\_irqsave](#)(\_\_lock, \_\_context) ((\_\_context) = \_\_rtdm\_lock\_get\_irqsave(\_\_↵  
lock))  
  
*Acquire lock and disable preemption, by stalling the head domain.*
- #define [rtdm\\_lock\\_irqsave](#)(\_\_context) [splhigh](#)(\_\_context)  
  
*Disable preemption locally*
- #define [rtdm\\_lock\\_irqrestore](#)(\_\_context) [splexit](#)(\_\_context)  
  
*Restore preemption state*

### Typedefs

- typedef ipipe\_spinlock\_t [rtdm\\_lock\\_t](#)  
*Lock variable.*
- typedef unsigned long [rtdm\\_lockctx\\_t](#)  
*Variable to save the context while holding a lock.*

### Functions

- static void [rtdm\\_lock\\_init](#) ([rtdm\\_lock\\_t](#) \*lock)  
  
*Dynamic lock initialisation*
- static void [rtdm\\_lock\\_get](#) ([rtdm\\_lock\\_t](#) \*lock)  
  
*Acquire lock from non-preemptible contexts*
- static void [rtdm\\_lock\\_put](#) ([rtdm\\_lock\\_t](#) \*lock)  
  
*Release lock without preemption restoration*
- static void [rtdm\\_lock\\_put\\_irqrestore](#) ([rtdm\\_lock\\_t](#) \*lock, [rtdm\\_lockctx\\_t](#) context)  
  
*Release lock and restore preemption state*

## 6.3.1 Detailed Description

## 6.3.2 Macro Definition Documentation

6.3.2.1 `rtm_lock_get_irqsave`

```
#define rtm_lock_get_irqsave(  
    __lock,  
    __context ) ((__context) = __rtm_lock_get_irqsave(__lock))
```

Acquire lock and disable preemption, by stalling the head domain.

Parameters

<code>__lock</code>	Address of lock variable
<code>__context</code>	name of local variable to store the context in

Tags

[unrestricted](#)

6.3.2.2 `rtm_lock_irqrestore`

```
#define rtm_lock_irqrestore(  
    __context ) splexit(__context)
```

Restore preemption state

Parameters

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

Tags

[unrestricted](#)

6.3.2.3 `rtm_lock_irqsave`

```
#define rtm_lock_irqsave(  
    __context ) splhigh(__context)
```

Disable preemption locally

Parameters

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

Tags

[unrestricted](#)

### 6.3.3 Function Documentation

#### 6.3.3.1 `rtm_lock_get()`

```
static void rtm_lock_get (  
    rtm_lock_t * lock ) [inline], [static]
```

Acquire lock from non-preemptible contexts

Parameters

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

Tags

[unrestricted](#)

#### 6.3.3.2 `rtm_lock_init()`

```
static void rtm_lock_init (  
    rtm_lock_t * lock ) [inline], [static]
```

Dynamic lock initialisation

Parameters

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

Tags

[task-unrestricted](#)

### 6.3.3.3 rtdm\_lock\_put()

```
static void rtdm_lock_put (  
    rtdm_lock_t * lock ) [inline], [static]
```

Release lock without preemption restoration

Parameters

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

Tags

[unrestricted](#), [might-switch](#)

### 6.3.3.4 rtdm\_lock\_put\_irqrestore()

```
static void rtdm_lock_put_irqrestore (  
    rtdm_lock_t * lock,  
    rtdm_lockctx_t context ) [inline], [static]
```

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


Tags

[unrestricted](#)

## 6.4 User-space driver core

This profile includes all mini-drivers sitting on top of the User-space Device Driver framework (UDD).

Collaboration diagram for User-space driver core:



### Data Structures

- struct [udd\\_memregion](#)
- struct [udd\\_device](#)
- struct [udd\\_signotify](#)  
*UDD event notification descriptor.*

### Functions

- int [udd\\_register\\_device](#) (struct [udd\\_device](#) \*udd)  
*Register a UDD device*
- int [udd\\_unregister\\_device](#) (struct [udd\\_device](#) \*udd)  
*Unregister a UDD device*
- struct [udd\\_device](#) \* [udd\\_get\\_device](#) (struct [rtm\\_fd](#) \*fd)  
*RTDM file descriptor to target UDD device*
- void [udd\\_notify\\_event](#) (struct [udd\\_device](#) \*udd)  
*Notify an IRQ event for an unmanaged interrupt*
- void [udd\\_enable\\_irq](#) (struct [udd\\_device](#) \*udd, [rtm\\_event\\_t](#) \*done)  
*Enable the device IRQ line*
- void [udd\\_disable\\_irq](#) (struct [udd\\_device](#) \*udd, [rtm\\_event\\_t](#) \*done)  
*Disable the device IRQ line*
- #define [UDD\\_IRQ\\_NONE](#) 0  
*No IRQ managed.*
- #define [UDD\\_IRQ\\_CUSTOM](#) (-1)  
*IRQ directly managed from the mini-driver on top of the UDD core.*

## Memory types for mapping

### Types of memory for mapping

The UDD core implements a default `->mmap()` handler which first attempts to hand over the request to the corresponding handler defined by the mini-driver. If not present, the UDD core establishes the mapping automatically, depending on the memory type defined for the region.

- `#define UDD_MEM_NONE 0`  
*No memory region.*
- `#define UDD_MEM_PHYS 1`  
*Physical I/O memory region.*
- `#define UDD_MEM_LOGICAL 2`  
*Kernel logical memory region (e.g.*
- `#define UDD_MEM_VIRTUAL 3`  
*Virtual memory region with no direct physical mapping (e.g.*

## UDD\_IOCTL

### IOCTL requests

- `#define UDD_RTIOC_IRQEN _IO(RTDM_CLASS_UDD, 0)`  
*Enable the interrupt line.*
- `#define UDD_RTIOC_IRQDIS _IO(RTDM_CLASS_UDD, 1)`  
*Disable the interrupt line.*
- `#define UDD_RTIOC_IRQSIG _IOW(RTDM_CLASS_UDD, 2, struct udd_notify)`  
*Enable/Disable signal notification upon interrupt event.*

### 6.4.1 Detailed Description

This profile includes all mini-drivers sitting on top of the User-space Device Driver framework (UDD).

The generic UDD core driver enables interrupt control and I/O memory access interfaces to user-space device drivers, as defined by the mini-drivers when registering.

A mini-driver supplements the UDD core with ancillary functions for dealing with [memory mappings](#) and [interrupt control](#) for a particular I/O card/device.

UDD-compliant mini-drivers only have to provide the basic support for dealing with the interrupt sources present in the device, so that most part of the device requests can be handled from a Xenomai application running in user-space. Typically, a mini-driver would handle the interrupt top-half, and the user-space application would handle the bottom-half.

This profile is reminiscent of the UIO framework available with the Linux kernel, adapted to the dual kernel Cobalt environment.

### 6.4.2 Macro Definition Documentation

#### 6.4.2.1 UDD\_IRQ\_CUSTOM

```
#define UDD_IRQ_CUSTOM (-1)
```

IRQ directly managed from the mini-driver on top of the UDD core.

The mini-driver is in charge of attaching the handler(s) to the IRQ(s) it manages, notifying the Cobalt threads waiting for IRQ events by calling the [udd\\_notify\\_event\(\)](#) service.

Referenced by [udd\\_register\\_device\(\)](#).

#### 6.4.2.2 UDD\_IRQ\_NONE

```
#define UDD_IRQ_NONE 0
```

No IRQ managed.

Special IRQ values for [udd\\_device.irq](#) Passing this code implicitly disables all interrupt-related services, including control (disable/enable) and notification.

Referenced by [udd\\_register\\_device\(\)](#).

#### 6.4.2.3 UDD\_MEM\_LOGICAL

```
#define UDD_MEM_LOGICAL 2
```

Kernel logical memory region (e.g.

[kmalloct\(\)](#)). By default, the UDD core maps such memory to a virtual user range by calling the [rtdm\\_mmap\\_kmem\(\)](#) service.

#### 6.4.2.4 UDD\_MEM\_NONE

```
#define UDD_MEM_NONE 0
```

No memory region.

Use this type code to disable an entry in the array of memory mappings, i.e. [udd\\_device.mem\\_regions\[\]](#).

#### 6.4.2.5 UDD\_MEM\_PHYS

```
#define UDD_MEM_PHYS 1
```

Physical I/O memory region.

By default, the UDD core maps such memory to a virtual user range by calling the [rtdm\\_mmap\\_iomem\(\)](#) service.



#### 6.4.2.6 UDD\_MEM\_VIRTUAL

```
#define UDD_MEM_VIRTUAL 3
```

Virtual memory region with no direct physical mapping (e.g.

`vmalloc()`). By default, the UDD core maps such memory to a virtual user range by calling the `rtm_mmap_vmem()` service.

#### 6.4.2.7 UDD\_RTIOC\_IRQDIS

```
#define UDD_RTIOC_IRQDIS _IO(RTDM_CLASS_UDD, 1)
```

Disable the interrupt line.

The UDD-class mini-driver should handle this request when received through its `->iocctl()` handler if provided. Otherwise, the UDD core disables the interrupt line in the interrupt controller before returning to the caller.

Note

The mini-driver must handle the `UDD_RTIOC_IRQEN` request for a custom IRQ from its `->iocctl()` handler, otherwise such request receives `-EIO` from the UDD core.

#### 6.4.2.8 UDD\_RTIOC\_IRQEN

```
#define UDD_RTIOC_IRQEN _IO(RTDM_CLASS_UDD, 0)
```

Enable the interrupt line.

The UDD-class mini-driver should handle this request when received through its `->iocctl()` handler if provided. Otherwise, the UDD core enables the interrupt line in the interrupt controller before returning to the caller.

#### 6.4.2.9 UDD\_RTIOC\_IRQSIG

```
#define UDD_RTIOC_IRQSIG _IOW(RTDM_CLASS_UDD, 2, struct udd_notify)
```

Enable/Disable signal notification upon interrupt event.

A valid [notification descriptor](#) must be passed along with this request, which is handled by the UDD core directly.

Note

The mini-driver must handle the `UDD_RTIOC_IRQDIS` request for a custom IRQ from its `->iocctl()` handler, otherwise such request receives `-EIO` from the UDD core.

### 6.4.3 Function Documentation

#### 6.4.3.1 udd\_disable\_irq()

```
void udd_disable_irq (
    struct udd_device * udd,
    rtdm_event_t * done )
```

Disable the device IRQ line

This service issues a request to the regular kernel for disabling the IRQ line registered by the driver. If the caller runs in primary mode, the request is scheduled but deferred until the current CPU leaves the real-time domain (see note). Otherwise, the request is immediately handled.

Parameters

<i>udd</i>	The UDD driver handling the IRQ to disable. If no IRQ was registered by the driver at the UDD core, this routine has no effect.
<i>done</i>	Optional event to signal upon completion. If non-NULL, <i>done</i> will be posted by a call to <a href="#">rtdm_event_signal()</a> after the interrupt line is disabled.

Tags

[unrestricted](#)

Note

The deferral is required as some interrupt management code involved in disabling interrupt lines may not be safely executed from primary mode. By passing a valid *done* object address, the caller can wait for the request to complete, by sleeping on [rtdm\\_event\\_wait\(\)](#).

#### 6.4.3.2 udd\_enable\_irq()

```
void udd_enable_irq (
    struct udd_device * udd,
    rtdm_event_t * done )
```

Enable the device IRQ line

This service issues a request to the regular kernel for enabling the IRQ line registered by the driver. If the caller runs in primary mode, the request is scheduled but deferred until the current CPU leaves the real-time domain (see note). Otherwise, the request is immediately handled.

## Parameters

<i>udd</i>	The UDD driver handling the IRQ to disable. If no IRQ was registered by the driver at the UDD core, this routine has no effect.
<i>done</i>	Optional event to signal upon completion. If non-NULL, <i>done</i> will be posted by a call to <a href="#">rtm_event_signal()</a> after the interrupt line is enabled.

## Tags

[unrestricted](#)

## Note

The deferral is required as some interrupt management code involved in enabling interrupt lines may not be safely executed from primary mode. By passing a valid *done* object address, the caller can wait for the request to complete, by sleeping on [rtm\\_event\\_wait\(\)](#).

6.4.3.3 `udd_get_device()`

```
struct udd\_device* udd_get_device (
    struct rtm\_fd * fd )
```

RTDM file descriptor to target UDD device

Retrieves the UDD device from a RTDM file descriptor.

## Parameters

<i>fd</i>	File descriptor received by an ancillary I/O handler from a mini-driver based on the UDD core.
-----------	--

## Returns

A pointer to the UDD device to which *fd* refers to.

## Note

This service is intended for use by mini-drivers based on the UDD core exclusively. Passing file descriptors referring to other RTDM devices will certainly lead to invalid results.

## Tags

[mode-unrestricted](#)

#### 6.4.3.4 udd\_notify\_event()

```
void udd_notify_event (
    struct udd_device * udd )
```

Notify an IRQ event for an unmanaged interrupt

When the UDD core shall hand over the interrupt management for a device to the mini-driver (see `UDD_IRQ_CUSTOM`), the latter should notify the UDD core when IRQ events are received by calling this service.

As a result, the UDD core wakes up any Cobalt thread waiting for interrupts on the device via a `read(2)` or `select(2)` call.

Parameters

<code>udd</code>	UDD device descriptor receiving the IRQ.
------------------	--

Tags

`coreirq-only`

Note

In case the [IRQ handler](#) from the mini-driver requested the UDD core not to re-enable the interrupt line, the application may later request the unmasking by issuing the `UDD_RTIOC_IRQEN` `ioctl(2)` command. Writing a non-zero integer to the device via the `write(2)` system call has the same effect.

#### 6.4.3.5 udd\_register\_device()

```
int udd_register_device (
    struct udd_device * udd )
```

Register a UDD device

This routine registers a mini-driver at the UDD core.

Parameters

<code>udd</code>	<a href="#">UDD device descriptor</a> which should describe the new device properties.
------------------	--

Returns

Zero is returned upon success, otherwise a negative error code is received, from the set of error codes defined by [rtdm\\_dev\\_register\(\)](#). In addition, the following error codes can be returned:

- -EINVAL, some of the memory regions declared in the `udd_device.mem_regions[]` array have invalid properties, i.e. bad type, NULL name, zero length or address. Any undeclared region entry from the array must bear the `UDD_MEM_NONE` type.
- -EINVAL, if `udd_device.irq` is different from `UDD_IRQ_CUSTOM` and `UDD_IRQ_NONE` but invalid, causing `rtdm_irq_request()` to fail.
- -EINVAL, if `udd_device.device_flags` contains invalid flags.
- -ENOSYS, if this service is called while the real-time core is disabled.

Tags

[secondary-only](#)

References `udd_device::device_flags`, `udd_device::interrupt`, `udd_device::irq`, `RTDM_PROTOCOL_DEVICE`, `UDD_IRQ_CUSTOM`, and `UDD_IRQ_NONE`.

#### 6.4.3.6 `udd_unregister_device()`

```
int udd_unregister_device (
    struct udd_device * udd )
```

Unregister a UDD device

This routine unregisters a mini-driver from the UDD core. This routine waits until all connections to `udd` have been closed prior to unregistering.

Parameters

<code>udd</code>	UDD device descriptor
------------------	-----------------------

Returns


Zero is returned upon success, otherwise `-ENXIO` is received if this service is called while the Cobalt kernel is disabled.

Tags

[secondary-only](#)

## 6.5 Thread runtime statistics

Collaboration diagram for Thread runtime statistics:



### 6.5.1 Detailed Description

## 6.6 Thread state flags

Bits reporting permanent or transient states of threads.

Collaboration diagram for Thread state flags:



### Macros

- #define `XNSUSP` 0x00000001  
*Suspended.*
- #define `XNPEND` 0x00000002  
*Sleep-wait for a resource.*
- #define `XNDELAY` 0x00000004  
*Delayed.*
- #define `XNREADY` 0x00000008  
*Linked to the ready queue.*
- #define `XNDORMANT` 0x00000010  
*Not started yet.*
- #define `XNZOMBIE` 0x00000020  
*Zombie thread in deletion process.*
- #define `XNMAPPED` 0x00000040  
*Thread is mapped to a linux task.*
- #define `XNRELAX` 0x00000080  
*Relaxed shadow thread (blocking bit)*
- #define `XNMIGRATE` 0x00000100  
*Thread is currently migrating to another CPU.*
- #define `XNHELD` 0x00000200  
*Thread is held to process emergency.*
- #define `XNBOOST` 0x00000400  
*PI/PP boost undergoing.*
- #define `XNSSTEP` 0x00000800  
*Single-stepped by debugger.*
- #define `XNLOCK` 0x00001000  
*Scheduler lock control (pseudo-bit, not in ->state)*
- #define `XNRRB` 0x00002000  
*Undergoes a round-robin scheduling.*
- #define `XNWARN` 0x00004000  
*Issue SIGDEBUG on error detection.*
- #define `XNFPU` 0x00008000  
*Thread uses FPU.*
- #define `XNROOT` 0x00010000

- *Root thread (that is, Linux/IDLE)*  
• #define `XNWEAK` 0x00020000  
*Non real-time shadow (from the WEAK class)*
- #define `XNUSER` 0x00040000  
*Shadow thread running in userland.*
- #define `XNJOINED` 0x00080000  
*Another thread waits for joining this thread.*
- #define `XNTRAPLB` 0x00100000  
*Trap lock break (i.e.*
- #define `XNDEBUG` 0x00200000  
*User-level debugging enabled.*
- #define `XNDBGSTOP` 0x00400000  
*Stopped for synchronous debugging.*

## 6.6.1 Detailed Description

Bits reporting permanent or transient states of threads.

## 6.6.2 Macro Definition Documentation

### 6.6.2.1 XNHELD

```
#define XNHELD 0x00000200
```

Thread is held to process emergency.

### 6.6.2.2 XNMIGRATE

```
#define XNMIGRATE 0x00000100
```

Thread is currently migrating to another CPU.

### 6.6.2.3 XNPEND

```
#define XNPEND 0x00000002
```

Sleep-wait for a resource.



**6.6.2.4 XNREADY**

```
#define XNREADY 0x00000008
```

Linked to the ready queue.

**6.6.2.5 XNSUSP**

```
#define XNSUSP 0x00000001
```

Suspended.

Referenced by `xnthread_init()`.

**6.6.2.6 XNTRAPLB**

```
#define XNTRAPLB 0x00100000
```

Trap lock break (i.e.

may not sleep with sched lock)

## 6.7 Thread information flags

Bits reporting events notified to threads.

Collaboration diagram for Thread information flags:



### Macros

- #define [XNTIMEO](#) 0x00000001  
*Woken up due to a timeout condition.*
- #define [XNRMID](#) 0x00000002  
*Pending on a removed resource.*
- #define [XNBREAK](#) 0x00000004  
*Forcibly awoken from a wait state.*
- #define [XNKICKED](#) 0x00000008  
*Forced out of primary mode.*
- #define [XNWAKEN](#) 0x00000010  
*Thread waken up upon resource availability.*
- #define [XNROBBED](#) 0x00000020  
*Robbed from resource ownership.*
- #define [XNCANCELLED](#) 0x00000040  
*Cancellation request is pending.*
- #define [XNPIALERT](#) 0x00000080  
*Priority inversion alert (SIGDEBUG sent)*
- #define [XNSCHEDP](#) 0x00000100  
*schedparam propagation is pending*
- #define [XNCONTHI](#) 0x00000200  
*Continue in primary mode after debugging.*
- #define [XNMOVED](#) 0x00000001  
*CPU migration in primary mode occurred.*
- #define [XNLBALERT](#) 0x00000002  
*Scheduler lock break alert (SIGDEBUG sent)*
- #define [XNDESCENT](#) 0x00000004  
*Adaptive transitioning to secondary mode.*
- #define [XNSYSRST](#) 0x00000008  
*Thread awaiting syscall restart after signal.*
- #define [XNHICCUP](#) 0x00000010  
*Just left from ptracing.*


#### 6.7.1 Detailed Description

Bits reporting events notified to threads.

## 6.8 CAN Devices

This is the common interface a RTDM-compliant CAN device has to provide.

Collaboration diagram for CAN Devices:



### 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.*
- struct [can\\_ifreq](#)  
*CAN interface request descriptor.*

### 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_ACTIVE` = 0, `CAN_STATE_ERROR_WARNING`  
= 1, `CAN_STATE_BUS_WARNING` = 1,  
`CAN_STATE_ERROR_PASSIVE` = 2, `CAN_STATE_BUS_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

**Deprecated** Passing struct `ifreq` as a request descriptor for CAN IOCTLs is still accepted for backward compatibility, however it is recommended to switch to struct `can_ifreq` at the first opportunity.

- #define `SIOCGIFINDEX` *defined\_by\_kernel\_header\_file*
  - *Get CAN interface index by name*
- #define `SIOCSCANBAUDRATE` `_IOW(RTIOC_TYPE_CAN, 0x01, struct can_ifreq)`
  - *Set baud rate*
- #define `SIOCGCANBAUDRATE` `_IOWR(RTIOC_TYPE_CAN, 0x02, struct can_ifreq)`
  - *Get baud rate*
- #define `SIOCSCANCUSTOMBITTIME` `_IOW(RTIOC_TYPE_CAN, 0x03, struct can_ifreq)`
  - *Set custom bit time parameter*
- #define `SIOGCANCUSTOMBITTIME` `_IOWR(RTIOC_TYPE_CAN, 0x04, struct can_ifreq)`
  - *Get custom bit-time parameters*
- #define `SIOCSCANMODE` `_IOW(RTIOC_TYPE_CAN, 0x05, struct can_ifreq)`
  - *Set operation mode of CAN controller*
- #define `SIOGCANSTATE` `_IOWR(RTIOC_TYPE_CAN, 0x06, struct can_ifreq)`
  - *Get current state of CAN controller*
- #define `SIOCSCANCTRLMODE` `_IOW(RTIOC_TYPE_CAN, 0x07, struct can_ifreq)`
  - *Set special controller modes*
- #define `SIOGCANCTRLMODE` `_IOWR(RTIOC_TYPE_CAN, 0x08, struct can_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`

### 6.8.1 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

#### Tags

`secondary-only`

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.

Tags

[secondary-only](#)

Specific return values: none

### IOCTL

Tags

[task-unrestricted](#). 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 [Sockopts](#)), 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.

Tags

[secondary-only](#)

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

Tags

[task-unrestricted](#) 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 `msg_hdr` 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

Tags

[mode-unrestricted](#)

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](#).)

Tags

[mode-unrestricted](#)

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.

## 6.8.2 Macro Definition Documentation

### 6.8.2.1 CAN\_CTRLMODE\_3\_SAMPLES

```
#define CAN_CTRLMODE_3_SAMPLES 0x4
```

Triple sampling mode

In this mode the CAN controller uses Triple sampling.

### 6.8.2.2 CAN\_CTRLMODE\_LISTENONLY

```
#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](#).

### 6.8.2.3 CAN\_CTRLMODE\_LOOPBACK

```
#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](#).

### 6.8.2.4 CAN\_ERR\_LOSTARB\_UNSPEC

```
#define CAN_ERR_LOSTARB_UNSPEC 0x00
```

unspecified

else bit number in bitstream

### 6.8.2.5 CAN\_RAW\_ERR\_FILTER

```
#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

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

## Tags

[task-unrestricted](#)

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](#).

## 6.8.2.6 CAN\_RAW\_FILTER

```
#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

in	<i>level</i>	<b>SOL_CAN_RAW</b>
in	<i>optname</i>	<b>CAN_RAW_FILTER</b>
in	<i>optval</i>	Pointer to array of struct <a href="#">can_filter</a> .
in	<i>optlen</i>	Size of filter list: <code>count * sizeof( struct can_filter)</code> .

## Tags

[task-unrestricted](#)

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 length "optlen")
- -ENOSPC (No space to store filter list, check RT-Socket-CAN kernel parameters)

Examples:

[can-rtt.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

### 6.8.2.7 CAN\_RAW\_LOOPBACK

```
#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

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

Tags

[task-unrestricted](#)

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](#).



## 6.8.2.8 CAN\_RAW\_RECV\_OWN\_MSGS

```
#define CAN_RAW_RECV_OWN_MSGS 0x4
```

CAN receive own messages.

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

## 6.8.2.9 RTCAN\_RTIOC\_RCV\_TIMEOUT

```
#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

in	<i>arg</i>	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.

## Tags

[task-unrestricted](#)

## Examples:

[rtcanrecv.c](#).

## 6.8.2.10 RTCAN\_RTIOC\_SND\_TIMEOUT

```
#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

in	<i>arg</i>	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.

## Tags

[task-unrestricted](#)

## Examples:

[rtcansend.c](#).

## 6.8.2.11 RTCAN\_RTIOC\_TAKE\_TIMESTAMP

```
#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

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

## Returns

0 on success.

## Tags

[task-unrestricted](#)

## 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.

## Examples:

[rtcanrecv.c](#).

## 6.8.2.12 SIOCGCANBAUDRATE

```
#define SIOCGCANBAUDRATE _IOWR(RTIIOC_TYPE_CAN, 0x02, struct can_ifreq)
```

## Get baud rate

## Parameters

in,out	<i>arg</i>	Pointer to interface request structure buffer (struct <a href="#">can_ifreq</a> ). <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.

## Tags

[task-unrestricted](#)

## 6.8.2.13 SIOCGCANCTRLMODE

```
#define SIOCGCANCTRLMODE _IOWR(RTIIOC_TYPE_CAN, 0x08, struct can_ifreq)
```

## Get special controller modes

## Parameters

in	<i>arg</i>	Pointer to interface request structure buffer (struct <a href="#">can_ifreq</a> ). <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.

## Tags

[task-unrestricted](#), [might-switch](#)

## 6.8.2.14 SIOCGCANCUSTOMBITTIME

```
#define SIOCGCANCUSTOMBITTIME _IOWR(RTIIOC_TYPE_CAN, 0x04, struct can_ifreq)
```

Get custom bit-time parameters

## Parameters

in,out	arg	
		Pointer to interface request structure buffer (struct <a href="#">can_ifreq</a> ). <code>ifr_name</code> must hold a valid CAN interface name, <code>ifr_ifru</code> 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.

## Tags

[task-unrestricted](#)

## 6.8.2.15 SIOCGCANSTATE

```
#define SIOCGCANSTATE _IOWR(RTIIOC_TYPE_CAN, 0x06, struct can_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 <a href="#">can_ifreq</a> ). <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_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.

## Tags

[task-unrestricted](#), [might-switch](#)

## 6.8.2.16 SIOCGIFINDEX

```
#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 <a href="#">can_ifreq</a> ). If <code>ifr_name</code> holds a valid CAN interface name <code>ifr_ifindex</code> 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.

## Tags

[task-unrestricted](#)

## Examples:

[can-rtt.c](#), [rtcanconfig.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

## 6.8.2.17 SIOCSCANBAUDRATE

```
#define SIOCSCANBAUDRATE _IOW(RTIOC_TYPE_CAN, 0x01, struct can_ifreq)
```

## Set baud rate

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

## Parameters

in	arg	Pointer to interface request structure buffer (struct <a href="#">can_ifreq</a> ). <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_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 successully fulfilled. Try again.

## Tags

[task-unrestricted](#), [might-switch](#)

## Note

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

## Examples:

[rtcanconfig.c](#).

## 6.8.2.18 SIOCSCANCTRLMODE

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

## Set special controller modes

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

## Parameters

in	<i>arg</i>	Pointer to interface request structure buffer (struct <a href="#">can_ifreq</a> ). <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 valid baud rate, see [can\\_baudrate\\_t](#).
- -EAGAIN: Request could not be successfully fulfilled. Try again.

## Tags

[task-unrestricted](#), [might-switch](#)

## Note

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

## Examples:

[rtcanconfig.c](#).

## 6.8.2.19 SIOCSCANCUSTOMBITTIME

```
#define SIOCSCANCUSTOMBITTIME _IOW(RTIIOC_TYPE_CAN, 0x03, struct can_ifreq)
```

## Set custom bit time parameter

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

## Parameters

in	<i>arg</i>	Pointer to interface request structure buffer (struct <a href="#">can_ifreq</a> ). <code>ifr_name</code> must hold a valid CAN interface name, <code>ifr_ifru</code> 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.

## Tags

[task-unrestricted](#), [might-switch](#)

## Note

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

## Examples:

[rtcanconfig.c](#).

## 6.8.2.20 SIOCSCANMODE

```
#define SIOCSCANMODE _IOW(RTIIOC_TYPE_CAN, 0x05, struct can_ifreq)
```

Set operation mode of CAN controller

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

## Parameters

in	arg	Pointer to interface request structure buffer (struct <a href="#">can_ifreq</a> ). <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_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

## Tags

[task-unrestricted](#), [might-switch](#)



## 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.

## Examples:

[rtcanconfig.c](#).

## 6.8.2.21 SOL\_CAN\_RAW

```
#define SOL_CAN_RAW 103
```

CAN socket levels.

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

## Examples:

[can-rtt.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

## 6.8.3 Typedef Documentation

## 6.8.3.1 can\_filter\_t

```
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.

6.8.3.2 `can_frame_t`

```
typedef struct can\_frame can\_frame\_t
```

Raw CAN frame.

Central structure for receiving and sending CAN frames.

Examples:

[rtcanrecv.c](#).

## 6.8.4 Enumeration Type Documentation

6.8.4.1 `CAN_BITTIME_TYPE`

```
enum CAN\_BITTIME\_TYPE
```

Supported CAN bit-time types.

Enumerator

<code>CAN_BITTIME_STD</code>	Standard bit-time definition according to Bosch.
<code>CAN_BITTIME_BTR</code>	Hardware-specific BTR bit-time definition.

6.8.4.2 `CAN_MODE`

```
enum CAN\_MODE
```

Enumerator

<code>CAN_MODE_STOP</code>	Set controller in Stop mode (no reception / transmission possible)
<code>CAN_MODE_START</code>	Set controller into normal operation. Coming from stopped mode or bus off, the controller begins with no errors in <a href="#">CAN_STATE_ACTIVE</a> .
<code>CAN_MODE_SLEEP</code>	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 <code>CAN_MODE_START</code> .

## 6.8.4.3 CAN\_STATE

enum [CAN\\_STATE](#)


Enumerator

CAN_STATE_ERROR_ACTIVE	CAN controller is error active.
CAN_STATE_ACTIVE	CAN controller is active.
CAN_STATE_ERROR_WARNING	CAN controller is error active, warning level is reached.
CAN_STATE_BUS_WARNING	CAN controller is error active, warning level is reached.
CAN_STATE_ERROR_PASSIVE	CAN controller is error passive.
CAN_STATE_BUS_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.

## 6.9 RTDM

The Real-Time Driver Model (RTDM) provides a unified interface to both users and developers of real-time device drivers.

Collaboration diagram for RTDM:



### Modules

- [RTDM User API](#)  
*Application interface to RTDM services.*
- [Driver programming interface](#)  
*RTDM driver programming interface.*
- [Device Profiles](#)  
*Pre-defined classes of real-time devices.*

### 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](#) 9  
*Common user and driver API version.*
- #define [RTDM\\_API\\_MIN\\_COMPAT\\_VER](#) 9  
*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.*

### 6.9.1 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**

### 6.9.2 Macro Definition Documentation

#### 6.9.2.1 RTDM\_TIMEOUT\_INFINITE

```
#define RTDM_TIMEOUT_INFINITE 0
```

Block forever.

#### 6.9.2.2 RTDM\_TIMEOUT\_NONE

```
#define RTDM_TIMEOUT_NONE (-1)
```

Any negative timeout means non-blocking.

### 6.9.3 Typedef Documentation

### 6.9.3.1 nanosecs\_abs\_t

```
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](#).

### 6.9.3.2 nanosecs\_rel\_t

```
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](#).

## 6.10 RTDM User API

Application interface to RTDM services.

Collaboration diagram for RTDM User API:



### Files

- file [rtdm.h](#)

*Real-Time Driver Model for Xenomai, user API header.*

### 6.10.1 Detailed Description

Application interface to RTDM services.

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.

## 6.11 Serial Devices

This is the common interface a RTDM-compliant serial device has to provide.

Collaboration diagram for Serial Devices:



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 Class:** RTDM\_CLASS\_SERIAL

**Device Name:** "/dev/rtdm/rtser<N>", N >= 0

### Supported Operations

**Open**

### Tags

**secondary-only** Specific return values: none

**Close**



Tags

[secondary-only](#) Specific return values: none

## IOCTL

Tags

[task-unrestricted](#). See [below](#)

Specific return values: see [below](#)

## Read

Tags

[mode-unrestricted](#) 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

Tags

[mode-unrestricted](#) 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)

## 6.12 Testing Devices

This group of devices is intended to provide in-kernel testing results.

Collaboration diagram for Testing Devices:



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 Class:** RTDM\_CLASS\_TESTING

### Supported Operations

#### Open

#### Tags

**secondary-only** Specific return values: none

#### Close

#### Tags

**secondary-only** Specific return values: none

#### IOCTL


#### Tags

**task-unrestricted.** See **TSTIOCTLs** below  
Specific return values: see **TSTIOCTLs** below

## 6.13 Real-time IPC

### Profile Revision: 1

Collaboration diagram for Real-time IPC:



### 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 msghdr \*msg, int flags)  
*Send a message on a RTIPC socket.*
- ssize\_t [recvmsg\\_\\_AF\\_RTIPC](#) (int sockfd, struct 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_SNDBUFSZ` defined\_by\_kernel\_header\_file  
*IPCPROTO\_IDDP and IPCPROTO\_BUFP protocols support the standard SO\_SNDBUFSZ socket option, from the SOL\_SOCKET level.*
- #define `SO_RCVBUFSZ` defined\_by\_kernel\_header\_file  
*All RTIPC protocols support the standard SO\_RCVBUFSZ socket option, from the SOL\_SOCKET level.*

### 6.13.1 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

### 6.13.2 Macro Definition Documentation

#### 6.13.2.1 BUFP\_BUFSZ

```
#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	<i>level</i>	<a href="#">SOL_BUF</a>
in	<i>optname</i>	<b>BUFP_BUFSZ</b>
in	<i>optval</i>	Pointer to a variable of type <code>size_t</code> , containing the required size of the buffer to reserve at binding time
in	<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* is invalid or *\*optval* is zero)

## Calling context:

RT/non-RT

## Examples:

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

## 6.13.2.2 BUFP\_LABEL

```
#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_BUF</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>	<code>sizeof(struct rtipc_port_label)</code>

## 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](#).

## 6.13.2.3 IDDP\_LABEL

```
#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

in	<i>level</i>	<a href="#">SOL_IDDP</a>
in	<i>optname</i>	<b>IDDP_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:

[iddp-label.c](#).

#### 6.13.2.4 IDDP\_POOLSZ

```
#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

in	<i>level</i>	<a href="#">SOL_IDDP</a>
in	<i>optname</i>	<b>IDDP_POOLSZ</b>
in	<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
in	<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* is invalid or *\*optval* is zero)

Calling context:

RT/non-RT

Examples:

[iddp-sendrecv.c](#).



### 6.13.2.5 SO\_RCVTIMEO

```
#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](#).

### 6.13.2.6 SO\_SNDTIMEO

```
#define SO_SNDTIMEO defined_by_kernel_header_file
```

IPPROTO\_IPDD 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/>

### 6.13.2.7 XDDP\_BUFSZ

```
#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](#).

## 6.13.2.8 XDDP\_EVTDOWN

```
#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.

## 6.13.2.9 XDDP\_EVTIN

```
#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.

## 6.13.2.10 XDDP\_EVTNOBUF

```
#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.

## 6.13.2.11 XDDP\_EVTOUT

```
#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.

## 6.13.2.12 XDDP\_LABEL

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

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](#).

### 6.13.2.13 XDDP\_MONITOR

```
#define XDDP_MONITOR 4
```

XDDP monitoring callback.

Other RTDM drivers may install a user-defined callback via the `rtdm_setssockopt` 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	<i>level</i>	<a href="#">SOL_XDDP</a>
in	<i>optname</i>	<b>XDDP_MONITOR</b>
in	<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.
in	<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

## 6.13.2.14 XDDP\_POOLSZ

```
#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

in	<i>level</i>	<a href="#">SOL_XDDP</a>
in	<i>optname</i>	<b>XDDP_POOLSZ</b>
in	<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
in	<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](#).

## 6.13.3 Enumeration Type Documentation

### 6.13.3.1 anonymous enum

anonymous enum

## Enumerator

IPCPROTO_IPC	Default protocol (IDDP)
IPCPROTO_XDDP	<p>Cross-domain datagram protocol (RT &lt;-&gt; 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.</p> <p>On the Linux domain side, pseudo-device files named <code>/dev/rtp&lt;minor&gt;</code> 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 <code>/dev/rtp7</code>.</p> <p>All data sent through a bound/connected XDDP socket via <code>sendto(2)</code> or <code>write(2)</code> will be passed to the peer endpoint in the Linux domain, and made available for reading via the standard <code>read(2)</code> system call. Conversely, all data sent using <code>write(2)</code> through the non real-time endpoint will be conveyed to the real-time socket endpoint, and made available to the <code>recvfrom(2)</code> or <code>read(2)</code> system calls.</p>
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_BUFP	<p>Buffer protocol (RT &lt;-&gt; 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.</p> <p>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).</p>

## 6.13.4 Function Documentation

6.13.4.1 `bind__AF_RTIPC()`

```
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

in	<code>sockfd</code>	The RTDM file descriptor obtained from the socket creation call.
in	<code>addr</code>	The address to bind the socket to (see struct <code>sockaddr_ipc</code> ). 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/rtpc/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/rtpc/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 [BUFPP\\_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/rtpc/bufpp/*label*, where *label* is the label string passed to setsockopt() for the [BUFPP\\_LABEL](#) option.

#### Parameters

in	<i>addrlen</i>	The size in bytes of the structure pointed to by <i>addr</i> .
----	----------------	--

#### 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)

- -EAGAIN (no registry slot available, check/raise CONFIG\_XENO\_OPT\_REGISTRY\_NRSL↔OTS) .

Calling context:

non-RT

#### 6.13.4.2 close\_\_AF\_RTIPC()

```
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.

Parameters

in	<i>sockfd</i>	The socket descriptor to close.
----	---------------	---------------------------------

Returns

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

Calling context:

non-RT

#### 6.13.4.3 connect\_\_AF\_RTIPC()

```
int connect__AF_RTIPC (
    int sockfd,
    const struct sockaddr\_ipc * addr,
    socklen_t addrlen )
```

Initiate a connection on a RTIPC socket.

Parameters

in	<i>sockfd</i>	The RTDM file descriptor obtained from the socket creation call.
in	<i>addr</i>	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)`.

## Parameters

in	<i>addrlen</i>	The size in bytes of the structure pointed to by <i>addr</i> .
----	----------------	--

## 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

6.13.4.4 `getpeername__AF_RTIPC()`

```
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

#### 6.13.4.5 getsockname\_\_AF\_RTIPC()

```
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

#### 6.13.4.6 getsockopt\_\_AF\_RTIPC()

```
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\\_BUFPP](#)

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

#### 6.13.4.7 recvmsg\_\_AF\_RTIPC()

```
ssize_t recvmsg__AF_RTIPC (
    int sockfd,
    struct msghdr * msg,
    int flags )
```

Receive a message from a RTIPC socket.

## Parameters

in	<i>sockfd</i>	The RTDM file descriptor obtained from the socket creation call.
out	<i>msg</i>	The address the message header will be copied at.
in	<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

**IPPROTO\_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

6.13.4.8 `sendmsg__AF_RTIPC()`

```
ssize_t sendmsg__AF_RTIPC (
    int sockfd,
    const struct msghdr * msg,
    int flags )
```

Send a message on a RTIPC socket.

## Parameters

in	<i>sockfd</i>	The RTDM file descriptor obtained from the socket creation call.
in	<i>msg</i>	The address of the message header conveying the datagram.
in	<i>flags</i>	Operation flags:

- **MSG\_OOB** Send out-of-band message. For all RTIPC protocols except **IPPROTO\_BUF**, 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. **IPPROTO\_BUF** 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

#### 6.13.4.9 `setsockopt__AF_RTIPC()`

```
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

6.13.4.10 `socket__AF_RTIPC()`

```
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

in	<i>domain</i>	The communication domain. Must be AF_RTIPC.
in	<i>type</i>	The socket type. Must be SOCK_DGRAM.
in	<i>protocol</i>	Any of <a href="#">IPCPROTO_XDDP</a> , <a href="#">IPCPROTO_IDDP</a> , or <a href="#">IPCPROTO_BUFP</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

## 6.14 Asynchronous Procedure Calls

Services for scheduling function calls in the Linux domain.

Collaboration diagram for Asynchronous Procedure Calls:



### Functions

- int `xnipc_alloc` (const char \*name, void(\*handler)(void \*cookie), void \*cookie)

*Allocate an APC slot.*

- void `xnipc_free` (int apc)

*Releases an APC slot.*

- static void `xnipc_schedule` (int apc)

*Schedule an APC invocation.*

### 6.14.1 Detailed Description

Services for scheduling function calls in the Linux domain.

APC is the acronym for Asynchronous Procedure Call, a mean by which activities from the Xenomai domain can schedule deferred invocations of handlers to be run into the Linux domain, as soon as possible when the Linux kernel gets back in control.

Up to `BITS_PER_LONG` APC slots can be active at any point in time.

APC support is built upon the interrupt pipeline's virtual interrupt support.

### 6.14.2 Function Documentation

6.14.2.1 `xnapc_alloc()`

```
int xnapc_alloc (
    const char * name,
    void(*)(void *cookie) handler,
    void * cookie )
```

Allocate an APC slot.

APC is the acronym for Asynchronous Procedure Call, a mean by which activities from the Xenomai domain can schedule deferred invocations of handlers to be run into the Linux domain, as soon as possible when the Linux kernel gets back in control. Up to `BITS_PER_LONG` APC slots can be active at any point in time. APC support is built upon the interrupt pipeline's virtual interrupt support.

Any Linux kernel service which is callable from a regular Linux interrupt handler is in essence available to APC handlers.

## Parameters

<i>name</i>	is a symbolic name identifying the APC which will get reported through the <code>/proc/xenomai/apc</code> interface. Passing <code>NULL</code> to create an anonymous APC is allowed.
<i>handler</i>	The address of the fault handler to call upon exception condition. The handle will be passed the <i>cookie</i> value unmodified.
<i>cookie</i>	A user-defined opaque pointer the APC handler receives as its sole argument.

## Returns

a valid APC identifier is returned upon success, or a negative error code otherwise:

- `-EINVAL` is returned if *handler* is invalid.
- `-EBUSY` is returned if no more APC slots are available.

## Tags

[unrestricted](#)

6.14.2.2 `xnapc_free()`

```
void xnapc_free (
    int apc )
```

Releases an APC slot.

This service deallocates an APC slot obtained by `xnapc_alloc()`.

## Parameters

<i>apc</i>	The APC id. to release, as returned by a successful call to the <a href="#">xnapc_alloc()</a> service.
------------	--

## Tags

[unrestricted](#)

## 6.14.2.3 xnapc\_schedule()

```
static inline int xnapc_schedule (  
    int apc ) [inline], [static]
```

Schedule an APC invocation.

This service marks the APC as pending for the Linux domain, so that its handler will be called as soon as possible, when the Linux domain gets back in control.

When posted from the Linux domain, the APC handler is fired as soon as the interrupt mask is explicitly cleared by some kernel code. When posted from the Xenomai domain, the APC handler is fired as soon as the Linux domain is resumed, i.e. after Xenomai has completed all its pending duties.

## Parameters

<i>apc</i>	The APC id. to schedule.
------------	--------------------------

This service can be called from:

- Any domain context, albeit the usual calling place is from the Xenomai domain.



## 6.15 In-kernel arithmetics

A collection of helpers performing arithmetics not implicitly available from kernel context via GCC helpers.

Collaboration diagram for In-kernel arithmetics:



### Functions

- unsigned long long [xnarch\\_generic\\_full\\_divmod64](#) (unsigned long long a, unsigned long long b, unsigned long long \*rem)

*Architecture-independent div64 operation with remainder.*

#### 6.15.1 Detailed Description

A collection of helpers performing arithmetics not implicitly available from kernel context via GCC helpers.

Many of these routines enable 64bit arithmetics on 32bit systems. Xenomai architecture ports normally implement the performance critical ones in hand-crafted assembly code (see kernel/cobalt/arch/<arch>/include/asm/xenomai/uapi/arith.h).

#### 6.15.2 Function Documentation

##### 6.15.2.1 xnarch\_generic\_full\_divmod64()

```

unsigned long long xnarch_generic_full_divmod64 (
    unsigned long long a,
    unsigned long long b,
    unsigned long long * rem )
  
```

Architecture-independent div64 operation with remainder.


##### Parameters

<i>a</i>	dividend
<i>b</i>	divisor
<i>rem</i>	if non-NULL, a pointer to a 64bit variable for collecting the remainder from the division.

## 6.16 Buffer descriptor

Abstraction for copying data to/from different address spaces.

Collaboration diagram for Buffer descriptor:



### Functions

- static void [xnbufd\\_map\\_uread](#) (struct xnbufd \*bufd, const void \_\_user \*ptr, size\_t len)

*Initialize a buffer descriptor for reading from user memory.*

- static void [xnbufd\\_map\\_uwrite](#) (struct xnbufd \*bufd, void \_\_user \*ptr, size\_t len)

*Initialize a buffer descriptor for writing to user memory.*

- ssize\_t [xnbufd\\_unmap\\_uread](#) (struct xnbufd \*bufd)

*Finalize a buffer descriptor obtained from [xnbufd\\_map\\_uread\(\)](#).*

- ssize\_t [xnbufd\\_unmap\\_uwrite](#) (struct xnbufd \*bufd)

*Finalize a buffer descriptor obtained from [xnbufd\\_map\\_uwrite\(\)](#).*

- static void [xnbufd\\_map\\_kread](#) (struct xnbufd \*bufd, const void \*ptr, size\_t len)

*Initialize a buffer descriptor for reading from kernel memory.*

- static void [xnbufd\\_map\\_kwrite](#) (struct xnbufd \*bufd, void \*ptr, size\_t len)

*Initialize a buffer descriptor for writing to kernel memory.*

- ssize\_t [xnbufd\\_unmap\\_kread](#) (struct xnbufd \*bufd)

*Finalize a buffer descriptor obtained from [xnbufd\\_map\\_kread\(\)](#).*

- ssize\_t [xnbufd\\_unmap\\_kwrite](#) (struct xnbufd \*bufd)

*Finalize a buffer descriptor obtained from [xnbufd\\_map\\_kwrite\(\)](#).*

- ssize\_t [xnbufd\\_copy\\_to\\_kmem](#) (void \*ptr, struct xnbufd \*bufd, size\_t len)

*Copy memory covered by a buffer descriptor to kernel memory.*

- ssize\_t [xnbufd\\_copy\\_from\\_kmem](#) (struct xnbufd \*bufd, void \*from, size\_t len)

*Copy kernel memory to the area covered by a buffer descriptor.*

- void [xnbufd\\_invalidate](#) (struct xnbufd \*bufd)

*Invalidate a buffer descriptor.*

- static void [xnbufd\\_reset](#) (struct xnbufd \*bufd)

*Reset a buffer descriptor.*

### 6.16.1 Detailed Description

Abstraction for copying data to/from different address spaces.

A buffer descriptor is a simple abstraction dealing with copy operations to/from memory buffers which may belong to different address spaces.

To this end, the buffer descriptor library provides a small set of copy routines which are aware of address space restrictions when moving data, and a generic container type which can hold a reference to - or cover - a particular memory area, either present in kernel space, or in any of the existing user memory contexts.

The goal of the buffer descriptor abstraction is to hide address space specifics from Xenomai services dealing with memory areas, allowing them to operate on multiple address spaces seamlessly.

The common usage patterns are as follows:

- Implementing a Xenomai syscall returning a bulk of data to the caller, which may have to be copied back to either kernel or user space:

```
[Syscall implementation]
ssize_t rt_bulk_read_inner(struct xnbuid *bufd)
{
    ssize_t ret;
    size_t len;
    void *bulk;

    bulk = get_next_readable_bulk(&len);
    ret = xnbuid_copy_from_kmem(bufd, bulk, min(bufd->b_len, len));
    free_bulk(bulk);

    ret = this_may_fail();
    if (ret)
        xnbuid_invalidate(bufd);

    return ret;
}

[Kernel wrapper for in-kernel calls]
int rt_bulk_read(void *ptr, size_t len)
{
    struct xnbuid bufd;
    ssize_t ret;

    xnbuid_map_kwrite(&bufd, ptr, len);
    ret = rt_bulk_read_inner(&bufd);
    xnbuid_unmap_kwrite(&bufd);

    return ret;
}

[Userland trampoline for user syscalls]
int __rt_bulk_read(struct pt_regs *regs)
{
    struct xnbuid bufd;
    void __user *ptr;
    ssize_t ret;
    size_t len;

    ptr = (void __user *)__xn_reg_arg1(regs);
    len = __xn_reg_arg2(regs);

    xnbuid_map_uwrite(&bufd, ptr, len);
    ret = rt_bulk_read_inner(&bufd);
    xnbuid_unmap_uwrite(&bufd);

    return ret;
}
```

- Implementing a Xenomai syscall receiving a bulk of data from the caller, which may have to be read from either kernel or user space:

```
[Syscall implementation]
ssize_t rt_bulk_write_inner(struct xbufd *bufd)
{
    void *bulk = get_free_bulk(bufd->b_len);
    return xbufd_copy_to_kmem(bulk, bufd, bufd->b_len);
}

[Kernel wrapper for in-kernel calls]
int rt_bulk_write(const void *ptr, size_t len)
{
    struct xbufd bufd;
    ssize_t ret;

    xbufd_map_kread(&bufd, ptr, len);
    ret = rt_bulk_write_inner(&bufd);
    xbufd_unmap_kread(&bufd);

    return ret;
}

[Userland trampoline for user syscalls]
int __rt_bulk_write(struct pt_regs *regs)
{
    struct xbufd bufd;
    void __user *ptr;
    ssize_t ret;
    size_t len;

    ptr = (void __user *)__xn_reg_arg1(regs);
    len = __xn_reg_arg2(regs);

    xbufd_map_uread(&bufd, ptr, len);
    ret = rt_bulk_write_inner(&bufd);
    xbufd_unmap_uread(&bufd);

    return ret;
}
```

## 6.16.2 Function Documentation

### 6.16.2.1 xbufd\_copy\_from\_kmem()

```
ssize_t xbufd_copy_from_kmem (
    struct xbufd * bufd,
    void * from,
    size_t len )
```

Copy kernel memory to the area covered by a buffer descriptor.

This routine copies *len* bytes from the kernel memory starting at *from* to the area referred to by the buffer descriptor *bufd*. `xbufd_copy_from_kmem()` tracks the write offset within the destination memory internally, so that it may be called several times in a loop, until the entire memory area is stored.

The destination address space is dealt with, according to the following rules:

- if *bufd* refers to a writable kernel area (i.e. see `xbufd_map_kwrite()`), the copy is immediately and fully performed with no restriction.
- if *bufd* refers to a writable user area (i.e. see `xbufd_map_uwrite()`), the copy is performed only if that area lives in the currently active address space, and only if the caller may sleep Linux-wise to process any potential page fault which may arise while writing to that memory.
- if *bufd* refers to a user area which may not be immediately written to from the current context, the copy is postponed until `xbufd_unmap_uwrite()` is invoked for *ubufd*, at which point the copy will take place. In such a case, the source memory is transferred to a carry over buffer allocated internally; this operation may lead to request dynamic memory from the nucleus heap if *len* is greater than 64 bytes.

## Parameters

<i>bufd</i>	The address of the buffer descriptor covering the user memory to copy data to.
<i>from</i>	The start address of the kernel memory to copy from.
<i>len</i>	The length of the kernel memory to copy to <i>bufd</i> .

## Returns

The number of bytes written so far to the memory area covered by *ubufd*. Otherwise,

- -ENOMEM is returned when no memory is available from the nucleus heap to allocate the carry over buffer.

## Tags

[unrestricted](#)

## Note

Calling this routine while holding the nklock and/or running with interrupts disabled is invalid, and doing so will trigger a debug assertion.

This routine may switch the caller to secondary mode if a page fault occurs while reading from the user area. For that reason, [xnbufd\\_copy\\_to\\_kmem\(\)](#) may only be called from a preemptible section (Linux-wise).

## 6.16.2.2 xnbufd\_copy\_to\_kmem()

```
ssize_t xnbufd_copy_to_kmem (  
    void * to,  
    struct xnbufd * bufd,  
    size_t len )
```

Copy memory covered by a buffer descriptor to kernel memory.

This routine copies *len* bytes from the area referred to by the buffer descriptor *bufd* to the kernel memory area *to*. [xnbufd\\_copy\\_to\\_kmem\(\)](#) tracks the read offset within the source memory internally, so that it may be called several times in a loop, until the entire memory area is loaded.

The source address space is dealt with, according to the following rules:

- if *bufd* refers to readable kernel area (i.e. see [xnbufd\\_map\\_kread\(\)](#)), the copy is immediately and fully performed with no restriction.
- if *bufd* refers to a readable user area (i.e. see [xnbufd\\_map\\_uread\(\)](#)), the copy is performed only if that area lives in the currently active address space, and only if the caller may sleep Linux-wise to process any potential page fault which may arise while reading from that memory.
- any attempt to read from *bufd* from a non-suitable context is considered as a bug, and will raise a panic assertion when the nucleus is compiled in debug mode.

## Parameters

<i>to</i>	The start address of the kernel memory to copy to.
<i>bufd</i>	The address of the buffer descriptor covering the user memory to copy data from.
<i>len</i>	The length of the user memory to copy from <i>bufd</i> .

## Returns

The number of bytes read so far from the memory area covered by *ubufd*. Otherwise:

- -EINVAL is returned upon attempt to read from the user area from an invalid context. This error is only returned when the debug mode is disabled; otherwise a panic assertion is raised.

## Tags

[task-unrestricted](#)

## Note

Calling this routine while holding the nklock and/or running with interrupts disabled is invalid, and doing so will trigger a debug assertion.

This routine may switch the caller to secondary mode if a page fault occurs while reading from the user area. For that reason, [xnbufd\\_copy\\_to\\_kmem\(\)](#) may only be called from a preemptible section (Linux-wise).

## 6.16.2.3 xnbufd\_invalidate()

```
void xnbufd_invalidate (
    struct xnbufd * bufd )
```

Invalidate a buffer descriptor.

The buffer descriptor is invalidated, making it unusable for further copy operations. If an outstanding carry over buffer was allocated by a previous call to [xnbufd\\_copy\\_from\\_kmem\(\)](#), it is immediately freed so that no data transfer will happen when the descriptor is finalized.

The only action that may subsequently be performed on an invalidated descriptor is calling the relevant unmapping routine for it. For that reason, [xnbufd\\_invalidate\(\)](#) should be invoked on the error path when data may have been transferred to the carry over buffer.

## Parameters

<i>bufd</i>	The address of the buffer descriptor to invalidate.
-------------	---

## Tags

[unrestricted](#)

6.16.2.4 `xnbufd_map_kread()`

```
void xnbufd_map_kread (
    struct xnbufd * bufd,
    const void * ptr,
    size_t len ) [inline], [static]
```

Initialize a buffer descriptor for reading from kernel memory.

The new buffer descriptor may be used to copy data from kernel memory. This routine should be used in pair with [xnbufd\\_unmap\\_kread\(\)](#).

## Parameters

<i>bufd</i>	The address of the buffer descriptor which will map a <i>len</i> bytes kernel memory area, starting from <i>ptr</i> .
<i>ptr</i>	The start of the kernel buffer to map.
<i>len</i>	The length of the kernel buffer starting at <i>ptr</i> .

## Tags

[unrestricted](#)

6.16.2.5 `xnbufd_map_kwrite()`

```
void xnbufd_map_kwrite (
    struct xnbufd * bufd,
    void * ptr,
    size_t len ) [inline], [static]
```

Initialize a buffer descriptor for writing to kernel memory.

The new buffer descriptor may be used to copy data to kernel memory. This routine should be used in pair with [xnbufd\\_unmap\\_kwrite\(\)](#).

## Parameters

<i>bufd</i>	The address of the buffer descriptor which will map a <i>len</i> bytes kernel memory area, starting from <i>ptr</i> .
<i>ptr</i>	The start of the kernel buffer to map.
<i>len</i>	The length of the kernel buffer starting at <i>ptr</i> .

Tags

[unrestricted](#)

#### 6.16.2.6 `xnbufd_map_uread()`

```
void xnbufd_map_uread (
    struct xnbufd * bufd,
    const void __user * ptr,
    size_t len ) [inline], [static]
```

Initialize a buffer descriptor for reading from user memory.

The new buffer descriptor may be used to copy data from user memory. This routine should be used in pair with [xnbufd\\_unmap\\_uread\(\)](#).

Parameters

<i>bufd</i>	The address of the buffer descriptor which will map a <i>len</i> bytes user memory area, starting from <i>ptr</i> . <i>ptr</i> is never dereferenced directly, since it may refer to a buffer that lives in another address space.
<i>ptr</i>	The start of the user buffer to map.
<i>len</i>	The length of the user buffer starting at <i>ptr</i> .

Tags

[task-unrestricted](#)

#### 6.16.2.7 `xnbufd_map_uwrite()`

```
void xnbufd_map_uwrite (
    struct xnbufd * bufd,
    void __user * ptr,
    size_t len ) [inline], [static]
```

Initialize a buffer descriptor for writing to user memory.

The new buffer descriptor may be used to copy data to user memory. This routine should be used in pair with [xnbufd\\_unmap\\_uwrite\(\)](#).

Parameters

<i>bufd</i>	The address of the buffer descriptor which will map a <i>len</i> bytes user memory area, starting from <i>ptr</i> . <i>ptr</i> is never dereferenced directly, since it may refer to a buffer that lives in another address space.
<i>ptr</i>	The start of the user buffer to map.
<i>len</i>	The length of the user buffer starting at <i>ptr</i> .



Tags

[task-unrestricted](#)

#### 6.16.2.8 `xnbufd_reset()`

```
void xnbufd_reset (
    struct xnbufd * bufd ) [inline], [static]
```

Reset a buffer descriptor.

The buffer descriptor is reset, so that all data already copied is forgotten. Any carry over buffer allocated is kept, though.

Parameters

<i>bufd</i>	The address of the buffer descriptor to reset.
-------------	--

Tags

[unrestricted](#)

#### 6.16.2.9 `xnbufd_unmap_kread()`

```
ssize_t xnbufd_unmap_kread (
    struct xnbufd * bufd )
```

Finalize a buffer descriptor obtained from [xnbufd\\_map\\_kread\(\)](#).

This routine finalizes a buffer descriptor previously initialized by a call to [xnbufd\\_map\\_kread\(\)](#), to read data from a kernel area.

Parameters

<i>bufd</i>	The address of the buffer descriptor to finalize.
-------------	---

Returns

The number of bytes read so far from the memory area covered by *ubufd*.

Tags

[task-unrestricted](#)

6.16.2.10 `xnbufd_unmap_kwrite()`

```
ssize_t xnbufd_unmap_kwrite (  
    struct xnbufd * bufd )
```

Finalize a buffer descriptor obtained from `xnbufd_map_kwrite()`.

This routine finalizes a buffer descriptor previously initialized by a call to `xnbufd_map_kwrite()`, to write data to a kernel area.

## Parameters

<i>bufd</i>	The address of the buffer descriptor to finalize.
-------------	---

## Returns

The number of bytes written so far to the memory area covered by *ubufd*.

## Tags

[task-unrestricted](#)

6.16.2.11 `xnbufd_unmap_uread()`

```
ssize_t xnbufd_unmap_uread (  
    struct xnbufd * bufd )
```

Finalize a buffer descriptor obtained from `xnbufd_map_uread()`.

This routine finalizes a buffer descriptor previously initialized by a call to `xnbufd_map_uread()`, to read data from a user area.

## Parameters

<i>bufd</i>	The address of the buffer descriptor to finalize.
-------------	---

## Returns

The number of bytes read so far from the memory area covered by *ubufd*.

## Tags

[task-unrestricted](#)

## Note

Calling this routine while holding the `nklock` and/or running with interrupts disabled is invalid, and doing so will trigger a debug assertion.

6.16.2.12 `xnbufd_unmap_uwrite()`

```
ssize_t xnbufd_unmap_uwrite (  
    struct xnbufd * bufd )
```

Finalize a buffer descriptor obtained from `xnbufd_map_uwrite()`.

This routine finalizes a buffer descriptor previously initialized by a call to `xnbufd_map_uwrite()`, to write data to a user area.

The main action taken is to write the contents of the kernel memory area passed to `xnbufd_copy_from_kmem()` whenever the copy operation was postponed at that time; the carry over buffer is eventually released as needed. If `xnbufd_copy_from_kmem()` was allowed to copy to the destination user memory at once, then `xnbufd_unmap_uwrite()` leads to a no-op.

## Parameters

<i>bufd</i>	The address of the buffer descriptor to finalize.
-------------	---

## Returns

The number of bytes written so far to the memory area covered by *ubufd*.

## Tags

[task-unrestricted](#)

## Note

Calling this routine while holding the `nklock` and/or running with interrupts disabled is invalid, and doing so will trigger a debug assertion.

## 6.17 Clock services

Collaboration diagram for Clock services:



### Functions

- int [xnclock\\_register](#) (struct xnclock \*clock, const cpumask\_t \*affinity)

*Register a Xenomai clock.*

- void [xnclock\\_deregister](#) (struct xnclock \*clock)

*Deregister a Xenomai clock.*

- void [xnclock\\_tick](#) (struct xnclock \*clock)

*Process a clock tick.*

- void [xnclock\\_adjust](#) (struct xnclock \*clock, xnsticks\_t delta)

*Adjust a clock time.*

### 6.17.1 Detailed Description

### 6.17.2 Function Documentation

#### 6.17.2.1 xnclock\_adjust()

```
void xnclock_adjust (  
    struct xnclock * clock,  
    xnsticks_t delta )
```

Adjust a clock time.

This service changes the epoch for the given clock by applying the specified tick delta on its wallclock offset.

Parameters

<i>clock</i>	The clock to adjust.
<i>delta</i>	The adjustment value expressed in nanoseconds.

Tags

[task-unrestricted](#), [atomic-entry](#)

Note

Xenomai tracks the system time in *nkclock*, as a monotonously increasing count of ticks since the epoch. The epoch is initially the same as the underlying machine time.

#### 6.17.2.2 `xnclock_deregister()`

```
void xnclock_deregister (
    struct xnclock * clock )
```

Deregister a Xenomai clock.

This service uninstalls a Xenomai clock previously registered with [xnclock\\_register\(\)](#).

This service may be called once all timers driven by *clock* have been stopped.

Parameters

<i>clock</i>	The clock to deregister.
--------------	--------------------------

Tags

[secondary-only](#)

#### 6.17.2.3 `xnclock_register()`

```
int xnclock_register (
    struct xnclock * clock,
    const cpumask_t * affinity )
```

Register a Xenomai clock.

This service installs a new clock which may be used to drive Xenomai timers.

## Parameters

<i>clock</i>	The new clock to register.
<i>affinity</i>	The set of CPUs we may expect the backing clock device to tick on. As a special case, passing a NULL affinity mask means that timer IRQs cannot be seen as percpu events, in which case all outstanding timers will be maintained into a single global queue instead of percpu timer queues.

## Tags

[secondary-only](#)

## 6.17.2.4 xnclock\_tick()

```
void xnclock_tick (  
    struct xnclock * clock )
```

Process a clock tick.

This routine processes an incoming *clock* event, firing elapsed timers as appropriate.

## Parameters

<i>clock</i>	The clock for which a new event was received.
--------------	---

## Tags


[coreirq-only](#), [atomic-entry](#)

## Note

The current CPU must be part of the real-time affinity set unless the clock device has no particular IRQ affinity, otherwise weird things may happen.

## 6.18 Debugging services

Collaboration diagram for Debugging services:




### 6.18.1 Detailed Description

## 6.19 Dynamic memory allocation services

This code implements a variant of the allocator described in "Design of a General Purpose Memory Allocator for the 4.3BSD Unix Kernel" by Marshall K.

Collaboration diagram for Dynamic memory allocation services:



### Functions

- void [xnheap\\_destroy](#) (struct xnheap \*heap)
  - Destroys a memory heap.*
- void \* [xnheap\\_alloc](#) (struct xnheap \*heap, size\_t size)
  - Allocate a memory block from a memory heap.*
- void [xnheap\\_free](#) (struct xnheap \*heap, void \*block)
  - Release a block to a memory heap.*
- void [xnheap\\_set\\_name](#) (struct xnheap \*heap, const char \*name,...)
  - Set the heap's name string.*

### 6.19.1 Detailed Description

This code implements a variant of the allocator described in "Design of a General Purpose Memory Allocator for the 4.3BSD Unix Kernel" by Marshall K.

McKusick and Michael J. Karels (USENIX 1988), see <http://docs.FreeBSD.org/44doc/papers/kernmalloc.pdf>. The free page list is maintained in rbtrees for fast lookups of multi-page memory ranges, and pages holding bucketed memory have a fast allocation bitmap to manage their blocks internally.

### 6.19.2 Function Documentation

#### 6.19.2.1 xnheap\_alloc()

```
void * xnheap_alloc (
    struct xnheap * heap,
    size_t size )
```

Allocate a memory block from a memory heap.

Allocates a contiguous region of memory from an active memory heap. Such allocation is guaranteed to be time-bounded.



## Parameters

<i>heap</i>	The descriptor address of the heap to get memory from.
<i>size</i>	The size in bytes of the requested block.

## Returns

The address of the allocated region upon success, or NULL if no memory is available from the specified heap.

## Tags

[unrestricted](#)

## 6.19.2.2 xnheap\_destroy()

```
void xnheap_destroy (
    struct xnheap * heap )
```

Destroys a memory heap.

Destroys a memory heap.

## Parameters

<i>heap</i>	The heap descriptor.
-------------	----------------------

## Tags

[secondary-only](#)

## 6.19.2.3 xnheap\_free()

```
void xnheap_free (
    struct xnheap * heap,
    void * block )
```

Release a block to a memory heap.

Releases a memory block to a heap.

## Parameters

<i>heap</i>	The heap descriptor.
<i>block</i>	The block to be returned to the heap.

## Tags

[unrestricted](#)

## 6.19.2.4 xnheap\_set\_name()

```
void xnheap_set_name (  
    struct xnheap * heap,  
    const char * name,  
    ... )
```

Set the heap's name string.

Set the heap name that will be used in statistic outputs.

## Parameters

<i>heap</i>	The address of a heap descriptor.
<i>name</i>	Name displayed in statistic outputs. This parameter can be a printf()-like format argument list.


## Tags

[task-unrestricted](#)

## 6.20 Cobalt

Cobalt supplements the native Linux kernel in dual kernel configurations.

Collaboration diagram for Cobalt:



### Modules

- [Cobalt kernel](#)

*The Cobalt core is a co-kernel which supplements the Linux kernel for delivering real-time services with very low latency.*

- [Analogy framework](#)

*A RTDM-based interface for implementing DAQ card drivers.*

- [POSIX interface](#)

### 6.20.1 Detailed Description

Cobalt supplements the native Linux kernel in dual kernel configurations.

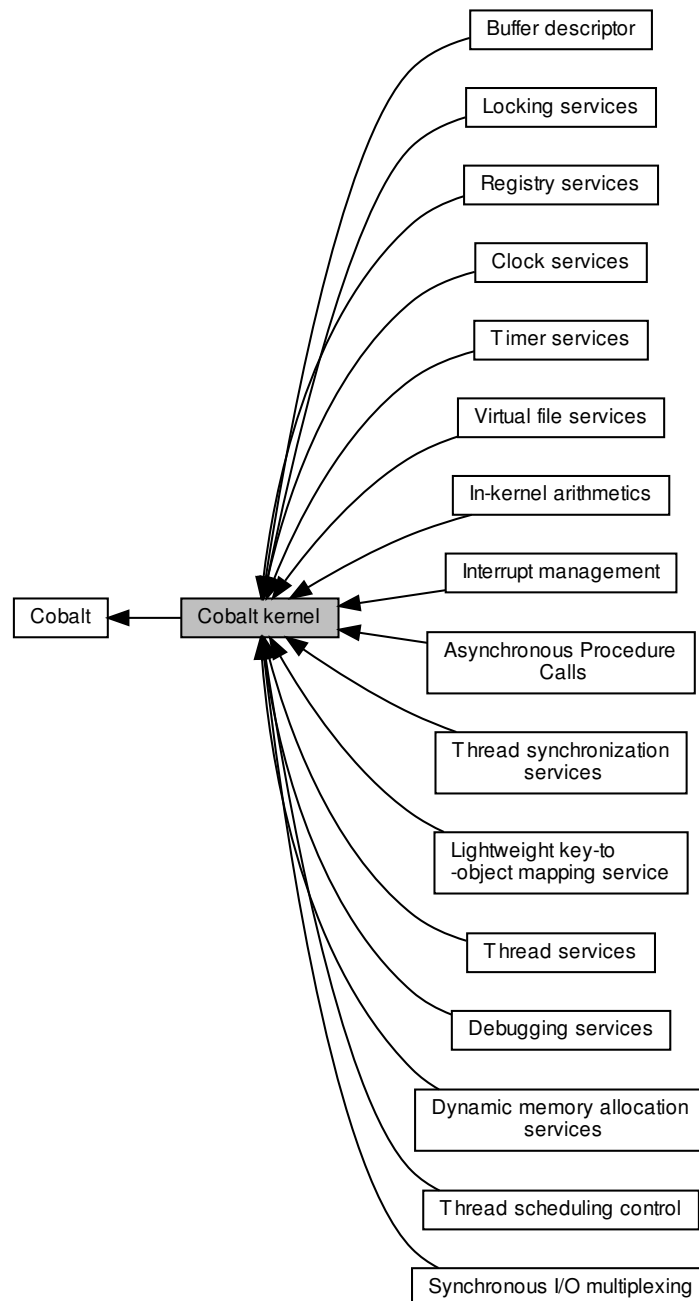
It deals with all time-critical activities, such as handling interrupts, and scheduling real-time threads. The Cobalt kernel has higher priority over all the native kernel activities.

Cobalt provides an implementation of the POSIX and RTDM interfaces based on a set of generic RTOS building blocks.

## 6.21 Cobalt kernel

The Cobalt core is a co-kernel which supplements the Linux kernel for delivering real-time services with very low latency.

Collaboration diagram for Cobalt kernel:



### Modules

- [Asynchronous Procedure Calls](#)

*Services for scheduling function calls in the Linux domain.*

- [In-kernel arithmetics](#)

*A collection of helpers performing arithmetics not implicitly available from kernel context via GCC helpers.*

- [Buffer descriptor](#)

*Abstraction for copying data to/from different address spaces.*

- [Clock services](#)

- [Debugging services](#)

- [Dynamic memory allocation services](#)

*This code implements a variant of the allocator described in "Design of a General Purpose Memory Allocator for the 4.3BSD Unix Kernel" by Marshall K.*

- [Interrupt management](#)

- [Locking services](#)

*The Xenomai core deals with concurrent activities from two distinct kernels running side-by-side.*

- [Lightweight key-to-object mapping service](#)

*A map is a simple indexing structure which associates unique integer keys with pointers to objects.*

- [Registry services](#)

*The registry provides a mean to index object descriptors on unique alphanumeric keys.*

- [Thread scheduling control](#)

- [Synchronous I/O multiplexing](#)

*This module implements the services needed for implementing the POSIX select() service, or any other event multiplexing services.*

- [Thread synchronization services](#)

- [Thread services](#)

- [Timer services](#)

*The Xenomai timer facility depends on a clock source (xnclock) for scheduling the next activation times.*

- [Virtual file services](#)

*Virtual files provide a mean to export Xenomai object states to user-space, based on common kernel interfaces.*

### 6.21.1 Detailed Description

The Cobalt core is a co-kernel which supplements the Linux kernel for delivering real-time services with very low latency.

It implements a set of generic RTOS building blocks, which the Cobalt/POSIX and Cobalt/RTDM APIs are based on. Cobalt has higher priority over the Linux kernel activities.

#### 6.21.1.1 Dual kernel service tags

The Cobalt kernel services may be restricted to particular calling contexts, or entail specific side-effects.

To describe this information, each service documented by this section bears a set of tags when applicable.

The table below matches the tags used throughout the documentation with the description of their meaning for the caller.

#### Context tags


Tag	Context on entry
primary-only	Must be called from a Cobalt task in primary mode
primary-timed	Requires a Cobalt task in primary mode if timed
coreirq-only	Must be called from a Cobalt IRQ handler
secondary-only	Must be called from a Cobalt task in secondary mode or regular Linux task
rtm-task	Must be called from a RTDM driver task
mode-unrestricted	May be called from a Cobalt task in either primary or secondary mode
task-unrestricted	May be called from a Cobalt or regular Linux task indifferently
unrestricted	May be called from any context previously described
atomic-entry	Caller must currently hold the big Cobalt kernel lock (nklock)

### Possible side-effects

Tag	Description
might-switch	The Cobalt kernel may switch context

## 6.22 Interrupt management

Collaboration diagram for Interrupt management:



### Functions

- void [xnintr\\_destroy](#) (struct xnintr \*intr)

*Destroy an interrupt descriptor.*

- int [xnintr\\_attach](#) (struct xnintr \*intr, void \*cookie)

*Attach an interrupt descriptor.*

- void [xnintr\\_detach](#) (struct xnintr \*intr)

*Detach an interrupt descriptor.*

- void [xnintr\\_enable](#) (struct xnintr \*intr)

*Enable an interrupt line.*

- void [xnintr\\_disable](#) (struct xnintr \*intr)

*Disable an interrupt line.*

- void [xnintr\\_affinity](#) (struct xnintr \*intr, cpumask\_t cpumask)

*Set processor affinity of interrupt.*

- int [xnintr\\_init](#) (struct xnintr \*intr, const char \*name, unsigned int irq, xnintr\_t isr, xniack\_t iack, int flags)

*Initialize an interrupt descriptor.*

### 6.22.1 Detailed Description

### 6.22.2 Function Documentation

### 6.22.2.1 xnintr\_affinity()

```
void xnintr_affinity (
    struct xnintr * intr,
    cpumask_t cpumask )
```

Set processor affinity of interrupt.

Restricts the IRQ line associated with the interrupt descriptor *intr* to be received only on processors which bits are set in *cpumask*.



## Parameters

<i>intr</i>	The address of the interrupt descriptor.
<i>cpumask</i>	The new processor affinity.

## Note

Depending on architectures, setting more than one bit in *cpumask* could be meaningless.

## Tags

[secondary-only](#)

## 6.22.2.2 xnintr\_attach()

```
int xnintr_attach (
    struct xnintr * intr,
    void * cookie )
```

Attach an interrupt descriptor.

Attach an interrupt descriptor previously initialized by [xnintr\\_init\(\)](#). This operation registers the descriptor at the interrupt pipeline, but does not enable the interrupt line yet. A call to [xnintr\\_enable\(\)](#) is required to start receiving IRQs from the interrupt line associated to the descriptor.

## Parameters

<i>intr</i>	The address of the interrupt descriptor to attach.
<i>cookie</i>	A user-defined opaque value which is stored into the descriptor for further retrieval by the interrupt handler.

## Returns

0 is returned on success. Otherwise:

- -EINVAL is returned if an error occurred while attaching the descriptor.
- -EBUSY is returned if the descriptor was already attached.

## Note

The caller **must not** hold `nklock` when invoking this service, this would cause deadlocks.

## Tags

[secondary-only](#)

## Note

Attaching an interrupt descriptor resets the tracked number of IRQ receipts to zero.

### 6.22.2.3 xnintr\_destroy()

```
void xnintr_destroy (
    struct xnintr * intr )
```

Destroy an interrupt descriptor.

Destroys an interrupt descriptor previously initialized by [xnintr\\_init\(\)](#). The descriptor is automatically detached by a call to [xnintr\\_detach\(\)](#). No more IRQs will be received through this descriptor after this service has returned.

Parameters

<i>intr</i>	The address of the interrupt descriptor to destroy.
-------------	---

Tags

[secondary-only](#)

### 6.22.2.4 xnintr\_detach()

```
void xnintr_detach (
    struct xnintr * intr )
```

Detach an interrupt descriptor.

This call unregisters an interrupt descriptor previously attached by [xnintr\\_attach\(\)](#) from the interrupt pipeline. Once detached, the associated interrupt line is disabled, but the descriptor remains valid. The descriptor can be attached anew by a call to [xnintr\\_attach\(\)](#).

Parameters

<i>intr</i>	The address of the interrupt descriptor to detach.
-------------	--

Note

The caller **must not** hold `nklock` when invoking this service, this would cause deadlocks.

Tags

[secondary-only](#)

## 6.22.2.5 xnintr\_disable()

```
void xnintr_disable (  
    struct xnintr * intr )
```

Disable an interrupt line.

Disables the interrupt line associated with an interrupt descriptor.

Parameters

<i>intr</i>	The address of the interrupt descriptor.
-------------	--

Tags

[secondary-only](#)

## 6.22.2.6 xnintr\_enable()

```
void xnintr_enable (  
    struct xnintr * intr )
```

Enable an interrupt line.

Enables the interrupt line associated with an interrupt descriptor.

Parameters

<i>intr</i>	The address of the interrupt descriptor.
-------------	--

Tags

[secondary-only](#)

## 6.22.2.7 xnintr\_init()

```
int xnintr_init (  
    struct xnintr * intr,  
    const char * name,  
    unsigned int irq,  
    xnintr_t isr,
```

```
xniack_t iack,
int flags )
```

Initialize an interrupt descriptor.

When an interrupt occurs on the given *irq* line, the interrupt service routine *isr* is fired in order to deal with the hardware event. The interrupt handler may call any non-blocking service from the Cobalt core.

Upon receipt of an IRQ, the interrupt handler *isr* is immediately called on behalf of the interrupted stack context, the rescheduling procedure is locked, and the interrupt line is masked in the system interrupt controller chip. Upon return, the status of the interrupt handler is checked for the following bits:

- XN\_IRQ\_HANDLED indicates that the interrupt request was successfully handled.
- XN\_IRQ\_NONE indicates the opposite to XN\_IRQ\_HANDLED, meaning that no interrupt source could be identified for the ongoing request by the handler.

In addition, one of the following bits may be present in the status:

- XN\_IRQ\_DISABLE tells the Cobalt core to disable the interrupt line before returning from the interrupt context.
- XN\_IRQ\_PROPAGATE propagates the IRQ event down the interrupt pipeline to Linux. Using this flag is strongly discouraged, unless you fully understand the implications of such propagation.

#### Warning

The handler should not use these bits if it shares the interrupt line with other handlers in the real-time domain. When any of these bits is detected, the interrupt line is left masked.

A count of interrupt receipts is tracked into the interrupt descriptor, and reset to zero each time such descriptor is attached. Since this count could wrap around, it should be used as an indication of interrupt activity only.

#### Parameters

<i>intr</i>	The address of a descriptor the Cobalt core will use to store the interrupt-specific data.
<i>name</i>	An ASCII string standing for the symbolic name of the interrupt or NULL.
<i>irq</i>	The IRQ line number associated with the interrupt descriptor. This value is architecture-dependent. An interrupt descriptor must be attached to the system by a call to <a href="#">xnintr_attach()</a> before <i>irq</i> events can be received.
<i>isr</i>	The address of an interrupt handler, which is passed the address of the interrupt descriptor receiving the IRQ.
<i>iack</i>	The address of an optional interrupt acknowledge routine, aimed at replacing the default one. Only very specific situations actually require to override the default setting for this parameter, like having to acknowledge non-standard PIC hardware. <i>iack</i> should return a non-zero value to indicate that the interrupt has been properly acknowledged. If <i>iack</i> is NULL, the default routine will be used instead.
<i>flags</i>	A set of creation flags affecting the operation. The valid flags are:

- `XN_IRQTYPE_SHARED` enables IRQ-sharing with other interrupt objects.
- `XN_IRQTYPE_EDGE` is an additional flag need to be set together with `XN_IRQTYPE_SHARED` to enable IRQ-sharing of edge-triggered interrupts.

#### Returns

0 is returned on success. Otherwise, `-EINVAL` is returned if *irq* is not a valid interrupt number.


#### Tags

[secondary-only](#)

## 6.23 Locking services

The Xenomai core deals with concurrent activities from two distinct kernels running side-by-side.

Collaboration diagram for Locking services:



### Macros

- `#define splhigh(x) ((x) = ipipe_test_and_stall_head() & 1)`  
*Hard disable interrupts on the local processor, saving previous state.*
- `#define splexit(x) ipipe_restore_head(x & 1)`  
*Restore the saved hard interrupt state on the local processor.*
- `#define splmax() ipipe_stall_head()`  
*Hard disable interrupts on the local processor.*
- `#define splnone() ipipe_unstall_head()`  
*Hard enable interrupts on the local processor.*
- `#define spltest() ipipe_test_head()`  
*Test hard interrupt state on the local processor.*

### 6.23.1 Detailed Description

The Xenomai core deals with concurrent activities from two distinct kernels running side-by-side.

When interrupts are involved, the services from this section control the **hard** interrupt state exclusively, for protecting against processor-local or SMP concurrency.

#### Note

In a dual kernel configuration, *hard interrupts* are gated by the CPU. When enabled, hard interrupts are immediately delivered to the Xenomai core if they belong to a real-time source, or deferred until enabled by a second-stage virtual interrupt mask, if they belong to regular Linux devices/sources.

### 6.23.2 Macro Definition Documentation

#### 6.23.2.1 splexit

```
#define splexit(
    x ) ipipe_restore_head(x & 1)
```

Restore the saved hard interrupt state on the local processor.

Parameters

in	x	The context variable previously updated by <a href="#">splhigh()</a>
----	---	--

### 6.23.2.2 splhigh

```
#define splhigh(  
    x ) ((x) = ipipe_test_and_stall_head() & 1)
```

Hard disable interrupts on the local processor, saving previous state.

Parameters

out	x	An unsigned long integer context variable
-----	---	---

### 6.23.2.3 spltest

```
#define spltest( ) ipipe_test_head()
```

Test hard interrupt state on the local processor.


Returns

Zero if the local processor currently accepts interrupts, non-zero otherwise.

## 6.24 Lightweight key-to-object mapping service

A map is a simple indexing structure which associates unique integer keys with pointers to objects.

Collaboration diagram for Lightweight key-to-object mapping service:



### Functions

- struct xnmap \* [xnmap\\_create](#) (int nkeys, int reserve, int offset)

*Create a map.*

- void [xnmap\\_delete](#) (struct xnmap \*map)

*Delete a map.*

- int [xnmap\\_enter](#) (struct xnmap \*map, int key, void \*objaddr)

*Index an object into a map.*

- int [xnmap\\_remove](#) (struct xnmap \*map, int key)

*Remove an object reference from a map.*

- static void \* [xnmap\\_fetch\\_nocheck](#) (struct xnmap \*map, int key)

*Search an object into a map - unchecked form.*

- static void \* [xnmap\\_fetch](#) (struct xnmap \*map, int key)

*Search an object into a map.*

### 6.24.1 Detailed Description

A map is a simple indexing structure which associates unique integer keys with pointers to objects.

The current implementation supports reservation, for naming/indexing objects, either on a fixed, user-provided integer (i.e. a reserved key value), or by drawing the next available key internally if the caller did not specify any fixed key. For instance, in some given map, the key space ranging from 0 to 255 could be reserved for fixed keys, whilst the range from 256 to 511 could be available for drawing free keys dynamically.

A maximum of 1024 unique keys per map is supported on 32bit machines.

(This implementation should not be confused with C++ STL maps, which are dynamically expandable and allow arbitrary key types; Xenomai maps don't).



## 6.24.2 Function Documentation

### 6.24.2.1 xnmap\_create()

```
struct xnmap * xnmap_create (
    int nkeys,
    int reserve,
    int offset )
```

Create a map.

Allocates a new map with the specified addressing capabilities. The memory is obtained from the Xenomai system heap.

Parameters

<i>nkeys</i>	The maximum number of unique keys the map will be able to hold. This value cannot exceed the static limit represented by <code>XNMAP_MAX_KEYS</code> , and must be a power of two.
<i>reserve</i>	The number of keys which should be kept for reservation within the index space. Reserving a key means to specify a valid key to the <code>xnmap_enter()</code> service, which will then attempt to register this exact key, instead of drawing the next available key from the unreserved index space. When reservation is in effect, the unreserved index space will hold key values greater than <i>reserve</i> , keeping the low key values for the reserved space. For instance, passing <i>reserve</i> = 32 would cause the index range [ 0 .. 31 ] to be kept for reserved keys. When non-zero, <i>reserve</i> is rounded to the next multiple of <code>BITS_PER_LONG</code> . If <i>reserve</i> is zero no reservation will be available from the map.
<i>offset</i>	The lowest key value <code>xnmap_enter()</code> will return to the caller. Key values will be in the range [ 0 + offset .. <i>nkeys</i> + offset - 1 ]. Negative offsets are valid.

Returns

the address of the new map is returned on success; otherwise, NULL is returned if *nkeys* is invalid.

Tags

[task-unrestricted](#)

### 6.24.2.2 xnmap\_delete()

```
void xnmap_delete (
    struct xnmap * map )
```

Delete a map.

Deletes a map, freeing any associated memory back to the Xenomai system heap.

## Parameters

<i>map</i>	The address of the map to delete.
------------	-----------------------------------

## Tags

[task-unrestricted](#)

## 6.24.2.3 xnmap\_enter()

```
int xnmap_enter (
    struct xnmap * map,
    int key,
    void * objaddr )
```

Index an object into a map.

Insert a new object into the given map.

## Parameters

<i>map</i>	The address of the map to insert into.
<i>key</i>	The key to index the object on. If this key is within the valid index range [ 0 - offset .. nkeys - offset - 1 ], then an attempt to reserve this exact key is made. If <i>key</i> has an out-of-range value lower or equal to 0 - offset - 1, then an attempt is made to draw a free key from the unreserved index space.
<i>objaddr</i>	The address of the object to index on the key. This value will be returned by a successful call to <a href="#">xnmap_fetch()</a> with the same key.

## Returns

a valid key is returned on success, either *key* if reserved, or the next free key. Otherwise:

- -EEXIST is returned upon attempt to reserve a busy key.
- -ENOSPC when no more free key is available.

## Tags

[unrestricted](#)

6.24.2.4 `xnmap_fetch()`

```
void xnmap_fetch (  
    struct xnmap * map,  
    int key ) [inline], [static]
```

Search an object into a map.

Retrieve an object reference from the given map by its index key.

Parameters

<i>map</i>	The address of the map to retrieve from.
<i>key</i>	The key to be searched for in the map index.

Returns

The indexed object address is returned on success, otherwise NULL is returned when *key* is invalid or no object is currently indexed on it.

Tags

[unrestricted](#)

6.24.2.5 `xnmap_fetch_nocheck()`

```
void xnmap_fetch_nocheck (  
    struct xnmap * map,  
    int key ) [inline], [static]
```

Search an object into a map - unchecked form.

Retrieve an object reference from the given map by its index key, but does not perform any sanity check on the provided key.

Parameters

<i>map</i>	The address of the map to retrieve from.
<i>key</i>	The key to be searched for in the map index.

Returns

The indexed object address is returned on success, otherwise NULL is returned when no object is currently indexed on *key*.

Tags

[unrestricted](#)

### 6.24.2.6 xnmap\_remove()

```
int xnmap_remove (
    struct xnmap * map,
    int key )
```

Remove an object reference from a map.

Removes an object reference from the given map, releasing the associated key.

Parameters

<i>map</i>	The address of the map to remove from.
<i>key</i>	The key the object reference to be removed is indexed on.

Returns

0 is returned on success. Otherwise:

- -ESRCH is returned if *key* is invalid.


Tags

[unrestricted](#)

## 6.25 Registry services

The registry provides a mean to index object descriptors on unique alphanumeric keys.

Collaboration diagram for Registry services:



### Functions

- int [xnregistry\\_enter](#) (const char \*key, void \*objaddr, xnhandle\_t \*phandle, struct xnnode \*pnode)

*Register a real-time object.*

- int [xnregistry\\_bind](#) (const char \*key, xnticks\_t timeout, int timeout\_mode, xnhandle\_t \*phandle)

*Bind to a real-time object.*

- int [xnregistry\\_remove](#) (xnhandle\_t handle)

*Forcibly unregister a real-time object.*

- static void \* [xnregistry\\_lookup](#) (xnhandle\_t handle, unsigned long \*cstamp\_r)

*Find a real-time object into the registry.*

- int [xnregistry\\_unlink](#) (const char \*key)

*Turn a named object into an anonymous object*

### 6.25.1 Detailed Description

The registry provides a mean to index object descriptors on unique alphanumeric keys.

When labeled this way, an object is globally exported; it can be searched for, and its descriptor returned to the caller for further use; the latter operation is called a "binding". When no object has been registered under the given name yet, the registry can be asked to set up a rendez-vous, blocking the caller until the object is eventually registered.

### 6.25.2 Function Documentation

## 6.25.2.1 xnregistry\_bind()

```
int xnregistry_bind (
    const char * key,
    xnticks_t timeout,
    int timeout_mode,
    xnhandle_t * phandle )
```

Bind to a real-time object.

This service retrieves the registry handle of a given object identified by its key. Unless otherwise specified, this service will block the caller if the object is not registered yet, waiting for such registration to occur.

## Parameters

<i>key</i>	A valid NULL-terminated string which identifies the object to bind to.
<i>timeout</i>	The timeout which may be used to limit the time the thread wait for the object to be registered. This value is a wait time given as a count of nanoseconds. It can either be relative, absolute monotonic (XN_ABSOLUTE), or absolute adjustable (XN_REALTIME) depending on <i>timeout_mode</i> . Passing XN_INFINITE <b>and</b> setting <i>timeout_mode</i> to XN_RELATIVE specifies an unbounded wait. Passing XN_NONBLOCK causes the service to return immediately without waiting if the object is not registered on entry. All other values are used as a wait limit.
<i>timeout_mode</i>	The mode of the <i>timeout</i> parameter. It can either be set to XN_RELATIVE, XN_ABSOLUTE, or XN_REALTIME (see also <a href="#">xntimer_start()</a> ).
<i>phandle</i>	A pointer to a memory location which will be written upon success with the generic handle defined by the registry for the retrieved object. Contents of this memory is undefined upon failure.

## Returns

0 is returned upon success. Otherwise:

- -EINVAL is returned if *key* is NULL.
- -EINTR is returned if [xnthread\\_unblock\(\)](#) has been called for the waiting thread before the retrieval has completed.
- -EWOULDBLOCK is returned if *timeout* is equal to XN\_NONBLOCK and the searched object is not registered on entry. As a special exception, this error is also returned if this service should block, but was called from a context which cannot sleep (e.g. interrupt, non-realtime or scheduler locked).
- -ETIMEDOUT is returned if the object cannot be retrieved within the specified amount of time.

## Tags

[primary-only](#), [might-switch](#)

## Note

[xnregistry\\_bind\(\)](#) only returns the index portion of a handle, which might include other fixed bits to be complete (e.g. XNSYNCH\_PSHARED). The caller is responsible for completing the handle returned with those bits if applicable, depending on the context.

## 6.25.2.2 xnregistry\_enter()

```
int xnregistry_enter (
    const char * key,
    void * objaddr,
    xnhandle_t * phandle,
    struct xninode * pnode )
```

Register a real-time object.

This service allocates a new registry slot for an associated object, and indexes it by an alphanumeric key for later retrieval.

## Parameters

<i>key</i>	A valid NULL-terminated string by which the object will be indexed and later retrieved in the registry. Since it is assumed that such key is stored into the registered object, it will <i>not</i> be copied but only kept by reference in the registry. Pass an empty or NULL string if the object shall only occupy a registry slot for handle-based lookups. The slash character is not accepted in <i>key</i> if <i>pnode</i> is non-NULL.
<i>objaddr</i>	An opaque pointer to the object to index by <i>key</i> .
<i>phandle</i>	A pointer to a generic handle defined by the registry which will uniquely identify the indexed object, until the latter is unregistered using the <a href="#">xnregistry_remove()</a> service.
<i>pnode</i>	A pointer to an optional /proc node class descriptor. This structure provides the information needed to export all objects from the given class through the /proc filesystem, under the /proc/xenomai/registry entry. Passing NULL indicates that no /proc support is available for the newly registered object.

## Returns

0 is returned upon success. Otherwise:

- -EINVAL is returned if *objaddr* is NULL.
- -EINVAL if *pnode* is non-NULL, and *key* points to a valid string containing a '/' character.
- -ENOMEM is returned if the system fails to get enough dynamic memory from the global real-time heap in order to register the object.
- -EEXIST is returned if the *key* is already in use.

## Tags

[unrestricted](#), [might-switch](#), [atomic-entry](#)

## 6.25.2.3 xnregistry\_lookup()

```
void * xnregistry_lookup (
    xnhandle_t handle,
    unsigned long * cstamp_r ) [inline], [static]
```

Find a real-time object into the registry.

This service retrieves an object from its handle into the registry and returns the memory address of its descriptor. Optionally, it also copies back the object's creation stamp which is unique across object registration calls.

## Parameters

<i>handle</i>	The generic handle of the object to fetch.
<i>cstamp_r</i>	If not-NULL, the object's creation stamp will be copied to this memory area.

## Returns

The memory address of the object's descriptor is returned on success. Otherwise, NULL is returned if *handle* does not reference a registered object.

## Tags

[unrestricted](#)

## 6.25.2.4 xnregistry\_remove()

```
int xnregistry_remove (
    xnhandle_t handle )
```

Forcibly unregister a real-time object.

This service forcibly removes an object from the registry. The removal is performed regardless of the current object's locking status.

## Parameters

<i>handle</i>	The generic handle of the object to remove.
---------------	---

## Returns

0 is returned upon success. Otherwise:

- -ESRCH is returned if *handle* does not reference a registered object.



Tags

[unrestricted](#)

#### 6.25.2.5 xnregistry\_unlink()

```
int xnregistry_unlink (  
    const char * key )
```

Turn a named object into an anonymous object

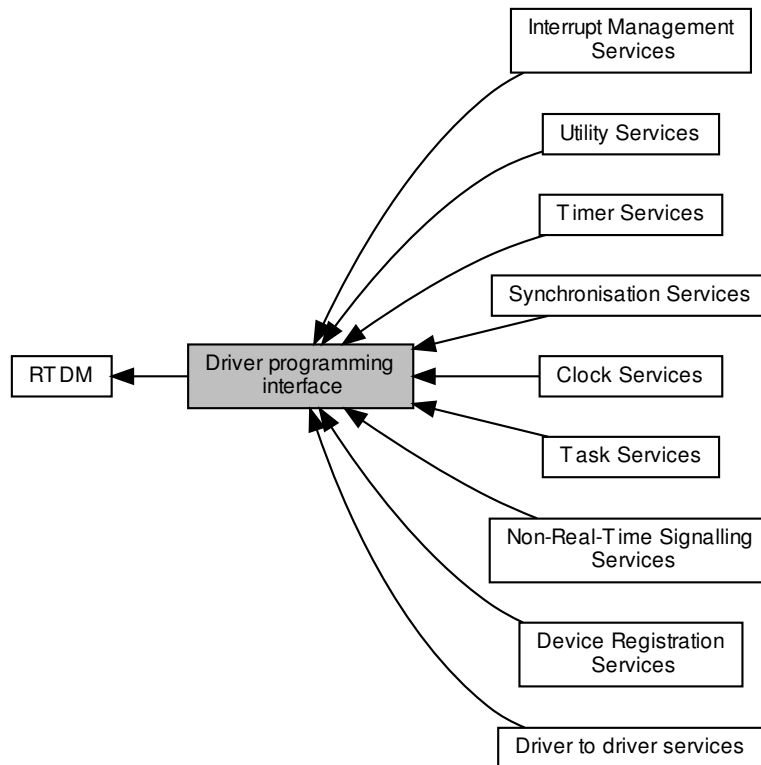
Tags

[unrestricted](#)

## 6.26 Driver programming interface

RTDM driver programming interface.

Collaboration diagram for Driver programming interface:



### Modules

- [Driver to driver services](#)  
*Inter-driver interface.*
- [Device Registration Services](#)
- [Clock Services](#)
- [Task Services](#)
- [Timer Services](#)
- [Synchronisation Services](#)
- [Interrupt Management Services](#)
- [Non-Real-Time Signalling Services](#)

*These services provide a mechanism to request the execution of a specified handler in non-real-time context.*

- [Utility Services](#)


### 6.26.1 Detailed Description

RTDM driver programming interface.

## 6.27 Driver to driver services

Inter-driver interface.

Collaboration diagram for Driver to driver services:



### Functions

- `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_rcvfrom` (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 to incoming connection requests*

- int `rtdm_accept` (int fd, struct sockaddr \*addr, socklen\_t \*addrlen)

*Accept a connection request*

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

*Shut down parts of a connection*

- int `rtdm_getsockopt` (int fd, int level, int optname, void \*optval, socklen\_t \*optlen)

*Get socket option*

- int `rtdm_setsockopt` (int fd, int level, int optname, const void \*optval, socklen\_t optlen)

*Set socket option*

- int `rtdm_getsockname` (int fd, struct sockaddr \*name, socklen\_t \*namelen)

*Get local socket address*

- int `rtdm_getpeername` (int fd, struct sockaddr \*name, socklen\_t \*namelen)

*Get socket destination address*

## 6.27.1 Detailed Description

Inter-driver interface.

## 6.27.2 Function Documentation

### 6.27.2.1 `rtdm_accept()`

```
int rtdm_accept (
    int fd,
    struct sockaddr * addr,
    socklen_t * addrlen )
```

Accept a connection request

Accept connection requests

Refer to `rtdm_accept()` for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[mode-unrestricted](#), [might-switch](#)

Parameters

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

Returns

0 on success, otherwise negative error code

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

See also

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

Tags

[mode-unrestricted](#), [might-switch](#)

### 6.27.2.2 `rtdm_bind()`

```
int rtdm_bind (
    int fd,
    const struct sockaddr * my_addr,
    socklen_t addrlen )
```

Bind to local address

Refer to `rtdm_bind()` for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[task-unrestricted](#), [might-switch](#)

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtdm_socket()</a>
in	<i>my_addr</i>	Address buffer
in	<i>addrlen</i>	Address buffer size

## Returns

0 on success, otherwise negative error code

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

## See also

[bind\(\)](#) in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

## Tags

[mode-unrestricted](#), [might-switch](#)

6.27.2.3 `rtdm_close()`

```
int rtdm_close (  
    int fd )
```

Close a device or socket

Refer to [rtdm\\_close\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

## Tags

[secondary-only](#), [might-switch](#)

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtdm_open()</a> or <a href="#">rtdm_socket()</a>
----	-----------	---

## Returns

0 on success, otherwise a negative error code.

## Note

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

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

## See also

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

## Tags

[secondary-only](#), [might-switch](#)

6.27.2.4 `rtm_connect()`

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

Connect to remote address

Refer to `rtm_connect()` for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

## Tags

[mode-unrestricted](#), [might-switch](#)

## Parameters

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

## Returns

0 on success, otherwise negative error code

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

## See also

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

Tags

[mode-unrestricted](#), [might-switch](#)

#### 6.27.2.5 rtdm\_getpeername()

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

Get socket destination address

Refer to [rtdm\\_getpeername\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[task-unrestricted](#), [might-switch](#)

Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtdm_socket()</a>
out	<i>name</i>	Address buffer
in,out	<i>namelen</i>	Address buffer size

Returns

0 on success, otherwise negative error code

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

See also

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

Tags

[task-unrestricted](#), [might-switch](#)



## 6.27.2.6 rtdm\_getsockname()

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

Get local socket address

Refer to [rtdm\\_getsockname\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[task-unrestricted](#), [might-switch](#)

Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtdm_socket()</a>
out	<i>name</i>	Address buffer
in,out	<i>namelen</i>	Address buffer size

Returns

0 on success, otherwise negative error code

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

See also

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

Tags

[task-unrestricted](#), [might-switch](#)

## 6.27.2.7 rtdm\_getsockopt()

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

Get socket option

Refer to [rtdm\\_getsockopt\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[task-unrestricted](#), [might-switch](#)

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtm_socket()</a>
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

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

## See also

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

## Tags

[task-unrestricted](#), [might-switch](#)

6.27.2.8 `rtm_ioctl()`

```
int rtm_ioctl (
    int fd,
    int request,
    ... )
```

## Issue an IOCTL

Refer to [rtm\\_ioctl\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

## Tags

[task-unrestricted](#), [might-switch](#)

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtm_open()</a> or <a href="#">rtm_socket()</a>
in	<i>request</i>	IOCTL code
	...	Optional third argument, depending on IOCTL function (void * or unsigned long)

## Returns

Positiv value on success, otherwise negative error code

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

## See also

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

## Tags

[task-unrestricted](#), [might-switch](#)

6.27.2.9 `rtdm_listen()`

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

Listen to incoming connection requests

Listen for incomming connection requests

Refer to `rtdm_listen()` for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

## Tags

[task-unrestricted](#), [might-switch](#)

## Parameters

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

## Returns

0 on success, otherwise negative error code

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

## See also

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

Tags

[task-unrestricted](#), [might-switch](#)

#### 6.27.2.10 rtdm\_open()

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

Open a device

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

Tags

[secondary-only](#), [might-switch](#)

Parameters

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

Returns

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

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

See also

[open\(\)](#) in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

Tags

[secondary-only](#), [might-switch](#)

6.27.2.11 `rtdm_read()`

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

Read from device

Refer to `rtdm_read()` for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[mode-unrestricted](#), [might-switch](#)

Parameters

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

Returns

Number of bytes read, otherwise negative error code

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

See also

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

Tags

[mode-unrestricted](#), [might-switch](#)

6.27.2.12 `rtdm_recv()`

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

Receive message from socket

Refer to `rtdm_recv()` for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[mode-unrestricted](#), [might-switch](#)

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtm_socket()</a>
out	<i>buf</i>	Message buffer
in	<i>len</i>	Message buffer size
in	<i>flags</i>	Message flags

## Returns

Number of bytes received, otherwise negative error code

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

## See also

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

## Tags

[mode-unrestricted](#), [might-switch](#)

6.27.2.13 `rtm_rcvfrom()`

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

## Receive message from socket

Refer to [rtm\\_rcvfrom\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

## Tags

[mode-unrestricted](#), [might-switch](#)

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtm_socket()</a>
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

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

## See also

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

## Tags

[mode-unrestricted](#), [might-switch](#)

6.27.2.14 `rtdm_recvmsg()`

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

Receive message from socket

Refer to `rtdm_recvmsg()` for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

## Tags

[mode-unrestricted](#), [might-switch](#)

## Parameters

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

## Returns

Number of bytes received, otherwise negative error code

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

## See also

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

## Tags

[mode-unrestricted](#), [might-switch](#)

## 6.27.2.15 rtdm\_send()

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

Transmit message to socket

Refer to [rtdm\\_send\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[mode-unrestricted](#), [might-switch](#)

Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtdm_socket()</a>
in	<i>buf</i>	Message buffer
in	<i>len</i>	Message buffer size
in	<i>flags</i>	Message flags

Returns

Number of bytes sent, otherwise negative error code

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

See also

[send\(\)](#) in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

Tags

[mode-unrestricted](#), [might-switch](#)

## 6.27.2.16 rtdm\_sendmsg()

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

Transmit message to socket

Refer to [rtdm\\_sendmsg\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[mode-unrestricted](#), [might-switch](#)



## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtm_socket()</a>
in	<i>msg</i>	Message descriptor
in	<i>flags</i>	Message flags

## Returns

Number of bytes sent, otherwise negative error code

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

## See also

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

## Tags

[mode-unrestricted](#), [might-switch](#)

6.27.2.17 `rtm_sendto()`

```
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

Refer to [rtm\\_sendto\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

## Tags

[mode-unrestricted](#), [might-switch](#)

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtm_socket()</a>
in	<i>buf</i>	Message buffer
in	<i>len</i>	Message buffer size
in	<i>flags</i>	Message flags
in	<i>to</i>	Buffer for message destination address
in	<i>tolen</i>	Address buffer size

## Returns

Number of bytes sent, otherwise negative error code

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

## See also

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

## Tags

[mode-unrestricted](#), [might-switch](#)

6.27.2.18 `rtm_setsockopt()`

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

## Set socket option

Refer to `rtm_setsockopt()` for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

## Tags

[task-unrestricted](#), [might-switch](#)

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtm_socket()</a>
in	<i>level</i>	Addressed stack level
in	<i>optname</i>	Option name ID
in	<i>optval</i>	Value buffer
in	<i>optlen</i>	Value buffer size

## Returns

0 on success, otherwise negative error code

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

See also

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

Tags

[task-unrestricted](#), [might-switch](#)

#### 6.27.2.19 `rtdm_shutdown()`

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

Shut down parts of a connection

Refer to `rtdm_shutdown()` for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[secondary-only](#), [might-switch](#)

Parameters

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

Returns

0 on success, otherwise negative error code

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

See also

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

Tags

[secondary-only](#), [might-switch](#)

## 6.27.2.20 rtdm\_socket()

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

Create a socket

Refer to [rtdm\\_socket\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[secondary-only](#), [might-switch](#)

Parameters

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

Returns

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

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

See also

[socket\(\)](#) in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

Tags

[secondary-only](#), [might-switch](#)

## 6.27.2.21 rtdm\_write()

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

Write to device

Refer to [rtdm\\_write\(\)](#) for parameters and return values. Action depends on driver implementation, see [Device Profiles](#).

Tags

[mode-unrestricted](#), [might-switch](#)

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">rtm_open()</a>
in	<i>buf</i>	Output buffer
in	<i>nbyte</i>	Number of bytes to write

## Returns

Number of bytes written, otherwise negative error code

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

## See also

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

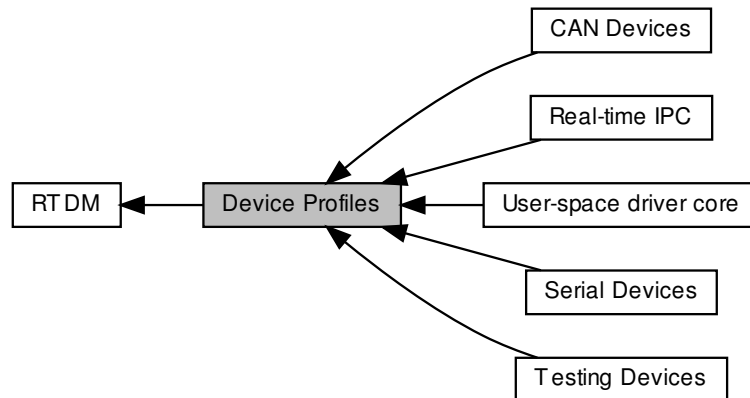
## Tags

[mode-unrestricted](#), [might-switch](#)

## 6.28 Device Profiles

Pre-defined classes of real-time devices.

Collaboration diagram for Device Profiles:



### Modules

- [User-space driver core](#)

*This profile includes all mini-drivers sitting on top of the User-space Device Driver framework (UDD).*

- [CAN Devices](#)

*This is the common interface a RTDM-compliant CAN device has to provide.*

- [Serial Devices](#)

*This is the common interface a RTDM-compliant serial device has to provide.*

- [Testing Devices](#)

*This group of devices is intended to provide in-kernel testing results.*

- [Real-time IPC](#)

**Profile Revision: 1**

### 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\_COBALT** 8
- #define **RTDM\_CLASS\_UDD** 9
- #define **RTDM\_CLASS\_MEMORY** 10
- #define **RTDM\_CLASS\_GPIO** 11
- #define **RTDM\_CLASS\_SPI** 12
- #define **RTDM\_CLASS\_PWM** 13
- #define **RTDM\_CLASS\_MISC** 223
- #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 `rtdm_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.*

### 6.28.1 Detailed Description

Pre-defined classes of real-time devices.

Device profiles define which operation handlers a driver of a certain class of devices 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.

### 6.28.2 Macro Definition Documentation

#### 6.28.2.1 RTIOC\_DEVICE\_INFO

```
#define RTIOC_DEVICE_INFO _IOR(RTIOC_TYPE_COMMON, 0x00, struct rtdm_device_info)
```

Retrieve information about a device or socket.

Parameters

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

#### 6.28.2.2 RTIOC\_PURGE

```
#define RTIOC_PURGE _IOW(RTIOC_TYPE_COMMON, 0x10, int)
```

Purge internal device or socket buffers.


Parameters

in	<i>arg</i>	Purge mask, see <a href="#">RTDM_PURGE_xxx_BUFFER</a>
----	------------	---



## 6.29 Device Registration Services

Collaboration diagram for Device Registration Services:



### Data Structures

- struct [rtm\\_fd\\_ops](#)  
*RTDM file operation descriptor.*

### Macros

- #define [RTDM\\_MAX\\_MINOR](#) 4096  
*Maximum number of named devices per driver.*

### Functions

- int [rtm\\_open\\_handler](#) (struct [rtm\\_fd](#) \*fd, int oflags)  
*Open handler for named devices.*
- int [rtm\\_socket\\_handler](#) (struct [rtm\\_fd](#) \*fd, int protocol)  
*Socket creation handler for protocol devices.*
- void [rtm\\_close\\_handler](#) (struct [rtm\\_fd](#) \*fd)  
*Close handler.*
- int [rtm\\_ioctl\\_handler](#) (struct [rtm\\_fd](#) \*fd, unsigned int request, void \_\_user \*arg)  
*IOCTL handler.*
- ssize\_t [rtm\\_read\\_handler](#) (struct [rtm\\_fd](#) \*fd, void \_\_user \*buf, size\_t size)  
*Read handler.*
- ssize\_t [rtm\\_write\\_handler](#) (struct [rtm\\_fd](#) \*fd, const void \_\_user \*buf, size\_t size)  
*Write handler.*
- ssize\_t [rtm\\_recvmmsg\\_handler](#) (struct [rtm\\_fd](#) \*fd, struct [user\\_mmsg](#) \*mmsg, int flags)  
*Receive message handler.*
- ssize\_t [rtm\\_sendmmsg\\_handler](#) (struct [rtm\\_fd](#) \*fd, const struct [user\\_mmsg](#) \*mmsg, int flags)  
*Transmit message handler.*
- int [rtm\\_select\\_handler](#) (struct [rtm\\_fd](#) \*fd, struct [xselector](#) \*selector, unsigned int type, unsigned int index)  
*Select handler.*
- int [rtm\\_mmap\\_handler](#) (struct [rtm\\_fd](#) \*fd, struct [vm\\_area\\_struct](#) \*vma)  
*Memory mapping handler.*

- unsigned long `rtdm_get_unmapped_area_handler` (struct `rtdm_fd` \*fd, unsigned long len, unsigned long pgoff, unsigned long flags)

*Allocate mapping region in address space.*

- int `rtdm_dev_register` (struct `rtdm_device` \*dev)

*Register a RTDM device*

- void `rtdm_dev_unregister` (struct `rtdm_device` \*dev)

*Unregister a RTDM device*

- int `rtdm_drv_set_sysclass` (struct `rtdm_driver` \*drv, struct class \*cls)

*Set the kernel device class of a RTDM driver.*

## 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_FIXED_MINOR` 0x0002  
*Use fixed minor provided in the `rtdm_device` description for registering.*
- #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.*
- #define `RTDM_SECURE_DEVICE` 0x80000000  
*Flag indicating a secure variant of RTDM (not supported here)*

### 6.29.1 Detailed Description

### 6.29.2 Macro Definition Documentation

#### 6.29.2.1 RTDM\_DEVICE\_TYPE\_MASK

```
#define RTDM_DEVICE_TYPE_MASK 0x00F0
```

Mask selecting the device type.

### 6.29.2.2 RTDM\_EXCLUSIVE

```
#define RTDM_EXCLUSIVE 0x0001
```

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

### 6.29.2.3 RTDM\_FIXED\_MINOR

```
#define RTDM_FIXED_MINOR 0x0002
```

Use fixed minor provided in the [rtdm\\_device](#) description for registering.

If this flag is absent, the RTDM core assigns minor numbers to devices managed by a driver in order of registration.

### 6.29.2.4 RTDM\_MAX\_MINOR

```
#define RTDM_MAX_MINOR 4096
```

Maximum number of named devices per driver.

### 6.29.2.5 RTDM\_NAMED\_DEVICE

```
#define RTDM_NAMED_DEVICE 0x0010
```

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

### 6.29.2.6 RTDM\_PROTOCOL\_DEVICE

```
#define RTDM_PROTOCOL_DEVICE 0x0020
```

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

Referenced by `udd_register_device()`.

## 6.29.3 Function Documentation

### 6.29.3.1 `rtdm_close_handler()`

```
void rtdm_close_handler (  
    struct rtdm_fd * fd )
```

Close handler.

## Parameters

in	<i>fd</i>	File descriptor associated with opened device instance.
----	-----------	---

## See also

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

6.29.3.2 `rtm_dev_register()`

```
int rtm_dev_register (  
    struct rtm_device * dev )
```

## Register a RTDM device

Registers a device in the RTDM namespace.

## Parameters

in	<i>dev</i>	Device descriptor.
----	------------	--------------------

## Returns

0 is returned upon success. Otherwise:

- -EINVAL is returned if the descriptor contains invalid entries. `RTDM_PROFILE_INFO()` must appear in the list of initializers for the driver properties.
- -EEXIST is returned if the specified device name or protocol ID is already in use.
- -ENOMEM is returned if a memory allocation failed in the process of registering the device.
- -EAGAIN is returned if no registry slot is available (check/raise `CONFIG_XENO_OPT_REGISTRY_NRSLOTS`).
- -ENOSYS is returned if the real-time core is disabled.
- -ENXIO is returned if no valid minor could be assigned

## Tags

[secondary-only](#)

6.29.3.3 `rtm_dev_unregister()`

```
void rtdm_dev_unregister (
    struct rtdm_device * dev )
```

Unregister a RTDM device

Removes the device from the RTDM namespace. This routine first attempts to teardown all active connections to the *device* prior to unregistering.

Parameters

in	<i>dev</i>	Device descriptor.
----	------------	--------------------

Tags

[secondary-only](#)

6.29.3.4 `rtm_drv_set_sysclass()`

```
int rtdm_drv_set_sysclass (
    struct rtdm_driver * drv,
    struct class * cls )
```

Set the kernel device class of a RTDM driver.

Set the kernel device class assigned to the RTDM driver. By default, RTDM drivers belong to Linux's "rtdm" device class, creating a device node hierarchy rooted at `/dev/rtdm`, and sysfs nodes under `/sys/class/rtdm`.

This call assigns a user-defined kernel device class to the RTDM driver, so that its devices are created into a different system hierarchy.

`rtdm_drv_set_sysclass()` is meaningful only before the first device which is attached to *drv* is registered by a call to `rtdm_dev_register()`.

Parameters

in	<i>drv</i>	Address of the RTDM driver descriptor.
in	<i>cls</i>	Pointer to the kernel device class. NULL is allowed to clear a previous setting, switching back to the default "rtdm" device class.

Returns

0 on success, otherwise:

- -EBUSY is returned if the kernel device class has already been set for *drv*, or some device(s) attached to *drv* are currently registered.

## Tags

[task-unrestricted](#)

## Attention

The kernel device class set by this call is not related to the RTDM class identification as defined by the [RTDM profiles](#) in any way. This is strictly related to the Linux kernel device hierarchy.

6.29.3.5 `rtdm_get_unmapped_area_handler()`

```
unsigned long rtdm_get_unmapped_area_handler (
    struct rtdm_fd * fd,
    unsigned long len,
    unsigned long pgoff,
    unsigned long flags )
```

Allocate mapping region in address space.

When present, this optional handler should return the start address of a free region in the process's address space, large enough to cover the ongoing `mmap()` operation. If unspecified, the default architecture-defined handler is invoked.

Most drivers can omit this handler, except on MMU-less platforms (see second note).

## Parameters

in	<i>fd</i>	File descriptor
in	<i>len</i>	Length of the requested region
in	<i>pgoff</i>	Page frame number to map to (see second note).
in	<i>flags</i>	Requested mapping flags

## Returns

The start address of the mapping region on success. On failure, a negative error code should be returned, with `-ENOSYS` meaning that the driver does not want to provide such information, in which case the ongoing `mmap()` operation will fail.

## Note

The address hint passed to the `mmap()` request is deliberately ignored by RTDM, and therefore not passed to this handler.

On MMU-less platforms, this handler is required because RTDM issues mapping requests over a shareable character device internally. In such context, the RTDM core may pass a null `pgoff` argument to the handler, for probing for the logical start address of the memory region to map to. Otherwise, when `pgoff` is non-zero, `pgoff << PAGE_SHIFT` is usually returned.

## 6.29.3.6 rtdm\_ioctl\_handler()

```
int rtdm_ioctl_handler (
    struct rtdm_fd * fd,
    unsigned int request,
    void __user * arg )
```

IOCTL handler.

Parameters

in	<i>fd</i>	File descriptor
in	<i>request</i>	Request number as passed by the user
in,out	<i>arg</i>	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>

## 6.29.3.7 rtdm\_mmap\_handler()

```
int rtdm_mmap_handler (
    struct rtdm_fd * fd,
    struct vm_area_struct * vma )
```

Memory mapping handler.

Parameters

in	<i>fd</i>	File descriptor
in	<i>vma</i>	Virtual memory area descriptor

Returns

0 on success. On failure, a negative error code is returned.

See also

mmap() in POSIX.1-2001, <http://pubs.opengroup.org/onlinepubs/7908799/xsh/mmap.html>

Note

The address hint passed to the mmap() request is deliberately ignored by RTDM.

6.29.3.8 `rtdm_open_handler()`

```
int rtdm_open_handler (
    struct rtdm_fd * fd,
    int oflags )
```

Open handler for named devices.

## Parameters

in	<i>fd</i>	File descriptor associated with opened device instance
in	<i>oflags</i>	Open flags as passed by the user

The file descriptor carries a device minor information which can be retrieved by a call to `rtdm_fd_get_minor(fd)`. The minor number can be used for distinguishing devices managed by a driver.

## Returns

0 on success. On failure, a negative error code is returned.

## See also

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

6.29.3.9 `rtdm_read_handler()`

```
ssize_t rtdm_read_handler (
    struct rtdm_fd * fd,
    void __user * buf,
    size_t size )
```

Read handler.

## Parameters

in	<i>fd</i>	File descriptor
out	<i>buf</i>	Input buffer as passed by the user
in	<i>size</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>

#### 6.29.3.10 `rtdm_recvmsg_handler()`

```
ssize_t rtdm_recvmsg_handler (
    struct rtdm_fd * fd,
    struct user_msghdr * msg,
    int flags )
```

Receive message handler.

Parameters

in	<i>fd</i>	File descriptor
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>

#### 6.29.3.11 `rtdm_select_handler()`

```
int rtdm_select_handler (
    struct rtdm_fd * fd,
    struct xnselector * selector,
    unsigned int type,
    unsigned int index )
```

Select handler.

Parameters

in	<i>fd</i>	File descriptor
	<i>selector</i>	Pointer to the selector structure
	<i>type</i>	Type of events ( <code>XNSELECT_READ</code> , <code>XNSELECT_WRITE</code> , or <code>XNSELECT_EXCEPT</code> )
	<i>index</i>	Index of the file descriptor

## Returns

0 on success. On failure, a negative error code is returned.

## See also

`select()` in POSIX.1-2001, <http://pubs.opengroup.org/onlinepubs/007908799/xsh/select.↵html>

6.29.3.12 `rtdm_sendmsg_handler()`

```
ssize_t rtdm_sendmsg_handler (
    struct rtdm_fd * fd,
    const struct user_msghdr * msg,
    int flags )
```

Transmit message handler.

## Parameters

in	<i>fd</i>	File descriptor
in	<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 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>

6.29.3.13 `rtdm_socket_handler()`

```
int rtdm_socket_handler (
    struct rtdm_fd * fd,
    int protocol )
```

Socket creation handler for protocol devices.

## Parameters

in	<i>fd</i>	File descriptor associated with opened device instance
in	<i>protocol</i>	Protocol number as passed by the user

## Returns

0 on success. On failure, a negative error code is returned.

## See also

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

6.29.3.14 `rtm_write_handler()`

```
ssize_t rtm_write_handler (  
    struct rtm_fd * fd,  
    const void __user * buf,  
    size_t size )
```

## Write handler.

## Parameters

in	<i>fd</i>	File descriptor
in	<i>buf</i>	Output buffer as passed by the user
in	<i>size</i>	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>

## 6.30 Clock Services

Collaboration diagram for Clock Services:



### Functions

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

*Get monotonic time*

#### 6.30.1 Detailed Description

#### 6.30.2 Function Documentation

##### 6.30.2.1 rtdm\_clock\_read()

```
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.

#### Tags

[unrestricted](#)

Referenced by [a4l\\_get\\_time\(\)](#).

## 6.30.2.2 rtdm\_clock\_read\_monotonic()

```
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.

Tags

[unrestricted](#)

## 6.31 Task Services

Collaboration diagram for 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*
- int [rtdm\\_task\\_should\\_stop](#) (void)  
*Check for pending termination request*
- 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\\_abs\\_t](#) start\_date, [nanosecs\\_rel\\_t](#) period)  
*Adjust real-time task period*
- int [rtdm\\_task\\_wait\\_period](#) (unsigned long \*overruns\_r)  
*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 `rtm_task_sleep_abs` (`nanosecs_abs_t` wakeup\_time, enum `rtm_timer_mode` mode)

*Sleep until a specified absolute time*

- int `rtm_task_busy_wait` (bool condition, `nanosecs_rel_t` spin\_ns, `nanosecs_rel_t` sleep\_ns)

*Safe busy waiting*

- void `rtm_wait_prepare` (struct `rtm_wait_context` \*wc)

*Register wait context.*

- void `rtm_wait_complete` (struct `rtm_wait_context` \*wc)

*Mark completion for a wait context.*

- int `rtm_wait_is_completed` (struct `rtm_wait_context` \*wc)

*Test completion of a wait context.*

- void `rtm_task_join` (`rtm_task_t` \*task)

*Wait on a real-time task to terminate*

- void `rtm_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 0`
- `#define RTDM_TASK_HIGHEST_PRIORITY 99`

## Task Priority Modification

Raise or lower task priorities by one level

- `#define RTDM_TASK_RAISE_PRIORITY (+1)`
- `#define RTDM_TASK_LOWER_PRIORITY (-1)`

### 6.31.1 Detailed Description

### 6.31.2 Typedef Documentation

#### 6.31.2.1 `rtm_task_proc_t`

```
typedef void(* rtm_task_proc_t) (void *arg)
```

Real-time task procedure.

## Parameters

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

## 6.31.3 Function Documentation

6.31.3.1 [rtm\\_task\\_busy\\_sleep\(\)](#)

```
void rtm_task_busy_sleep (
    nanosecs_rel_t delay )
```

Busy-wait a specified amount of time

This service does not schedule out the caller, but rather spins in a tight loop, burning CPU cycles until the timeout elapses.

## Parameters

in	<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.

## Tags

[unrestricted](#)

6.31.3.2 [rtm\\_task\\_busy\\_wait\(\)](#)

```
int rtm_task_busy_wait (
    bool condition,
    nanosecs_rel_t spin_ns,
    nanosecs_rel_t sleep_ns )
```

## Safe busy waiting

This service alternates active spinning and sleeping within a wait loop, until a condition is satisfied. While sleeping, a task is scheduled out and does not consume any CPU time.

[rtm\\_task\\_busy\\_wait\(\)](#) is particularly useful for waiting for a state change reading an I/O register, which usually happens shortly after the wait starts, without incurring the adverse effects of long busy waiting if it doesn't.



## Parameters

in	<i>condition</i>	The C expression to be tested for detecting completion.
in	<i>spin_ns</i>	The time to spin on <i>condition</i> before sleeping, expressed as a count of nanoseconds.
in	<i>sleep_ns</i>	The time to sleep for before spinning again, expressed as a count of nanoseconds.

## Returns

0 on success if *condition* is satisfied, otherwise:

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

## Tags

[primary-only](#), [might-switch](#)

## 6.31.3.3 rtm\_task\_current()

```
rtm_task_t* rtm_task_current (
    void )
```

Get current real-time task

## Returns

Pointer to task handle

## Tags

[mode-unrestricted](#)

## 6.31.3.4 rtm\_task\_destroy()

```
void rtm_task_destroy (
    rtm_task_t * task )
```

Destroy a real-time task

This call sends a termination request to *task*, then waits for it to exit. All RTDM task should check for pending termination requests by calling [rtm\\_task\\_should\\_stop\(\)](#) from their work loop.

If *task* is current, [rtm\\_task\\_destroy\(\)](#) terminates the current context, and does not return to the caller.

## Parameters

in,out	<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.

## Tags

[secondary-only](#), [might-switch](#)

## 6.31.3.5 rtdm\_task\_init()

```
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\(\)](#) was invoked.

## Parameters

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

## Returns

0 on success, otherwise negative error code

## Tags

[secondary-only](#), [might-switch](#)

6.31.3.6 `rtm_task_join()`

```
void rtdm_task_join (
    rtdm_task_t * task )
```

Wait on a real-time task to terminate

Parameters

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

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 [rtm\\_task\\_join\(\)](#) will never return.

Tags

[mode-unrestricted](#)

6.31.3.7 `rtm_task_set_period()`

```
int rtdm_task_set_period (
    rtdm_task_t * task,
    nanosecs_abs_t start_date,
    nanosecs_rel_t period )
```

Adjust real-time task period

Parameters

in,out	<i>task</i>	Task handle as returned by <a href="#">rtm_task_init()</a> , or NULL for referring to the current RTDM task or Cobalt thread.
in	<i>start_date</i>	The initial (absolute) date of the first release point, expressed in nanoseconds. <i>task</i> will be delayed by the first call to <a href="#">rtm_task_wait_period()</a> until this point is reached. If <i>start_date</i> is zero, the first release point is set to <i>period</i> nanoseconds after the current date.
in	<i>period</i>	New period in nanoseconds of a cyclic task, zero to disable cyclic mode for <i>task</i> .

Tags

[task-unrestricted](#)

6.31.3.8 `rtdm_task_set_priority()`

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

Adjust real-time task priority

Parameters

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

Tags

[task-unrestricted](#), [might-switch](#)

6.31.3.9 `rtdm_task_should_stop()`

```
int rtdm_task_should_stop (
    void )
```

Check for pending termination request

Check whether a termination request was received by the current RTDM task. Termination requests are sent by calling [rtdm\\_task\\_destroy\(\)](#).

Returns

Non-zero indicates that a termination request is pending, in which case the caller should wrap up and exit.

Tags

[rtdm-task](#), [might-switch](#)

6.31.3.10 `rtdm_task_sleep()`

```
int rtdm_task_sleep (
    nanosecs_rel_t delay )
```

Sleep a specified amount of time

## Parameters

in	<i>delay</i>	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.

## Tags

[primary-only](#), [might-switch](#)

6.31.3.11 `rtdm_task_sleep_abs()`

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

Sleep until a specified absolute time

## Parameters

in	<i>wakeup_time</i>	Absolute timeout in nanoseconds
in	<i>mode</i>	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.

## Tags

[primary-only](#), [might-switch](#)

6.31.3.12 `rtm_task_sleep_until()`

```
int rtdm_task_sleep_until (
    nanosecs_abs_t wakeup_time )
```

Sleep until a specified absolute time

**Deprecated** Use `rtm_task_sleep_abs` instead!

Parameters

in	<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 `rtm_task_unblock()`.
- -EPERM *may* be returned if an illegal invocation environment is detected.

Tags

[primary-only](#), [might-switch](#)

6.31.3.13 `rtm_task_unblock()`

```
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.

Tags

[unrestricted](#), [might-switch](#)

6.31.3.14 `rtm_task_wait_period()`

```
int rtdm_task_wait_period (
    unsigned long * overruns_r )
```

Wait on next real-time task period

## Parameters

in	<i>overruns</i> <i>_r</i>	Address of a long word receiving the count of overruns if -ETIMEDOUT is returned, or NULL if the caller don't need that information.
----	------------------------------	--

## 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.

## Tags

[primary-only](#), [might-switch](#)

6.31.3.15 `rtdm_wait_complete()`

```
void rtdm_wait_complete (
    struct rtdm_wait_context * wc )
```

Mark completion for a wait context.

`rtdm_complete_wait()` marks a wait context as completed, so that [rtdm\\_wait\\_is\\_completed\(\)](#) returns true for such context.

## Parameters

<i>wc</i>	Wait context to complete.
-----------	---------------------------

6.31.3.16 `rtdm_wait_is_completed()`

```
int rtdm_wait_is_completed (
    struct rtdm_wait_context * wc )
```

Test completion of a wait context.

[rtdm\\_wait\\_is\\_completed\(\)](#) returns true if `rtdm_complete_wait()` was called for *wc*. The completion mark is reset each time [rtdm\\_wait\\_prepare\(\)](#) is called for a wait context.

## Parameters

<i>wc</i>	Wait context to check for completion.
-----------	---------------------------------------

## Returns

non-zero/true if `rtdm_wait_complete()` was called for `wc`, zero otherwise.

6.31.3.17 `rtdm_wait_prepare()`

```
void rtdm_wait_prepare (
    struct rtdm_wait_context * wc )
```

Register wait context.

`rtdm_wait_prepare()` registers a wait context structure for the caller, which can be later retrieved by a call to `rtdm_wait_get_context()`. This call is normally issued before the current task blocks on a wait object, waiting for some (producer) code to wake it up. Arbitrary data can be exchanged between both sites via the wait context structure, which is allocated by the waiter (consumer) side.

`wc` is the address of an anchor object which is commonly embedded into a larger structure with arbitrary contents, which needs to be shared between the consumer (waiter) and the producer for implementing the wait code.

A typical implementation pattern for the wait side is:

```
struct rtdm_waitqueue wq;
struct some_wait_context {
    int input_value;
    int output_value;
    struct rtdm_wait_context wc;
} wait_context;

wait_context.input_value = 42;
rtdm_wait_prepare(&wait_context);
ret = rtdm_wait_condition(&wq, rtdm_wait_is_completed(&
    wait_context));
if (ret)
    goto wait_failed;
handle_event(wait_context.output_value);
```

On the producer side, the implementation would look like:

```
struct rtdm_waitqueue wq;
struct some_wait_context {
    int input_value;
    int output_value;
    struct rtdm_wait_context wc;
} *wait_context_ptr;
struct rtdm_wait_context *wc;
rtdm_task_t *task;

rtdm_for_each_waiter(task, &wq) {
    wc = rtdm_wait_get_context(task);
    wait_context_ptr = container_of(wc, struct some_wait_context, wc);
    wait_context_ptr->output_value = 12;
}
rtdm_waitqueue_broadcast(&wq);
```


## Parameters

<code>wc</code>	Wait context to register.
-----------------	---------------------------



## 6.32 Timer Services

Collaboration diagram for 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 }

### 6.32.1 Detailed Description

### 6.32.2 Typedef Documentation

#### 6.32.2.1 rtdm\_timer\_handler\_t

```
typedef void(* rtdm_timer_handler_t) (rtdm_timer_t *timer)
```

Timer handler.

Parameters

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

### 6.32.3 Enumeration Type Documentation

#### 6.32.3.1 rtdm\_timer\_mode

```
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.

### 6.32.4 Function Documentation

#### 6.32.4.1 rtdm\_timer\_destroy()

```
void rtdm_timer_destroy (
    rtdm_timer_t * timer )
```

Destroy a timer

Parameters

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

Tags

[task-unrestricted](#)

#### 6.32.4.2 rtdm\_timer\_init()

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

Initialise a timer

Parameters

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

Returns

0 on success, otherwise negative error code

Tags

[task-unrestricted](#)

#### 6.32.4.3 rtdm\_timer\_start()

```
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

in,out	<i>timer</i>	Timer handle as returned by <a href="#">rtm_timer_init()</a>
in	<i>expiry</i>	Firing time of the timer, <i>mode</i> defines if relative or absolute
in	<i>interval</i>	Relative reload value, > 0 if the timer shall work in periodic mode with the specific interval, 0 for one-shot timers
in	<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. In such an event, the timer is nevertheless armed for the next shot in the timeline if *interval* is non-zero.

## Tags

[unrestricted](#)

6.32.4.4 `rtm_timer_start_in_handler()`

```
int rtm_timer_start_in_handler (
    rtm_timer_t * timer,
    nanosecs_abs_t expiry,
    nanosecs_rel_t interval,
    enum rtm_timer_mode mode )
```

Start a timer from inside a timer handler

## Parameters

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

## Tags

[coreirq-only](#)

6.32.4.5 `rtm_timer_stop()`

```
void rtm_timer_stop (
    rtm_timer_t * timer )
```

Stop a timer

Parameters

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

Tags

[unrestricted](#)

6.32.4.6 `rtm_timer_stop_in_handler()`

```
void rtm_timer_stop_in_handler (
    rtm_timer_t * timer )
```

Stop a timer from inside a timer handler

Parameters

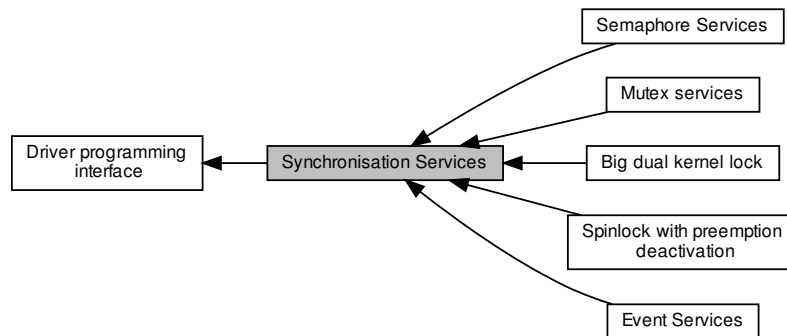
in,out	<i>timer</i>	Timer handle as returned by <a href="#">rtm_timer_init()</a>
--------	--------------	--

Tags

[coreirq-only](#)

## 6.33 Synchronisation Services

Collaboration diagram for Synchronisation Services:



### Modules

- [Big dual kernel lock](#)
- [Spinlock with preemption deactivation](#)
- [Event Services](#)
- [Semaphore Services](#)
- [Mutex services](#)

### Functions

- void [rtdm\\_waitqueue\\_init](#) (struct rtdm\_waitqueue \*wq)

*Initialize a RTDM wait queue*

- void [rtdm\\_waitqueue\\_destroy](#) (struct rtdm\_waitqueue \*wq)

*Deletes a RTDM wait queue*

- [rtdm\\_timedwait\\_condition\\_locked](#) (struct rtdm\_wait\_queue \*wq, C\_expr condition, [nanosecs\\_rel\\_t](#) timeout, rtdm\_toseq\_t \*toseq)

*Timed sleep on a locked waitqueue until a condition gets true*

- [rtdm\\_wait\\_condition\\_locked](#) (struct rtdm\_wait\_queue \*wq, C\_expr condition)

*Sleep on a locked waitqueue until a condition gets true*

- [rtdm\\_timedwait\\_condition](#) (struct rtdm\_wait\_queue \*wq, C\_expr condition, [nanosecs\\_rel\\_t](#) timeout, rtdm\_toseq\_t \*toseq)

*Timed sleep on a waitqueue until a condition gets true*

- void [rtdm\\_timedwait](#) (struct rtdm\_wait\_queue \*wq, [nanosecs\\_rel\\_t](#) timeout, rtdm\_toseq\_t \*toseq)

*Timed sleep on a waitqueue unconditionally*

- void [rtdm\\_timedwait\\_locked](#) (struct rtdm\_wait\_queue \*wq, [nanosecs\\_rel\\_t](#) timeout, rtdm\_toseq\_t \*toseq)

*Timed sleep on a locked waitqueue unconditionally*

- `rtdm_wait_condition` (struct `rtdm_wait_queue *wq`, C\_expr condition)

*Sleep on a waitqueue until a condition gets true*

- void `rtdm_wait` (struct `rtdm_wait_queue *wq`)

*Sleep on a waitqueue unconditionally*

- void `rtdm_wait_locked` (struct `rtdm_wait_queue *wq`)

*Sleep on a locked waitqueue unconditionally*

- void `rtdm_waitqueue_lock` (struct `rtdm_wait_queue *wq`, `rtdm_lockctx_t` context)

*Lock a waitqueue*

- void `rtdm_waitqueue_unlock` (struct `rtdm_wait_queue *wq`, `rtdm_lockctx_t` context)

*Unlock a waitqueue*

- void `rtdm_waitqueue_signal` (struct `rtdm_wait_queue *wq`)

*Signal a waitqueue*

- void `rtdm_waitqueue_broadcast` (struct `rtdm_wait_queue *wq`)

*Broadcast a waitqueue*

- void `rtdm_waitqueue_flush` (struct `rtdm_wait_queue *wq`)

*Flush a waitqueue*

- void `rtdm_waitqueue_wakeup` (struct `rtdm_wait_queue *wq`, `rtdm_task_t` waiter)

*Signal a particular waiter on a waitqueue*

- `rtdm_for_each_waiter` (`rtdm_task_t` pos, struct `rtdm_wait_queue *wq`)

*Simple iterator for waitqueues*

- `rtdm_for_each_waiter_safe` (`rtdm_task_t` pos, `rtdm_task_t` tmp, struct `rtdm_wait_queue *wq`)

*Safe iterator for waitqueues*

## 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 }

## Timeout Sequence Management

- void `rtdm_toseq_init` (`rtdm_toseq_t *timeout_seq`, `nanosecs_rel_t` timeout)

*Initialise a timeout sequence*

### 6.33.1 Detailed Description

### 6.33.2 Enumeration Type Documentation

#### 6.33.2.1 rtdm\_selecttype

enum [rtdm\\_selecttype](#)

Enumerator

RTDM_SELECTTYPE_READ	Select input data availability events.
RTDM_SELECTTYPE_WRITE	Select ouput buffer availability events.
RTDM_SELECTTYPE_EXCEPT	Select exceptional events.

### 6.33.3 Function Documentation

#### 6.33.3.1 rtdm\_for\_each\_waiter()

```
rtdm_for_each_waiter (
    rtdm_task_t pos,
    struct rtdm_wait_queue * wq )
```

Simple iterator for waitqueues

This construct traverses the wait list of a given waitqueue *wq*, assigning each RTDM task pointer to the cursor variable *pos*, which must be of type `rtdm_task_t`.

*wq* must have been locked by a call to [rtdm\\_waitqueue\\_lock\(\)](#) prior to traversing its wait list.

Parameters

<i>pos</i>	cursor variable holding a pointer to the RTDM task being fetched.
<i>wq</i>	waitqueue to scan.

Note

The waitqueue should not be signaled, broadcast or flushed during the traversal, unless the loop is aborted immediately after. Should multiple waiters be readied while iterating, the safe form [rtdm\\_for\\_each\\_waiter\\_safe\(\)](#) must be used for traversal instead.

Tags

[unrestricted](#)



6.33.3.2 `rtm_for_each_waiter_safe()`

```
rtm_for_each_waiter_safe (
    rtdm_task_t pos,
    rtdm_task_t tmp,
    struct rtdm_wait_queue * wq )
```

Safe iterator for waitqueues

This construct traverses the wait list of a given waitqueue `wq`, assigning each RTDM task pointer to the cursor variable `pos`, which must be of type `rtdm_task_t`.

Unlike with `rtm_for_each_waiter()`, the waitqueue may be signaled, broadcast or flushed during the traversal.

`wq` must have been locked by a call to `rtm_waitqueue_lock()` prior to traversing its wait list.

Parameters

<code>pos</code>	cursor variable holding a pointer to the RTDM task being fetched.
<code>tmp</code>	temporary cursor variable.
<code>wq</code>	waitqueue to scan.

Tags

[unrestricted](#)

6.33.3.3 `rtm_timedwait()`

```
void rtm_timedwait (
    struct rtdm_wait_queue * wq,
    nanosecs_rel_t timeout,
    rtdm_toseq_t * toseq )
```

Timed sleep on a waitqueue unconditionally

The calling task is put to sleep until the waitqueue is signaled by either `rtm_waitqueue_signal()` or `rtm_waitqueue_broadcast()`, or flushed by a call to `rtm_waitqueue_flush()`, or a timeout occurs.

Parameters

	<code>wq</code>	waitqueue to wait on.
	<code>timeout</code>	relative timeout in nanoseconds, see <code>RTDM_TIMEOUT_XXX</code> for special values.
<code>in,out</code>	<code>toseq</code>	handle of a timeout sequence as returned by <code>rtm_toseq_init()</code> or NULL.

## Returns

0 on success, otherwise:

- -EINTR is returned if the waitqueue has been flushed, or the calling task has received a Linux signal or has been forcibly unblocked by a call to [rtdm\\_task\\_unblock\(\)](#).
- -ETIMEDOUT is returned if the if the request has not been satisfied within the specified amount of time.

## Note

Passing RTDM\_TIMEOUT\_NONE to *timeout* makes no sense for such service, and might cause unexpected behavior.

## Tags

[primary-only](#), [might-switch](#)

## 6.33.3.4 rtdm\_timedwait\_condition()

```
rtdm_timedwait_condition (
    struct rtdm_wait_queue * wq,
    C_expr condition,
    nanosecs_rel_t timeout,
    rtdm_toseq_t * toseq )
```

Timed sleep on a waitqueue until a condition gets true

The calling task is put to sleep until *condition* evaluates to true or a timeout occurs. The condition is checked each time the waitqueue *wq* is signaled.

## Parameters

	<i>wq</i>	waitqueue to wait on.
	<i>condition</i>	C expression for the event to wait for.
	<i>timeout</i>	relative timeout in nanoseconds, see <a href="#">RTDM_TIMEOUT_xxx</a> for special values.
in, out	<i>toseq</i>	handle of a timeout sequence as returned by <a href="#">rtdm_toseq_init()</a> or NULL.

## Returns

0 on success, otherwise:

- -EINTR is returned if calling task has received a Linux signal or has been forcibly unblocked by a call to [rtdm\\_task\\_unblock\(\)](#).
- -ETIMEDOUT is returned if the if the request has not been satisfied within the specified amount of time.

## Note

[rtm\\_waitqueue\\_signal\(\)](#) has to be called after changing any variable that could change the result of the wait condition.

Passing `RTDM_TIMEOUT_NONE` to *timeout* makes no sense for such service, and might cause unexpected behavior.

## Tags

[primary-only](#), [might-switch](#)

## 6.33.3.5 rtdm\_timedwait\_condition\_locked()

```
rtdm_timedwait_condition_locked (
    struct rtdm_wait_queue * wq,
    C_expr condition,
    nanosecs_rel_t timeout,
    rtdm_toseq_t * toseq )
```

Timed sleep on a locked waitqueue until a condition gets true

The calling task is put to sleep until *condition* evaluates to true or a timeout occurs. The condition is checked each time the waitqueue *wq* is signaled.

The waitqueue must have been locked by a call to [rtm\\_waitqueue\\_lock\(\)](#) prior to calling this service.

## Parameters

	<i>wq</i>	locked waitqueue to wait on. The waitqueue lock is dropped when sleeping, then reacquired before this service returns to the caller.
	<i>condition</i>	C expression for the event to wait for.
	<i>timeout</i>	relative timeout in nanoseconds, see <a href="#">RTDM_TIMEOUT_xxx</a> for special values.
in, out	<i>toseq</i>	handle of a timeout sequence as returned by <a href="#">rtm_toseq_init()</a> or NULL.

## Returns

0 on success, otherwise:

- `-EINTR` is returned if calling task has received a Linux signal or has been forcibly unblocked by a call to [rtm\\_task\\_unblock\(\)](#).
- `-ETIMEDOUT` is returned if the request has not been satisfied within the specified amount of time.

## Note

[rtm\\_waitqueue\\_signal\(\)](#) has to be called after changing any variable that could change the result of the wait condition.

Passing `RTDM_TIMEOUT_NONE` to *timeout* makes no sense for such service, and might cause unexpected behavior.

Tags

[primary-only](#), [might-switch](#)

### 6.33.3.6 rtdm\_timedwait\_locked()

```
void rtdm_timedwait_locked (
    struct rtdm_wait_queue * wq,
    nanosecs_rel_t timeout,
    rtdm_toseq_t * toseq )
```

Timed sleep on a locked waitqueue unconditionally

The calling task is put to sleep until the waitqueue is signaled by either [rtdm\\_waitqueue\\_signal\(\)](#) or [rtdm\\_waitqueue\\_broadcast\(\)](#), or flushed by a call to [rtdm\\_waitqueue\\_flush\(\)](#), or a timeout occurs.

The waitqueue must have been locked by a call to [rtdm\\_waitqueue\\_lock\(\)](#) prior to calling this service.

Parameters

	<i>wq</i>	locked waitqueue to wait on. The waitqueue lock is dropped when sleeping, then reacquired before this service returns to the caller.
	<i>timeout</i>	relative timeout in nanoseconds, see <a href="#">RTDM_TIMEOUT_XXX</a> for special values.
in, out	<i>toseq</i>	handle of a timeout sequence as returned by <a href="#">rtdm_toseq_init()</a> or NULL.

Returns

0 on success, otherwise:

- -EINTR is returned if the waitqueue has been flushed, or the calling task has received a Linux signal or has been forcibly unblocked by a call to [rtdm\\_task\\_unblock\(\)](#).
- -ETIMEDOUT is returned if the if the request has not been satisfied within the specified amount of time.

Note

Passing [RTDM\\_TIMEOUT\\_NONE](#) to *timeout* makes no sense for such service, and might cause unexpected behavior.

Tags

[primary-only](#), [might-switch](#)

6.33.3.7 `rtm_toseq_init()`

```
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.

Tags

[task-unrestricted](#)

6.33.3.8 `rtm_wait()`

```
void rtdm_wait (
    struct rtdm_wait_queue * wq )
```

Sleep on a waitqueue unconditionally

The calling task is put to sleep until the waitqueue is signaled by either `rtdm_waitqueue_signal()` or `rtdm_waitqueue_broadcast()`, or flushed by a call to `rtdm_waitqueue_flush()`.

## Parameters

<i>wq</i>	waitqueue to wait on.
-----------	-----------------------

## Returns

0 on success, otherwise:

- -EINTR is returned if the waitqueue has been flushed, or the calling task has received a Linux signal or has been forcibly unblocked by a call to [rtm\\_task\\_unblock\(\)](#).

## Tags

[primary-only](#), [might-switch](#)

## 6.33.3.9 rtdm\_wait\_condition()

```
rtdm_wait_condition (
    struct rtdm_wait_queue * wq,
    C_expr condition )
```

Sleep on a waitqueue until a condition gets true

The calling task is put to sleep until *condition* evaluates to true. The condition is checked each time the waitqueue *wq* is signaled.

## Parameters

<i>wq</i>	waitqueue to wait on
<i>condition</i>	C expression for the event to wait for.

## Returns

0 on success, otherwise:

- -EINTR is returned if calling task has received a Linux signal or has been forcibly unblocked by a call to [rtm\\_task\\_unblock\(\)](#).

## Note

[rtdm\\_waitqueue\\_signal\(\)](#) has to be called after changing any variable that could change the result of the wait condition.

## Tags

[primary-only](#), [might-switch](#)

6.33.3.10 `rtm_wait_condition_locked()`

```
rtm_wait_condition_locked (
    struct rtdm_wait_queue * wq,
    C_expr condition )
```

Sleep on a locked waitqueue until a condition gets true

The calling task is put to sleep until *condition* evaluates to true. The condition is checked each time the waitqueue *wq* is signaled.

The waitqueue must have been locked by a call to [rtm\\_waitqueue\\_lock\(\)](#) prior to calling this service.

## Parameters

<i>wq</i>	locked waitqueue to wait on. The waitqueue lock is dropped when sleeping, then reacquired before this service returns to the caller.
<i>condition</i>	C expression for the event to wait for.

## Returns

0 on success, otherwise:

- -EINTR is returned if calling task has received a Linux signal or has been forcibly unblocked by a call to [rtm\\_task\\_unblock\(\)](#).

## Note

[rtm\\_waitqueue\\_signal\(\)](#) has to be called after changing any variable that could change the result of the wait condition.

## Tags

[primary-only](#), [might-switch](#)

6.33.3.11 `rtm_wait_locked()`

```
void rtdm_wait_locked (
    struct rtdm_wait_queue * wq )
```

Sleep on a locked waitqueue unconditionally

The calling task is put to sleep until the waitqueue is signaled by either [rtm\\_waitqueue\\_signal\(\)](#) or [rtm\\_waitqueue\\_broadcast\(\)](#), or flushed by a call to [rtm\\_waitqueue\\_flush\(\)](#).

The waitqueue must have been locked by a call to [rtm\\_waitqueue\\_lock\(\)](#) prior to calling this service.

## Parameters

<code>wq</code>	locked waitqueue to wait on. The waitqueue lock is dropped when sleeping, then reacquired before this service returns to the caller.
-----------------	--

## Returns

0 on success, otherwise:

- -EINTR is returned if the waitqueue has been flushed, or the calling task has received a Linux signal or has been forcibly unblocked by a call to [rtm\\_task\\_unblock\(\)](#).

## Tags

[primary-only](#), [might-switch](#)

6.33.3.12 `rtm_waitqueue_broadcast()`

```
void rtm_waitqueue_broadcast (
    struct rtm_wait_queue * wq )
```

## Broadcast a waitqueue

Broadcast the waitqueue `wq`, waking up all waiters. Each readied task may assume to have received the wake up event.

## Parameters

<code>wq</code>	waitqueue to broadcast.
-----------------	-------------------------

## Returns

non-zero if at least one task has been readied as a result of this call, zero otherwise.

## Tags

[unrestricted](#), [might-switch](#)

6.33.3.13 `rtm_waitqueue_destroy()`

```
void rtm_waitqueue_destroy (
    struct rtm_waitqueue * wq )
```

## Deletes a RTDM wait queue

Dismantles a wait queue structure, releasing all resources attached to it.



Parameters

<i>wq</i>	waitqueue to delete.
-----------	----------------------

Tags

[task-unrestricted](#)

#### 6.33.3.14 `rtm_waitqueue_flush()`

```
void rtm_waitqueue_flush (
    struct rtm_wait_queue * wq )
```

Flush a waitqueue

Flushes the waitqueue *wq*, unblocking all waiters with an error status (-EINTR).

Parameters

<i>wq</i>	waitqueue to flush.
-----------	---------------------

Returns

non-zero if at least one task has been readied as a result of this call, zero otherwise.

Tags

[unrestricted](#), [might-switch](#)

#### 6.33.3.15 `rtm_waitqueue_init()`

```
void rtm_waitqueue_init (
    struct rtm_waitqueue * wq )
```

Initialize a RTDM wait queue

Sets up a wait queue structure for further use.

Parameters

<i>wq</i>	waitqueue to initialize.
-----------	--------------------------

Tags

[task-unrestricted](#)

### 6.33.3.16 rtdm\_waitqueue\_lock()

```
void rtdm_waitqueue_lock (
    struct rtdm_wait_queue * wq,
    rtdm_lockctx_t context )
```

Lock a waitqueue

Acquires the lock on the waitqueue *wq*.

Parameters

<i>wq</i>	waitqueue to lock.
<i>context</i>	name of local variable to store the context in.

Note

Recursive locking might lead to unexpected behavior, including lock up.

Tags

[unrestricted](#)

### 6.33.3.17 rtdm\_waitqueue\_signal()

```
void rtdm_waitqueue_signal (
    struct rtdm_wait_queue * wq )
```

Signal a waitqueue

Signals the waitqueue *wq*, waking up a single waiter (if any).

Parameters

<i>wq</i>	waitqueue to signal.
-----------	----------------------

Returns

non-zero if a task has been readied as a result of this call, zero otherwise.

Tags

[unrestricted](#), [might-switch](#)

#### 6.33.3.18 `rtm_waitqueue_unlock()`

```
void rtdm_waitqueue_unlock (
    struct rtdm_wait_queue * wq,
    rtdm_lockctx_t context )
```

Unlock a waitqueue

Releases the lock on the waitqueue *wq*.

Parameters

<i>wq</i>	waitqueue to unlock.
<i>context</i>	name of local variable to retrieve the context from.

Tags

[unrestricted](#)

#### 6.33.3.19 `rtm_waitqueue_wakeup()`

```
void rtdm_waitqueue_wakeup (
    struct rtdm_wait_queue * wq,
    rtdm_task_t waiter )
```

Signal a particular waiter on a waitqueue

Signals the waitqueue *wq*, waking up waiter *waiter* only, which must be currently sleeping on the waitqueue.

Parameters


<i>wq</i>	waitqueue to signal.
<i>waiter</i>	RTDM task to wake up.

Tags

[unrestricted](#), [might-switch](#)

## 6.34 Event Services

Collaboration diagram for Event Services:



### Functions

- void [rtdm\\_event\\_init](#) (rtdm\_event\_t \*event, unsigned long pending)

*Initialise an event*

- void [rtdm\\_event\\_destroy](#) (rtdm\_event\_t \*event)

*Destroy an event*

- void [rtdm\\_event\\_pulse](#) (rtdm\_event\_t \*event)

*Signal an event occurrence to currently listening waiters*

- void [rtdm\\_event\\_signal](#) (rtdm\_event\_t \*event)

*Signal an event occurrence*

- int [rtdm\\_event\\_wait](#) (rtdm\_event\_t \*event)

*Wait on event occurrence*

- int [rtdm\\_event\\_timedwait](#) (rtdm\_event\_t \*event, [nanosecs\\_rel\\_t](#) timeout, rtdm\_toseq\_t \*timeout←\_seq)

*Wait on event occurrence with timeout*

- void [rtdm\\_event\\_clear](#) (rtdm\_event\_t \*event)

*Clear event state*

- int [rtdm\\_event\\_select](#) (rtdm\_event\_t \*event, rtdm\_selector\_t \*selector, enum [rtdm\\_selecttype](#) type, unsigned int fd\_index)

*Bind a selector to an event*

### 6.34.1 Detailed Description

### 6.34.2 Function Documentation

## 6.34.2.1 rtdm\_event\_clear()

```
void rtdm_event_clear (  
    rtdm_event_t * event )
```

Clear event state

## Parameters

in,out	<i>event</i>	Event handle as returned by <a href="#">rtm_event_init()</a>
--------	--------------	--

## Tags

[unrestricted](#)

## 6.34.2.2 rtdm\_event\_destroy()

```
void rtdm_event_destroy (
    rtdm_event_t * event )
```

## Destroy an event

## Parameters

in,out	<i>event</i>	Event handle as returned by <a href="#">rtm_event_init()</a>
--------	--------------	--

## Tags

[task-unrestricted](#), [might-switch](#)

## 6.34.2.3 rtdm\_event\_init()

```
void rtdm_event_init (
    rtdm_event_t * event,
    unsigned long pending )
```

## Initialise an event

## Parameters

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

## Tags

[task-unrestricted](#)

6.34.2.4 `rtm_event_pulse()`

```
void rtm_event_pulse (
    rtm_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 `rtm_event_wait()` or `rtm_event_timedwait()` will therefore be blocked first.

Parameters

in, out	<i>event</i>	Event handle as returned by <code>rtm_event_init()</code>
---------	--------------	---

Tags

[unrestricted](#), [might-switch](#)

6.34.2.5 `rtm_event_select()`

```
int rtm_event_select (
    rtm_event_t * event,
    rtm_selector_t * selector,
    enum rtm_selecttype type,
    unsigned int 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

in, out	<i>event</i>	Event handle as returned by <code>rtm_event_init()</code>
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.

Tags

[task-unrestricted](#)

#### 6.34.2.6 rtdm\_event\_signal()

```
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

in,out	<i>event</i>	Event handle as returned by <a href="#">rtdm_event_init()</a>
--------	--------------	---

Tags

[unrestricted](#), [might-switch](#)

#### 6.34.2.7 rtdm\_event\_timedwait()

```
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

in,out	<i>event</i>	Event handle as returned by <a href="#">rtdm_event_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="#">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 [rtm\\_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.

## Tags

[primary-timed](#), [might-switch](#)

## 6.34.2.8 rtdm\_event\_wait()

```
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

in, out	<i>event</i>	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 [rtm\\_task\\_unblock\(\)](#).
- -EIDRM is returned if *event* has been destroyed.
- -EPERM *may* be returned if an illegal invocation environment is detected.

## Tags

[primary-only](#), [might-switch](#)

## 6.35 Semaphore Services

Collaboration diagram for Semaphore Services:



### Functions

- void [rtm\\_sem\\_init](#) (rtm\_sem\_t \*sem, unsigned long value)

*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*

- 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*

- int [rtm\\_sem\\_select](#) (rtm\_sem\_t \*sem, rtm\_selector\_t \*selector, enum [rtm\\_selecttype](#) type, unsigned int fd\_index)

*Bind a selector to a semaphore*

### 6.35.1 Detailed Description

### 6.35.2 Function Documentation

#### 6.35.2.1 rtm\_sem\_destroy()

```
void rtm_sem_destroy (
    rtm_sem_t * sem )
```

Destroy a semaphore

Parameters

in,out	<i>sem</i>	Semaphore handle as returned by <a href="#">rt dm_sem_init()</a>
--------	------------	--

Tags

[task-unrestricted](#), [might-switch](#)

### 6.35.2.2 rtdm\_sem\_down()

```
int rtdm_sem_down (
    rtdm_sem_t * sem )
```

Decrement a semaphore

This is the light-weight version of [rtdm\\_sem\\_timeddown\(\)](#), implying an infinite timeout.

Parameters

in,out	<i>sem</i>	Semaphore handle as returned by <a href="#">rt dm_sem_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 *sem* has been destroyed.
- -EPERM *may* be returned if an illegal invocation environment is detected.

Tags

[primary-only](#), [might-switch](#)

### 6.35.2.3 rtdm\_sem\_init()

```
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

## Tags

[task-unrestricted](#)

## 6.35.2.4 rtdm\_sem\_select()

```
int rtdm_sem_select (
    rtdm_sem_t * sem,
    rtdm_selector_t * selector,
    enum rtdm_selecttype type,
    unsigned int 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.

## Tags

[task-unrestricted](#)

6.35.2.5 `rtdm_sem_timeddown()`

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

Tags

[primary-timed](#), [might-switch](#)

6.35.2.6 `rtdm_sem_up()`

```
void rtdm_sem_up (
    rtdm_sem_t * sem )
```

Increment a semaphore

This function increments the given semaphore's value, waking up a potential waiter which was blocked upon [rtdm\\_sem\\_down\(\)](#).

## Parameters


in,out	<i>sem</i>	Semaphore handle as returned by <a href="#">rt dm_sem_init()</a>
--------	------------	--

## Tags

[unrestricted](#), [might-switch](#)

## 6.36 Mutex services

Collaboration diagram for Mutex services:



### Functions

- void `rtm_mutex_init` (`rtm_mutex_t *mutex`)

*Initialise a mutex*

- void `rtm_mutex_destroy` (`rtm_mutex_t *mutex`)

*Destroy a mutex*

- void `rtm_mutex_unlock` (`rtm_mutex_t *mutex`)

*Release 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*

### 6.36.1 Detailed Description

### 6.36.2 Function Documentation

#### 6.36.2.1 `rtm_mutex_destroy()`

```
void rtm_mutex_destroy (
    rtm_mutex_t * mutex )
```

Destroy a mutex

## Parameters

in,out	<i>mutex</i>	Mutex handle as returned by <a href="#">rt dm_mutex_init()</a>
--------	--------------	--

## Tags

[task-unrestricted](#), [might-switch](#)

6.36.2.2 [rt dm\\_mutex\\_init\(\)](#)

```
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

in,out	<i>mutex</i>	Mutex handle
--------	--------------	--------------

## Tags

[task-unrestricted](#)

6.36.2.3 [rt dm\\_mutex\\_lock\(\)](#)

```
int rtdm_mutex_lock (  
    rtdm_mutex_t * mutex )
```

## Request a mutex

This is the light-weight version of [rt dm\\_mutex\\_timedlock\(\)](#), implying an infinite timeout.

## Parameters

in,out	<i>mutex</i>	Mutex handle as returned by <a href="#">rt dm_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.

## Tags

[primary-only](#), [might-switch](#)

## 6.36.2.4 rtdm\_mutex\_timedlock()

```
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

in,out	<i>mutex</i>	Mutex handle as returned by <a href="#">rtdm_mutex_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="#">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.

## Tags

[primary-only](#), [might-switch](#)

6.36.2.5 `rtm_mutex_unlock()`

```
void rtm_mutex_unlock (  
    rtm_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


in, out	<i>mutex</i>	Mutex handle as returned by <a href="#">rtm_mutex_init()</a>
---------	--------------	--

## Tags

[primary-only](#), [might-switch](#)

## 6.37 Interrupt Management Services

Collaboration diagram for 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_IRQTYPE_SHARED`

*Enable IRQ-sharing with other real-time drivers.*

- `#define RTDM\_IRQTYPE\_EDGE XN_IRQTYPE_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\_IRQ\_NONE  
*Unhandled interrupt.*
- #define [RTDM\\_IRQ\\_HANDLED](#) XN\_IRQ\_HANDLED  
*Denote handled interrupt.*
- #define [RTDM\\_IRQ\\_DISABLE](#) XN\_IRQ\_DISABLE  
*Request interrupt disabling on exit.*

### 6.37.1 Detailed Description

### 6.37.2 Macro Definition Documentation

#### 6.37.2.1 rtdm\_irq\_get\_arg

```
#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 [rtdm\\_irq\\_request\(\)](#) is returned, type-casted to the specified *type*.

Tags

[unrestricted](#)

### 6.37.3 Typedef Documentation

#### 6.37.3.1 rtdm\_irq\_handler\_t

```
typedef int(* rtdm_irq_handler_t) (rtdm_irq_t *irq_handle)
```

Interrupt handler.

Parameters

in	<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

## 6.37.4 Function Documentation

### 6.37.4.1 [rtm\\_irq\\_disable\(\)](#)

```
int rtm_irq_disable (  
    rtm_irq_t * irq_handle )
```

Disable interrupt line

Parameters

in, out	<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, such service is solely available from secondary mode. The caller is responsible for excluding such conflicts.

Tags

[secondary-only](#)

### 6.37.4.2 [rtm\\_irq\\_enable\(\)](#)

```
int rtm_irq_enable (  
    rtm_irq_t * irq_handle )
```

Enable interrupt line

## Parameters

in,out	<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, such service is solely available from secondary mode. The caller is responsible for excluding such conflicts.

## Tags

[secondary-only](#)

6.37.4.3 `rtm_irq_free()`

```
int rtm_irq_free (
    rtm_irq_t * irq_handle )
```

## Release an interrupt handler

## Parameters

in,out	<i>irq_handle</i>	IRQ handle as returned by <a href="#">rtm_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.

## Tags

[secondary-only](#)

## 6.37.4.4 rtdm\_irq\_request()

```
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.
- -ENOSYS is returned if the real-time core is disabled.


## Tags

[secondary-only](#)

## 6.38 Non-Real-Time Signalling Services

These services provide a mechanism to request the execution of a specified handler in non-real-time context.

Collaboration diagram for 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

- void [rtdm\\_schedule\\_nrt\\_work](#) (struct work\_struct \*lostage\_work)  
*Put a work task in Linux non real-time global workqueue from primary mode.*
- 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*

#### 6.38.1 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.

#### 6.38.2 Typedef Documentation

##### 6.38.2.1 rtdm\_nrtsig\_handler\_t

```
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 pointer as passed to <a href="#">rtm_nrtsig_init()</a>
in	<i>arg</i>	Argument as passed to <a href="#">rtm_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.

## 6.38.3 Function Documentation

6.38.3.1 `rtm_nrtsig_destroy()`

```
void rtm_nrtsig_destroy (
    rtm_nrtsig_t * nrt_sig )
```

Release a non-realtime signal handler

## Parameters

in,out	<i>nrt_sig</i>	Signal handle
--------	----------------	---------------

## Tags

[task-unrestricted](#)

6.38.3.2 `rtm_nrtsig_init()`

```
int rtm_nrtsig_init (
    rtm_nrtsig_t * nrt_sig,
    rtm_nrtsig_handler_t handler,
    void * arg )
```

Register a non-real-time signal handler

## Parameters

in,out	<i>nrt_sig</i>	Signal handle
in	<i>handler</i>	Non-real-time signal handler
in	<i>arg</i>	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.

Tags

[task-unrestricted](#)

### 6.38.3.3 rtdm\_nrtsig\_pend()

```
void rtdm_nrtsig_pend (  
    rtdm_nrtsig_t * nrt_sig )
```

Trigger non-real-time signal

Parameters

in,out	<i>nrt_sig</i>	Signal handle
--------	----------------	---------------

Tags

[unrestricted](#)

### 6.38.3.4 rtdm\_schedule\_nrt\_work()

```
void rtdm_schedule_nrt_work (  
    struct work_struct * lostage_work )
```


Put a work task in Linux non real-time global workqueue from primary mode.

Parameters

<i>lostage_work</i>	
---------------------	--

## 6.39 Utility Services

Collaboration diagram for Utility Services:



### Functions

- int `rtdm_mmap_to_user` (struct `rtdm_fd` \*fd, 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` (struct `rtdm_fd` \*fd, `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_mmap_kmem` (struct `vm_area_struct` \*vma, void \*va)

*Map a kernel logical memory range to a virtual user area.*

- int `rtdm_mmap_vmem` (struct `vm_area_struct` \*vma, void \*va)

*Map a virtual memory range to a virtual user area.*

- int `rtdm_mmap_iomem` (struct `vm_area_struct` \*vma, `phys_addr_t` pa)

*Map an I/O memory range to a virtual user area.*

- int `rtdm_munmap` (void \*ptr, size\_t len)

*Unmap a user memory range.*

- 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*

- void `rtdm_free` (void \*ptr)

*Release real-time memory block*

- int `rtdm_read_user_ok` (struct `rtdm_fd` \*fd, const void \_\_user \*ptr, size\_t size)

*Check if read access to user-space memory block is safe*

- int `rtdm_rw_user_ok` (struct `rtdm_fd` \*fd, const void \_\_user \*ptr, size\_t size)

*Check if read/write access to user-space memory block is safe*

- int [rtm\\_copy\\_from\\_user](#) (struct rtdm\_fd \*fd, void \*dst, const void \_\_user \*src, size\_t size)

*Copy user-space memory block to specified buffer*

- int [rtm\\_safe\\_copy\\_from\\_user](#) (struct rtdm\_fd \*fd, 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](#) (struct rtdm\_fd \*fd, void \_\_user \*dst, const void \*src, size\_t size)

*Copy specified buffer to user-space memory block*

- int [rtm\\_safe\\_copy\\_to\\_user](#) (struct rtdm\_fd \*fd, 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](#) (struct rtdm\_fd \*fd, 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](#) (struct rtdm\_fd \*fd)

*Test if the caller is capable of running in real-time context*

- bool [rtm\\_available](#) (void)

*Test if the real-time core is available*

### 6.39.1 Detailed Description

### 6.39.2 Function Documentation

#### 6.39.2.1 rtm\_available()

```
bool rtm_available (
    void )
```

Test if the real-time core is available

Returns

True if the real-time is available, false if it is disabled or in error state.

Note

Drivers should query the core state during initialization if they perform hardware setup operations or interact with RTDM services such as locks prior to calling an RTDM service that has a built-in state check of the real-time core (e.g. [rtm\\_dev\\_register\(\)](#) or [rtm\\_task\\_init\(\)](#)).

Tags

[unrestricted](#)

Referenced by [a4l\\_register\\_drv\(\)](#).

6.39.2.2 `rtdm_copy_from_user()`

```
int rtdm_copy_from_user (
    struct rtdm_fd * fd,
    void * dst,
    const void __user * src,
    size_t size )
```

Copy user-space memory block to specified buffer

## Parameters

in	<i>fd</i>	RTDM file descriptor as passed to the invoked device operation handler
in	<i>dst</i>	Destination buffer address
in	<i>src</i>	Address of the user-space memory block
in	<i>size</i>	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.

## Tags

[task-unrestricted](#)

6.39.2.3 `rtdm_copy_to_user()`

```
int rtdm_copy_to_user (
    struct rtdm_fd * fd,
    void __user * dst,
    const void * src,
    size_t size )
```

Copy specified buffer to user-space memory block

## Parameters

in	<i>fd</i>	RTDM file descriptor as passed to the invoked device operation handler
in	<i>dst</i>	Address of the user-space memory block
in	<i>src</i>	Source buffer address
in	<i>size</i>	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 [rtm\\_rw\\_user\\_ok\(\)](#) that the provided user-space address can securely be accessed.

## Tags

[task-unrestricted](#)

## 6.39.2.4 rtdm\_free()

```
void rtdm_free (  
    void * ptr )
```

Release real-time memory block

## Parameters

in	<i>ptr</i>	Pointer to memory block as returned by <a href="#">rtdm_malloc()</a>
----	------------	--

## Tags

[unrestricted](#)

## 6.39.2.5 rtdm\_in\_rt\_context()

```
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.

## Tags

[unrestricted](#)

6.39.2.6 `rtdm_iomap_to_user()`

```
int rtdm_iomap_to_user (
    struct rtdm_fd * fd,
    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>fd</i>	RTDM file descriptor 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 <code>PROT_READ</code> or <code>PROT_READ PROT_WRITE</code>
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>	<code>vm_operations</code> to be executed on the <code>vm_area</code> of the user memory range or NULL
in	<i>vm_private_data</i>	Private data to be stored in the <code>vm_area</code> , primarily useful for <code>vm_operation</code> 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 memory area:

- manual unmapping via `rtdm_munmap()`, which may be issued from a driver in response to an `IOCTL` call, or by a call to the regular `munmap()` call from the application.
- automatic unmapping, triggered by the termination of the process which owns the mapping. To track the number of references pending on the resource mapped, the driver can pass the address of a close handler for the `vm_area` considered, in the `vm_ops` descriptor. See the relevant Linux kernel programming documentation (e.g. Linux Device Drivers book) on virtual memory management for details.

## Tags

[secondary-only](#)

6.39.2.7 `rtdm_malloc()`

```
void* rtdm_malloc (
    size_t size )
```

Allocate memory block

Parameters

in	<i>size</i>	Requested size of the memory block
----	-------------	------------------------------------

Returns

The pointer to the allocated block is returned on success, NULL otherwise.

Tags

[unrestricted](#)

Referenced by `a4l_alloc_subd()`.

6.39.2.8 `rtdm_mmap_iomem()`

```
int rtdm_mmap_iomem (
    struct vm_area_struct * vma,
    phys_addr_t pa )
```

Map an I/O memory range to a virtual user area.

This routine is commonly used from a `->mmap()` handler of a RTDM driver, for mapping an I/O memory area over the user address space referred to by *vma*.

Parameters

in	<i>vma</i>	The VMA descriptor to receive the mapping.
in	<i>pa</i>	The physical I/O address to be mapped.

Returns

0 on success, otherwise a negated error code is returned.

Note

To map a chunk of logical space obtained from `kmalloc()`, or a purely virtual area with no direct physical mapping to a VMA, call `rtdm_mmap_kmem()` or `rtdm_mmap_vmem()` respectively instead.



Tags

[secondary-only](#)

### 6.39.2.9 rtdm\_mmap\_kmem()

```
int rtdm_mmap_kmem (
    struct vm_area_struct * vma,
    void * va )
```

Map a kernel logical memory range to a virtual user area.

This routine is commonly used from a `->mmap()` handler of a RTDM driver, for mapping a virtual memory area with a direct physical mapping over the user address space referred to by `vma`.

Parameters

in	<code>vma</code>	The VMA descriptor to receive the mapping.
in	<code>va</code>	The kernel logical address to be mapped.

Returns

0 on success, otherwise a negated error code is returned.

Note

This service works on memory regions allocated via `kmalloc()`. To map a chunk of virtual space with no direct physical mapping, or a physical I/O memory to a VMA, call [rtdm\\_mmap\\_vmem\(\)](#) or [rtdm\\_mmap\\_iomem\(\)](#) respectively instead.

Tags

[secondary-only](#)

### 6.39.2.10 rtdm\_mmap\_to\_user()

```
int rtdm_mmap_to_user (
    struct rtdm_fd * fd,
    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

in	<i>fd</i>	RTDM file descriptor as passed to the invoked device operation handler
in	<i>src_addr</i>	Kernel virtual 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 vm_area of the user memory range or NULL
in	<i>vm_private_data</i>	Private data to be stored in the vm_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 memory area:

- manual unmapping via `rtdm_munmap()`, which may be issued from a driver in response to an IOCTL call, or by a call to the regular `munmap()` call from the application.
- automatic unmapping, triggered by the termination of the process which owns the mapping. To track the number of references pending on the resource mapped, the driver can pass the address of a close handler for the vm\_area considered, in the `vm_ops` descriptor. See the relevant Linux kernel programming documentation (e.g. Linux Device Drivers book) on virtual memory management for details.

## Tags

[secondary-only](#)

## 6.39.2.11 rtdm\_mmap\_vmem()

```
int rtdm_mmap_vmem (
    struct vm_area_struct * vma,
    void * va )
```

Map a virtual memory range to a virtual user area.

This routine is commonly used from a `->mmap()` handler of a RTDM driver, for mapping a purely virtual memory area over the user address space referred to by `vma`.

## Parameters

in	<i>vma</i>	The VMA descriptor to receive the mapping.
in	<i>va</i>	The virtual address to be mapped.

## Returns

0 on success, otherwise a negated error code is returned.

## Note

This service works on memory regions allocated via `vmalloc()`. To map a chunk of logical space obtained from `kmalloc()`, or a physical I/O memory to a VMA, call `rdm_mmap_kmem()` or `rdm_mmap_iomem()` respectively instead.

## Tags

[secondary-only](#)

6.39.2.12 `rdm_munmap()`

```
int rdm_munmap (
    void * ptr,
    size_t len )
```

Unmap a user memory range.

## Parameters

in	<i>ptr</i>	User address or the memory range
in	<i>len</i>	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.

## Tags

[secondary-only](#)

## 6.39.2.13 rtdm\_printk()

```
void rtdm_printk (
    const char * format,
    ... )
```

Real-time safe message printing on kernel console

Parameters

in	<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.

Tags

[unrestricted](#)

## 6.39.2.14 rtdm\_printk\_ratelimited()

```
void rtdm_printk_ratelimited (
    const char * format,
    ... )
```

Real-time safe rate-limited message printing on kernel console

Parameters

in	<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.

Tags

[unrestricted](#)

## 6.39.2.15 rtdm\_read\_user\_ok()

```
int rtdm_read_user_ok (
    struct rtdm_fd * fd,
    const void __user * ptr,
    size_t size )
```

Check if read access to user-space memory block is safe

## Parameters

in	<i>fd</i>	RTDM file descriptor as passed to the invoked device operation handler
in	<i>ptr</i>	Address of the user-provided memory block
in	<i>size</i>	Size of the memory block

## Returns

Non-zero is return when it is safe to read from the specified memory block, 0 otherwise.

## Tags

[task-unrestricted](#)

## 6.39.2.16 rtdm\_rt\_capable()

```
int rtdm_rt_capable (
    struct rtdm_fd * fd )
```

Test if the caller is capable of running in real-time context

## Parameters

in	<i>fd</i>	RTDM file descriptor 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.

Tags

[unrestricted](#)

### 6.39.2.17 rtdm\_rw\_user\_ok()

```
int rtdm_rw_user_ok (
    struct rtdm_fd * fd,
    const void __user * ptr,
    size_t size )
```

Check if read/write access to user-space memory block is safe

Parameters

in	<i>fd</i>	RTDM file descriptor as passed to the invoked device operation handler
in	<i>ptr</i>	Address of the user-provided memory block
in	<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.

Tags

[task-unrestricted](#)

### 6.39.2.18 rtdm\_safe\_copy\_from\_user()

```
int rtdm_safe_copy_from_user (
    struct rtdm_fd * fd,
    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

in	<i>fd</i>	RTDM file descriptor as passed to the invoked device operation handler
in	<i>dst</i>	Destination buffer address
in	<i>src</i>	Address of the user-space memory block
in	<i>size</i>	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 `rtdm_read_user_ok` and `rtdm_copy_from_user`.

## Tags

[task-unrestricted](#)

6.39.2.19 `rtdm_safe_copy_to_user()`

```
int rtdm_safe_copy_to_user (
    struct rtdm_fd * fd,
    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

in	<i>fd</i>	RTDM file descriptor as passed to the invoked device operation handler
in	<i>dst</i>	Address of the user-space memory block
in	<i>src</i>	Source buffer address
in	<i>size</i>	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 `rtdm_rw_user_ok` and `rtdm_copy_to_user`.

## Tags

[task-unrestricted](#)

6.39.2.20 `rtdm_strncpy_from_user()`

```
int rtdm_strncpy_from_user (
    struct rtdm_fd * fd,
    char * dst,
    const char __user * src,
    size_t count )
```

Copy user-space string to specified buffer

## Parameters

in	<i>fd</i>	RTDM file descriptor as passed to the invoked device operation handler
in	<i>dst</i>	Destination buffer address
in	<i>src</i>	Address of the user-space string
in	<i>count</i>	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.

## Tags


[task-unrestricted](#)



## 6.40 SCHED\_QUOTA scheduling policy

The SCHED\_QUOTA policy enforces a limitation on the CPU consumption of threads over a globally defined period, known as the quota interval.

Collaboration diagram for SCHED\_QUOTA scheduling policy:



### 6.40.1 Detailed Description


The SCHED\_QUOTA policy enforces a limitation on the CPU consumption of threads over a globally defined period, known as the quota interval.

This is done by pooling threads with common requirements in groups, and giving each group a share of the global period (CONFIG\_XENO\_OPT\_SCHED\_QUOTA\_PERIOD).

When threads have entirely consumed the quota allotted to the group they belong to, the latter is suspended as a whole, until the next quota interval starts. At this point, a new runtime budget is given to each group, in accordance with its share.

## 6.41 Thread scheduling control

Collaboration diagram for Thread scheduling control:



### Modules

- [SCHED\\_QUOTA scheduling policy](#)

*The SCHED\_QUOTA policy enforces a limitation on the CPU consumption of threads over a globally defined period, known as the quota interval.*

### Data Structures

- struct [xnsched](#)

*Scheduling information structure.*

### Functions

- static int [xnsched\\_run](#) (void)

*The rescheduling procedure.*

- static void [xnsched\\_rotate](#) (struct [xnsched](#) \**sched*, struct [xnsched\\_class](#) \**sched\_class*, const union [xnsched\\_policy\\_param](#) \**sched\_param*)

*Rotate a scheduler runqueue.*

#### 6.41.1 Detailed Description

#### 6.41.2 Function Documentation

##### 6.41.2.1 xnsched\_rotate()

```

void xnsched_rotate (
    struct xnsched * sched,
    struct xnsched\_class * sched_class,
    const union xnsched\_policy\_param * sched_param ) [inline], [static]
  
```

Rotate a scheduler runqueue.

The specified scheduling class is requested to rotate its runqueue for the given scheduler. Rotation is performed according to the scheduling parameter specified by *sched\_param*.

#### Note

The nucleus supports round-robin scheduling for the members of the RT class.

## Parameters

<i>sched</i>	The per-CPU scheduler hosting the target scheduling class.
<i>sched_class</i>	The scheduling class which should rotate its runqueue.
<i>sched_param</i>	The scheduling parameter providing rotation information to the specified scheduling class.

## Tags

[unrestricted](#), [atomic-entry](#)

6.41.2.2 `xnsched_run()`

```
int xnsched_run (
    void ) [inline], [static]
```

The rescheduling procedure.

This is the central rescheduling routine which should be called to validate and apply changes which have previously been made to the nucleus scheduling state, such as suspending, resuming or changing the priority of threads. This call performs context switches as needed. `xnsched_run()` schedules out the current thread if:

- the current thread is about to block.
- a runnable thread from a higher priority scheduling class is waiting for the CPU.
- the current thread does not lead the runnable threads from its own scheduling class (i.e. round-robin).

The Cobalt core implements a lazy rescheduling scheme so that most of the services affecting the threads state MUST be followed by a call to the rescheduling procedure for the new scheduling state to be applied.

In other words, multiple changes on the scheduler state can be done in a row, waking threads up, blocking others, without being immediately translated into the corresponding context switches. When all changes have been applied, `xnsched_run()` should be called for considering those changes, and possibly switching context.

As a notable exception to the previous principle however, every action which ends up suspending the current thread begets an implicit call to the rescheduling procedure on behalf of the blocking service.

Typically, self-suspension or sleeping on a synchronization object automatically leads to a call to the rescheduling procedure, therefore the caller does not need to explicitly issue `xnsched_run()` after such operations.

The rescheduling procedure always leads to a null-effect if it is called on behalf of an interrupt service routine. Any outstanding scheduler lock held by the outgoing thread will be restored when the thread is scheduled back in.

Calling this procedure with no applicable context switch pending is harmless and simply leads to a null-effect.

## Returns

Non-zero is returned if a context switch actually happened, otherwise zero if the current thread was left running.


## Tags

[unrestricted](#)

## 6.42 Synchronous I/O multiplexing

This module implements the services needed for implementing the POSIX `select()` service, or any other event multiplexing services.

Collaboration diagram for Synchronous I/O multiplexing:



### Functions

- void `xnselect_init` (struct `xnselect` \*select\_block)

*Initialize a struct xnselect structure.*

- static int `xnselect_signal` (struct `xnselect` \*select\_block, unsigned int state)

*Signal a file descriptor state change.*

- void `xnselect_destroy` (struct `xnselect` \*select\_block)

*Destroy the xnselect structure associated with a file descriptor.*

- int `xnselector_init` (struct `xnselector` \*selector)

*Initialize a selector structure.*

- int `xnselect` (struct `xnselector` \*selector, fd\_set \*out\_fds[XNSELECT\_MAX\_TYPES], fd\_set \*in\_fds[XNSELECT\_MAX\_TYPES], int nfds, xnticks\_t timeout, xntmode\_t timeout\_mode)

*Check the state of a number of file descriptors, wait for a state change if no descriptor is ready.*

- void `xnselector_destroy` (struct `xnselector` \*selector)

*Destroy a selector block.*

- int `xnselect_bind` (struct `xnselect` \*select\_block, struct `xnselect_binding` \*binding, struct `xnselector` \*selector, unsigned type, unsigned index, unsigned state)

*Bind a file descriptor (represented by its xnselect structure) to a selector block.*

### 6.42.1 Detailed Description

This module implements the services needed for implementing the POSIX `select()` service, or any other event multiplexing services.

Following the implementation of the posix `select` service, this module defines three types of events:

- `XNSELECT_READ` meaning that a file descriptor is ready for reading;
- `XNSELECT_WRITE` meaning that a file descriptor is ready for writing;

- `XNSELECT_EXCEPT` meaning that a file descriptor received an exceptional event.

It works by defining two structures:

- a `struct xnselect` structure, which should be added to every file descriptor for every event type (read, write, or except);
- a `struct xnselector` structure, the selection structure, passed by the thread calling the `xnselect` service, where this service does all its housekeeping.

## 6.42.2 Function Documentation

### 6.42.2.1 `xnselect()`

```
int xnselect (
    struct xnselector * selector,
    fd_set * out_fds[XNSELECT_MAX_TYPES],
    fd_set * in_fds[XNSELECT_MAX_TYPES],
    int nfds,
    xnticks_t timeout,
    xntmode_t timeout_mode )
```

Check the state of a number of file descriptors, wait for a state change if no descriptor is ready.

#### Parameters

<code>selector</code>	structure to check for pending events
<code>out_fds</code>	The set of descriptors with pending events if a strictly positive number is returned, or the set of descriptors not yet bound if <code>-ECHRNG</code> is returned;
<code>in_fds</code>	the set of descriptors which events should be checked
<code>nfds</code>	the highest-numbered descriptor in any of the <code>in_fds</code> sets, plus 1;
<code>timeout</code>	the timeout, whose meaning depends on <code>timeout_mode</code> , note that <code>xnselect()</code> pass <code>timeout</code> and <code>timeout_mode</code> unchanged to <code>xnsynch_sleep_on</code> , so passing a relative value different from <code>XN_INFINITE</code> as a timeout with <code>timeout_mode</code> set to <code>XN_RELATIVE</code> , will cause a longer sleep than expected if the sleep is interrupted.
<code>timeout_mode</code>	the mode of <code>timeout</code> .

#### Return values

<code>-EINVAL</code>	if <code>nfds</code> is negative;
<code>-ECHRNG</code>	if some of the descriptors passed in <code>in_fds</code> have not yet been registered with <code>xnselect_bind()</code> , <code>out_fds</code> contains the set of such descriptors;
<code>-EINTR</code>	if <code>xnselect</code> was interrupted while waiting;
<code>0</code>	in case of timeout.
<code>the</code>	number of file descriptors having received an event.

Tags

[primary-only](#), [might-switch](#)

#### 6.42.2.2 xselect\_bind()

```
int xselect_bind (
    struct xselect * select_block,
    struct xselect_binding * binding,
    struct xselector * selector,
    unsigned type,
    unsigned index,
    unsigned state )
```

Bind a file descriptor (represented by its *xselect* structure) to a selector block.

Parameters

<i>select_block</i>	pointer to the <i>struct xselect</i> to be bound;
<i>binding</i>	pointer to a newly allocated (using <i>xmalloc</i> ) <i>struct xselect_binding</i> ;
<i>selector</i>	pointer to the selector structure;
<i>type</i>	type of events ( <i>XNSELECT_READ</i> , <i>XNSELECT_WRITE</i> , or <i>XNSELECT_EXCEPT</i> );
<i>index</i>	index of the file descriptor (represented by <i>select_block</i> ) in the bit fields used by the <i>selector</i> structure;
<i>state</i>	current state of the file descriptor.

*select\_block* must have been initialized with [xselect\\_init\(\)](#), the *xselector* structure must have been initialized with [xselector\\_init\(\)](#), *binding* may be uninitialized.

This service must be called with *nklock* locked, *irqs* off. For this reason, the *binding* parameter must have been allocated by the caller outside the locking section.

Return values

<i>-EINVAL</i>	if <i>type</i> or <i>index</i> is invalid;
<i>0</i>	otherwise.

Tags

[task-unrestricted](#), [might-switch](#), [atomic-entry](#)

#### 6.42.2.3 xselect\_destroy()

```
void xselect_destroy (
    struct xselect * select_block )
```

Destroy the *xnselect* structure associated with a file descriptor.

Any binding with a *xnselector* block is destroyed.

Parameters

<i>select_block</i>	pointer to the <i>xnselect</i> structure associated with a file descriptor
---------------------	--

Tags

[task-unrestricted](#), [might-switch](#)

#### 6.42.2.4 xnselect\_init()

```
void xnselect_init (
    struct xnselect * select_block )
```

Initialize a *struct xnselect* structure.

This service must be called to initialize a *struct xnselect* structure before it is bound to a selector by the means of [xnselect\\_bind\(\)](#).

Parameters

<i>select_block</i>	pointer to the <i>xnselect</i> structure to be initialized
---------------------	--

Tags

[task-unrestricted](#)

#### 6.42.2.5 xnselect\_signal()

```
static int xnselect_signal (
    struct xnselect * select_block,
    unsigned int state ) [inline], [static]
```

Signal a file descriptor state change.

Parameters

<i>select_block</i>	pointer to an <i>xnselect</i> structure representing the file descriptor whose state changed;
<i>state</i>	new value of the state.

Return values

1	if rescheduling is needed;
0	otherwise.

#### 6.42.2.6 xselector\_destroy()

```
void xselector_destroy (  
    struct xselector * selector )
```

Destroy a selector block.

All bindings with file descriptor are destroyed.

Parameters

<i>selector</i>	the selector block to be destroyed
-----------------	------------------------------------

Tags

[task-unrestricted](#)

#### 6.42.2.7 xselector\_init()

```
int xselector_init (  
    struct xselector * selector )
```

Initialize a selector structure.

Parameters

<i>selector</i>	The selector structure to be initialized.
-----------------	---

Return values

0	
---	--


Tags

[task-unrestricted](#)



## 6.43 Thread synchronization services

Collaboration diagram for Thread synchronization services:



### Functions

- void [xnsynch\\_init](#) (struct xnsynch \*synch, int flags, [atomic\\_t](#) \*fastlock)

*Initialize a synchronization object.*

- void [xnsynch\\_init\\_protect](#) (struct xnsynch \*synch, int flags, [atomic\\_t](#) \*fastlock, u32 \*ceiling\_ref)

*Initialize a synchronization object enforcing PP.*

- int [xnsynch\\_destroy](#) (struct xnsynch \*synch)

*Destroy a synchronization object.*

- int \_\_must\_check [xnsynch\\_sleep\\_on](#) (struct xnsynch \*synch, xnticks\_t timeout, xntmode\_t timeout\_mode)

*Sleep on an ownerless synchronization object.*

- struct xnthread \* [xnsynch\\_wakeup\\_one\\_sleeper](#) (struct xnsynch \*synch)

*Unblock the heading thread from wait.*

- void [xnsynch\\_wakeup\\_this\\_sleeper](#) (struct xnsynch \*synch, struct xnthread \*sleeper)

*Unblock a particular thread from wait.*

- int \_\_must\_check [xnsynch\\_acquire](#) (struct xnsynch \*synch, xnticks\_t timeout, xntmode\_t timeout\_mode)

*Acquire the ownership of a synchronization object.*

- int \_\_must\_check [xnsynch\\_try\\_acquire](#) (struct xnsynch \*synch)

*Try acquiring the ownership of a synchronization object.*

- bool [xnsynch\\_release](#) (struct xnsynch \*synch, struct xnthread \*thread)

*Release a resource and pass it to the next waiting thread.*

- struct xnthread \* [xnsynch\\_peek\\_pendq](#) (struct xnsynch \*synch)

*Access the thread leading a synch object wait queue.*

- int [xnsynch\\_flush](#) (struct xnsynch \*synch, int reason)

*Unblock all waiters pending on a resource.*

### 6.43.1 Detailed Description

### 6.43.2 Function Documentation

#### 6.43.2.1 xnsynch\_acquire()

```
int xnsynch_acquire (
    struct xnsynch * synch,
    xnticks_t timeout,
    xntmode_t timeout_mode )
```

Acquire the ownership of a synchronization object.

This service should be called by upper interfaces wanting the current thread to acquire the ownership of the given resource. If the resource is already assigned to another thread, the caller is suspended.

This service must be used only with synchronization objects that track ownership (XNSYNCH\_OWNER set).

#### Parameters

<i>synch</i>	The descriptor address of the synchronization object to acquire.
<i>timeout</i>	The timeout which may be used to limit the time the thread pends on the resource. This value is a wait time given as a count of nanoseconds. It can either be relative, absolute monotonic, or absolute adjustable depending on <i>timeout_mode</i> . Passing XN_INFINITE <b>and</b> setting <i>mode</i> to XN_RELATIVE specifies an unbounded wait. All other values are used to initialize a watchdog timer.
<i>timeout_mode</i>	The mode of the <i>timeout</i> parameter. It can either be set to XN_RELATIVE, XN_ABSOLUTE, or XN_REALTIME (see also <a href="#">xntimer_start()</a> ).

#### Returns

A bitmask which may include zero or one information bit among XNRMID, XNTIMEO and XNB←REAK, which should be tested by the caller, for detecting respectively: object deletion, timeout or signal/unblock conditions which might have happened while waiting.

#### Tags

[primary-only](#), [might-switch](#)

#### Note

Unlike [xnsynch\\_try\\_acquire\(\)](#), this call does NOT check for invalid recursive locking request, which means that such request will always cause a deadlock for the caller.

6.43.2.2 `xnsynch_destroy()`

```
int xnsynch_destroy (
    struct xnsynch * synch )
```

Destroy a synchronization object.

Destroys the synchronization object *synch*, unblocking all waiters with the XNRMID status.

Returns

XNSYNCH\_RESCHEDED is returned if at least one thread is unblocked, which means the caller should invoke [xnsched\\_run\(\)](#) for applying the new scheduling state. Otherwise, XNSYNCH\_DONE is returned.

Side effects

Same as [xnsynch\\_flush\(\)](#).

Tags

[task-unrestricted](#)

6.43.2.3 `xnsynch_flush()`

```
int xnsynch_flush (
    struct xnsynch * synch,
    int reason )
```

Unblock all waiters pending on a resource.

This service atomically releases all threads which currently sleep on a given resource. This service should be called by upper interfaces under circumstances requiring that the pending queue of a given resource is cleared, such as before the resource is deleted.

Parameters

<i>synch</i>	The descriptor address of the synchronization object to be flushed.
<i>reason</i>	Some flags to set in the information mask of every unblocked thread. Zero is an acceptable value. The following bits are pre-defined by Cobalt:

- XNRMID should be set to indicate that the synchronization object is about to be destroyed (see [xnthread\\_resume\(\)](#)).
- XNBREAK should be set to indicate that the wait has been forcibly interrupted (see [xnthread\\_unblock\(\)](#)).

## Returns

XNSYNCH\_RESCHEDED is returned if at least one thread is unblocked, which means the caller should invoke [xnsched\\_run\(\)](#) for applying the new scheduling state. Otherwise, XNSYNCH\_DONE is returned.

## Side effects

- The effective priority of the current resource owner might be lowered to its base priority value as a consequence of the priority inheritance boost being cleared.

## Tags

[unrestricted](#)

## 6.43.2.4 xnsynch\_init()

```
void xnsynch_init (
    struct xnsynch * synch,
    int flags,
    atomic_t * fastlock )
```

Initialize a synchronization object.

Initializes a synchronization object. Xenomai threads can wait on and signal such objects for serializing access to resources. This object has built-in support for priority inheritance.

## Parameters

<i>synch</i>	The address of a synchronization object descriptor Cobalt will use to store the object-specific data. This descriptor must always be valid while the object is active therefore it must be allocated in permanent memory.
<i>flags</i>	A set of creation flags affecting the operation. The valid flags are:

- XNSYNCH\_PRIO causes the threads waiting for the resource to pend in priority order. Otherwise, FIFO ordering is used (XNSYNCH\_FIFO).
- XNSYNCH\_OWNER indicates that the synchronization object shall track the resource ownership, allowing a single owner at most at any point in time. Note that setting this flag implies the use of [xnsynch\\_acquire\(\)](#) and [xnsynch\\_release\(\)](#) instead of [xnsynch\\_sleep\\_on\(\)](#) and [xnsynch\\_wakeup\\_\\*](#)().
- XNSYNCH\_PI enables priority inheritance when a priority inversion is detected among threads using this object. XNSYNCH\_PI implies XNSYNCH\_OWNER and XNSYNCH\_PRIO.
- XNSYNCH\_PP enables priority protect to prevent priority inversion. XNSYNCH\_PP implies XNSYNCH\_OWNER and XNSYNCH\_PRIO.

- XNSYNCH\_DREORD (Disable REORDERing) tells Cobalt not to reorder the wait list upon priority change of a waiter. Reordering is the default. Only applies when XNSYNCH\_PRIO is present.

## Parameters

<i>fastlock</i>	Address of the fast lock word to be associated with a synchronization object with ownership tracking. Therefore, a valid fast-lock address is required if XNSYNCH_OWNER is set in <i>flags</i> .
-----------------	--

## Tags

[task-unrestricted](#)

Referenced by `xnsynch_init_protect()`.

6.43.2.5 `xnsynch_init_protect()`

```
void xnsynch_init_protect (
    struct xnsynch * synch,
    int flags,
    atomic_t * fastlock,
    u32 * ceiling_ref )
```

Initialize a synchronization object enforcing PP.

This call is a variant of [xnsynch\\_init\(\)](#) for initializing synchronization objects enabling the priority protect protocol.

## Parameters

<i>synch</i>	The address of a synchronization object descriptor Cobalt will use to store the object-specific data. See <a href="#">xnsynch_init()</a> .
<i>flags</i>	A set of creation flags affecting the operation. See <a href="#">xnsynch_init()</a> . XNSYNCH_PI is mutually exclusive with XNSYNCH_PP, and won't be considered.
<i>fastlock</i>	Address of the fast lock word to be associated with a synchronization object with ownership tracking. See <a href="#">xnsynch_init()</a> .
<i>ceiling_ref</i>	The address of the variable holding the current priority ceiling value for this object.

## Tags

[task-unrestricted](#)

References `xnsynch_init()`.

6.43.2.6 `xnsynch_peek_pendq()`

```
struct xnthread * xnsynch_peek_pendq (
    struct xnsynch * synch )
```

Access the thread leading a synch object wait queue.

This services returns the descriptor address of to the thread leading a synchronization object wait queue.

## Parameters

<i>synch</i>	The descriptor address of the target synchronization object.
--------------	--

## Returns

The descriptor address of the unblocked thread.

## Tags

[unrestricted](#)

6.43.2.7 `xnsynch_release()`

```
bool xnsynch_release (
    struct xnsynch * synch,
    struct xnthread * curr )
```

Release a resource and pass it to the next waiting thread.

This service releases the ownership of the given synchronization object. The thread which is currently leading the object's pending list, if any, is unblocked from its pending state. However, no reschedule is performed.

This service must be used only with synchronization objects that track ownership (`XNSYNCH_OWNER` set).

## Parameters

<i>synch</i>	The descriptor address of the synchronization object whose ownership is changed.
<i>curr</i>	The descriptor address of the current thread, which must own the object at the time of calling.

## Returns

True if a reschedule is required.

## Side effects

- The effective priority of the previous resource owner might be lowered to its base priority value as a consequence of the priority boost being cleared.
- The synchronization object ownership is transferred to the unblocked thread.

Tags

[primary-only](#), [might-switch](#)

#### 6.43.2.8 xnsynch\_sleep\_on()

```
int xnsynch_sleep_on (
    struct xnsynch * synch,
    xnticks_t timeout,
    xntmode_t timeout_mode )
```

Sleep on an ownerless synchronization object.

Makes the calling thread sleep on the specified synchronization object, waiting for it to be signaled.

This service should be called by upper interfaces wanting the current thread to pend on the given resource. It must not be used with synchronization objects that are supposed to track ownership (XNSY↔NCH\_OWNER).

Parameters

<i>synch</i>	The descriptor address of the synchronization object to sleep on.
<i>timeout</i>	The timeout which may be used to limit the time the thread pends on the resource. This value is a wait time given as a count of nanoseconds. It can either be relative, absolute monotonic, or absolute adjustable depending on <i>timeout_mode</i> . Passing XN_INFINITE <b>and</b> setting <i>mode</i> to XN_RELATIVE specifies an unbounded wait. All other values are used to initialize a watchdog timer.
<i>timeout_mode</i>	The mode of the <i>timeout</i> parameter. It can either be set to XN_RELATIVE, XN_ABSOLUTE, or XN_REALTIME (see also <a href="#">xntimer_start()</a> ).

Returns

A bitmask which may include zero or one information bit among XNRMID, XNTIMEO and XNB↔REAK, which should be tested by the caller, for detecting respectively: object deletion, timeout or signal/unblock conditions which might have happened while waiting.

Tags

[primary-only](#), [might-switch](#)

## 6.43.2.9 xnsynch\_try\_acquire()

```
int xnsynch_try_acquire (
    struct xnsynch * synch )
```

Try acquiring the ownership of a synchronization object.

This service should be called by upper interfaces wanting the current thread to acquire the ownership of the given resource. If the resource is already assigned to another thread, the call returns with an error code.

This service must be used only with synchronization objects that track ownership (XNSYNCH\_OWNER set).

## Parameters

<i>synch</i>	The descriptor address of the synchronization object to acquire.
--------------	--

## Returns

Zero is returned if *synch* has been successfully acquired. Otherwise:

- -EDEADLK is returned if *synch* is currently held by the calling thread.
- -EBUSY is returned if *synch* is currently held by another thread.

## Tags

[primary-only](#)

## 6.43.2.10 xnsynch\_wakeup\_one\_sleeper()

```
struct xnthread * xnsynch_wakeup_one_sleeper (
    struct xnsynch * synch )
```

Unblock the heading thread from wait.

This service wakes up the thread which is currently leading the synchronization object's pending list. The sleeping thread is unblocked from its pending state, but no reschedule is performed.

This service should be called by upper interfaces wanting to signal the given resource so that a single waiter is resumed. It must not be used with synchronization objects that are supposed to track ownership (XNSYNCH\_OWNER not set).

## Parameters

<i>synch</i>	The descriptor address of the synchronization object whose ownership is changed.
--------------	--



## Returns

The descriptor address of the unblocked thread.

## Tags

[unrestricted](#)

## 6.43.2.11 xnsynch\_wakeup\_this\_sleeper()

```
void xnsynch_wakeup_this_sleeper (  
    struct xnsynch * synch,  
    struct xnthread * sleeper )
```

Unblock a particular thread from wait.

This service wakes up a specific thread which is currently pending on the given synchronization object. The sleeping thread is unblocked from its pending state, but no reschedule is performed.

This service should be called by upper interfaces wanting to signal the given resource so that a specific waiter is resumed. It must not be used with synchronization objects that are supposed to track ownership (XNSYNCH\_OWNER not set).

## Parameters


<i>synch</i>	The descriptor address of the synchronization object whose ownership is changed.
<i>sleeper</i>	The thread to unblock which MUST be currently linked to the synchronization object's pending queue (i.e. <i>synch-&gt;pendq</i> ).

## Tags

[unrestricted](#)

## 6.44 Thread services

Collaboration diagram for Thread services:



### Modules

- [Thread runtime statistics](#)
- [Thread state flags](#)  
*Bits reporting permanent or transient states of threads.*
- [Thread information flags](#)  
*Bits reporting events notified to threads.*

### Functions

- static struct xntthread \* [xntthread\\_current](#) (void)  
*Retrieve the current Cobalt core TCB.*
- static struct xntthread \* [xntthread\\_from\\_task](#) (struct task\_struct \*p)  
*Retrieve the Cobalt core TCB attached to a Linux task.*
- static void [xntthread\\_test\\_cancel](#) (void)  
*Introduce a thread cancellation point.*
- int [xntthread\\_init](#) (struct xntthread \*thread, const struct xntthread\_init\_attr \*attr, struct xnsched\_class \*sched\_class, const union xnsched\_policy\_param \*sched\_param)  
*Initialize a new thread.*
- int [xntthread\\_start](#) (struct xntthread \*thread, const struct xntthread\_start\_attr \*attr)  
*Start a newly created thread.*
- int [xntthread\\_set\\_mode](#) (int clrmask, int setmask)  
*Change control mode of the current thread.*
- void [xntthread\\_suspend](#) (struct xntthread \*thread, int mask, xnticks\_t timeout, xntmode\_t timeout\_mode, struct xnsynch \*wchan)  
*Suspend a thread.*
- void [xntthread\\_resume](#) (struct xntthread \*thread, int mask)

*Resume a thread.*

- int `xnthread_unblock` (struct `xnthread *thread`)

*Unblock a thread.*

- int `xnthread_set_periodic` (struct `xnthread *thread`, `xnticks_t` idate, `xntmode_t` timeout\_mode, `xnticks_t` period)

*Make a thread periodic.*

- int `xnthread_wait_period` (unsigned long `*overruns_r`)

*Wait for the next periodic release point.*

- int `xnthread_set_slice` (struct `xnthread *thread`, `xnticks_t` quantum)

*Set thread time-slicing information.*

- void `xnthread_cancel` (struct `xnthread *thread`)

*Cancel a thread.*

- int `xnthread_join` (struct `xnthread *thread`, bool uninterruptible)

*Join with a terminated thread.*

- int `xnthread_harden` (void)

*Migrate a Linux task to the Xenomai domain.*

- void `xnthread_relax` (int notify, int reason)

*Switch a shadow thread back to the Linux domain.*

- int `xnthread_map` (struct `xnthread *thread`, struct completion `*done`)

*Create a shadow thread context over a kernel task.*

- int `xnthread_set_schedparam` (struct `xnthread *thread`, struct `xnsched_class *sched_class`, const union `xnsched_policy_param *sched_param`)

*Change the base scheduling parameters of a thread.*

### 6.44.1 Detailed Description

### 6.44.2 Function Documentation

#### 6.44.2.1 `xnthread_cancel()`

```
void xnthread_cancel (
    struct xnthread * thread )
```

Cancel a thread.

Request cancellation of a thread. This service forces *thread* to exit from any blocking call, then to switch to secondary mode. *thread* will terminate as soon as it reaches a cancellation point. Cancellation points are defined for the following situations:

- *thread* self-cancels by a call to `xnthread_cancel()`.
- *thread* invokes a Linux syscall (user-space shadow only).
- *thread* receives a Linux signal (user-space shadow only).
- *thread* unblocks from a Xenomai syscall (user-space shadow only).
- *thread* attempts to block on a Xenomai syscall (user-space shadow only).
- *thread* explicitly calls `xnthread_test_cancel()`.

#### Parameters

<i>thread</i>	The descriptor address of the thread to terminate.
---------------	--

#### Tags

[task-unrestricted](#), [might-switch](#)

#### Note

In addition to the common actions taken upon cancellation, a thread which belongs to the SCHE↔D\_WEAK class is sent a regular SIGTERM signal.

#### 6.44.2.2 xnthread\_current()

```
struct xnthread * xnthread_current (
    void ) [inline], [static]
```

Retrieve the current Cobalt core TCB.

Returns the address of the current Cobalt core thread descriptor, or NULL if running over a regular Linux task. This call is not affected by the current runtime mode of the core thread.

#### Note

The returned value may differ from `xnsched_current_thread()` called from the same context, since the latter returns the root thread descriptor for the current CPU if the caller is running in secondary mode.

#### Tags

[unrestricted](#)

6.44.2.3 `xnthread_from_task()`

```
struct xnthread * xnthread_from_task (
    struct task_struct * p ) [inline], [static]
```

Retrieve the Cobalt core TCB attached to a Linux task.

Returns the address of the Cobalt core thread descriptor attached to the Linux task *p*, or NULL if *p* is a regular Linux task. This call is not affected by the current runtime mode of the core thread.

Tags

[unrestricted](#)

6.44.2.4 `xnthread_harden()`

```
int xnthread_harden (
    void )
```

Migrate a Linux task to the Xenomai domain.

This service causes the transition of "current" from the Linux domain to Xenomai. The shadow will resume in the Xenomai domain as returning from `schedule()`.

Tags

[secondary-only](#), [might-switch](#)

6.44.2.5 `xnthread_init()`

```
int xnthread_init (
    struct xnthread * thread,
    const struct xnthread_init_attr * attr,
    struct xnsched_class * sched_class,
    const union xnsched_policy_param * sched_param )
```

Initialize a new thread.

Initializes a new thread. The thread is left dormant until it is actually started by [xnthread\\_start\(\)](#).

## Parameters

<i>thread</i>	The address of a thread descriptor Cobalt will use to store the thread-specific data. This descriptor must always be valid while the thread is active therefore it must be allocated in permanent memory.
---------------	---

## Warning

Some architectures may require the descriptor to be properly aligned in memory; this is an additional reason for descriptors not to be laid in the program stack where alignment constraints might not always be satisfied.

## Parameters

<i>attr</i>	A pointer to an attribute block describing the initial properties of the new thread. Members of this structure are defined as follows:
-------------	--

- *name*: An ASCII string standing for the symbolic name of the thread. This name is copied to a safe place into the thread descriptor. This name might be used in various situations by Cobalt for issuing human-readable diagnostic messages, so it is usually a good idea to provide a sensible value here. NULL is fine though and means "anonymous".
- *flags*: A set of creation flags affecting the operation. The following flags can be part of this bitmask:
  - XNSUSP creates the thread in a suspended state. In such a case, the thread shall be explicitly resumed using the `xnthread_resume()` service for its execution to actually begin, additionally to issuing `xnthread_start()` for it. This flag can also be specified when invoking `xnthread_start()` as a starting mode.
- XNUSER shall be set if *thread* will be mapped over an existing user-space task. Otherwise, a new kernel host task is created, then paired with the new Xenomai thread.
- XNFPU (enable FPU) tells Cobalt that the new thread may use the floating-point unit. XNFPU is implicitly assumed for user-space threads even if not set in *flags*.
- *affinity*: The processor affinity of this thread. Passing CPU\_MASK\_ALL means "any cpu" from the allowed core affinity mask (`cobalt_cpu_affinity`). Passing an empty set is invalid.

## Parameters

<i>sched_class</i>	The initial scheduling class the new thread should be assigned to.
<i>sched_param</i>	The initial scheduling parameters to set for the new thread; <i>sched_param</i> must be valid within the context of <i>sched_class</i> .

## Returns

0 is returned on success. Otherwise, the following error code indicates the cause of the failure:

- -EINVAL is returned if *attr->flags* has invalid bits set, or *attr->affinity* is invalid (e.g. empty).

## Tags

[secondary-only](#)

References XNFPU, XNSUSP, and XNUSER.

#### 6.44.2.6 `xnthread_join()`

```
int xnthread_join (
    struct xnthread * thread,
    bool uninterruptible )
```

Join with a terminated thread.

This service waits for *thread* to terminate after a call to `xnthread_cancel()`. If that thread has already terminated or is dormant at the time of the call, then `xnthread_join()` returns immediately.

`xnthread_join()` adapts to the calling context (primary or secondary), switching to secondary mode if needed for the duration of the wait. Upon return, the original runtime mode is restored, unless a Linux signal is pending.

Parameters

<i>thread</i>	The descriptor address of the thread to join with.
<i>uninterruptible</i>	Boolean telling whether the service should wait for completion uninterruptible.

Returns

0 is returned on success. Otherwise, the following error codes indicate the cause of the failure:

- -EDEADLK is returned if the current thread attempts to join itself.
- -EINTR is returned if the current thread was unblocked while waiting for *thread* to terminate.
- -EBUSY indicates that another thread is already waiting for *thread* to terminate.

Tags

[task-unrestricted](#), [might-switch](#)

#### 6.44.2.7 `xnthread_map()`

```
int xnthread_map (
    struct xnthread * thread,
    struct completion * done )
```

Create a shadow thread context over a kernel task.

This call maps a Cobalt core thread to the "current" Linux task running in kernel space. The priority and scheduling class of the underlying Linux task are not affected; it is assumed that the caller did set them appropriately before issuing the shadow mapping request.

This call immediately moves the calling kernel thread to the Xenomai domain.

## Parameters

<i>thread</i>	The descriptor address of the new shadow thread to be mapped to "current". This descriptor must have been previously initialized by a call to <a href="#">xnthread_init()</a> .
<i>done</i>	A completion object to be signaled when <i>thread</i> is fully mapped over the current Linux context, waiting for <a href="#">xnthread_start()</a> .

## Returns

0 is returned on success. Otherwise:

- -ERESTARTSYS is returned if the current Linux task has received a signal, thus preventing the final migration to the Xenomai domain (i.e. in order to process the signal in the Linux domain). This error should not be considered as fatal.
- -EPERM is returned if the shadow thread has been killed before the current task had a chance to return to the caller. In such a case, the real-time mapping operation has failed globally, and no Xenomai resource remains attached to it.
- -EINVAL is returned if the thread control block bears the XNUSER bit.
- -EBUSY is returned if either the current Linux task or the associated shadow thread is already involved in a shadow mapping.

## Tags

[secondary-only](#), [might-switch](#)

## 6.44.2.8 xnthread\_relax()

```
void xnthread_relax (
    int notify,
    int reason )
```

Switch a shadow thread back to the Linux domain.

This service yields the control of the running shadow back to Linux. This is obtained by suspending the shadow and scheduling a wake up call for the mated user task inside the Linux domain. The Linux task will resume on return from [xnthread\\_suspend\(\)](#) on behalf of the root thread.

## Parameters

<i>notify</i>	A boolean flag indicating whether threads monitored from secondary mode switches should be sent a SIGDEBUG signal. For instance, some internal operations like task exit should not trigger such signal.
<i>reason</i>	The reason to report along with the SIGDEBUG signal.



Tags

[primary-only](#), [might-switch](#)

Note

"current" is valid here since the shadow runs with the properties of the Linux task.

#### 6.44.2.9 `xnthread_resume()`

```
void xnthread_resume (
    struct xnthread * thread,
    int mask )
```

Resume a thread.

Resumes the execution of a thread previously suspended by one or more calls to [xnthread\\_suspend\(\)](#). This call removes a suspensive condition affecting the target thread. When all suspensive conditions are gone, the thread is left in a READY state at which point it becomes eligible anew for scheduling.

Parameters

<i>thread</i>	The descriptor address of the resumed thread.
<i>mask</i>	The suspension mask specifying the suspensive condition to remove from the thread's wait mask. Possible values usable by the caller are:

- XNSUSP. This flag removes the explicit suspension condition. This condition might be additive to the XNPEND condition.
- XNDELAY. This flag removes the counted delay wait condition.
- XNPEND. This flag removes the resource wait condition. If a watchdog is armed, it is automatically disarmed by this call. Unlike the two previous conditions, only the current thread can set this condition for itself, i.e. no thread can force another one to pend on a resource.

When the thread is eventually resumed by one or more calls to [xnthread\\_resume\(\)](#), the caller of [xnthread\\_suspend\(\)](#) in the awakened thread that suspended itself should check for the following bits in its own information mask to determine what caused its wake up:

- XNRMID means that the caller must assume that the pended synchronization object has been destroyed (see [xnsynch\\_flush\(\)](#)).
- XNTIMEO means that the delay elapsed, or the watchdog went off before the corresponding synchronization object was signaled.
- XNBREAK means that the wait has been forcibly broken by a call to [xnthread\\_unblock\(\)](#).

Tags

[unrestricted](#), [might-switch](#)

6.44.2.10 `xnthread_set_mode()`

```
int xnthread_set_mode (
    int clrmask,
    int setmask )
```

Change control mode of the current thread.

Change the control mode of the current thread. The control mode affects several behaviours of the Cobalt core regarding this thread.

## Parameters

<i>clrmask</i>	Clears the corresponding bits from the control mode before <i>setmask</i> is applied. The scheduler lock held by the current thread can be forcibly released by passing the XNLOCK bit in this mask. In this case, the lock nesting count is also reset to zero.
<i>setmask</i>	The new thread mode. The following flags may be set in this bitmask:

- XNLOCK makes the current thread non-preemptible by other threads. Unless XNTRAPLB is also set for the thread, the latter may still block, dropping the lock temporarily, in which case, the lock will be reacquired automatically when the thread resumes execution.
- XNWARN enables debugging notifications for the current thread. A SIGDEBUG (Linux-originated) signal is sent when the following atypical or abnormal behavior is detected:
  - the current thread switches to secondary mode. Such notification comes in handy for detecting spurious relaxes.
  - CONFIG\_XENO\_OPT\_DEBUG\_MUTEX\_RELAXED is enabled in the kernel configuration, and the current thread is sleeping on a Cobalt mutex currently owned by a thread running in secondary mode, which reveals a priority inversion.
  - the current thread is about to sleep while holding a Cobalt mutex, and CONFIG\_XENO\_OPT\_DEBUG\_MUTEX\_SLEEP is enabled in the kernel configuration. Blocking for acquiring a mutex does not trigger such a signal though.
  - the current thread has both XNTRAPLB and XNLOCK set, and attempts to block on a Cobalt service, which would cause a lock break.
- XNTRAPLB disallows breaking the scheduler lock. In the default case, a thread which holds the scheduler lock is allowed to drop it temporarily for sleeping. If this mode bit is set, such thread would return immediately with XNBREAK set from `xnthread_suspend()`. If XNWARN is set for the current thread, SIGDEBUG is sent in addition to raising the break condition.

## Tags

[primary-only](#), [might-switch](#)

## Note

Setting *clrmask* and *setmask* to zero leads to a nop, in which case `xnthread_set_mode()` returns the current mode.

6.44.2.11 `xnthread_set_periodic()`

```
int xnthread_set_periodic (
    struct xnthread * thread,
    xnticks_t idate,
    xntmode_t timeout_mode,
    xnticks_t period )
```

Make a thread periodic.

Make a thread periodic by programming its first release point and its period in the processor time line. Subsequent calls to `xnthread_wait_period()` will delay the thread until the next periodic release point in the processor timeline is reached.

Parameters

<i>thread</i>	The core thread to make periodic. If NULL, the current thread is assumed.
<i>idate</i>	The initial (absolute) date of the first release point, expressed in nanoseconds. The affected thread will be delayed by the first call to <code>xnthread_wait_period()</code> until this point is reached. If <i>idate</i> is equal to XN_INFINITE, the first release point is set to <i>period</i> nanoseconds after the current date. In the latter case, <i>timeout_mode</i> is not considered and can have any valid value.
<i>timeout_mode</i>	The mode of the <i>idate</i> parameter. It can either be set to XN_ABSOLUTE or XN_REALTIME with <i>idate</i> different from XN_INFINITE (see also <code>xntimer_start()</code> ).
<i>period</i>	The period of the thread, expressed in nanoseconds. As a side-effect, passing XN_INFINITE attempts to stop the thread's periodic timer; in the latter case, the routine always exits successfully, regardless of the previous state of this timer.

Returns

0 is returned upon success. Otherwise:

- -ETIMEDOUT is returned if *idate* is different from XN\_INFINITE and represents a date in the past.
- -EINVAL is returned if *period* is different from XN\_INFINITE but shorter than the scheduling latency value for the target system, as available from `/proc/xenomai/latency`. -EINVAL is also returned if *timeout\_mode* is not compatible with *idate*, such as XN\_RELATIVE with *idate* different from XN\_INFINITE.
- -EPERM is returned if *thread* is NULL, but the caller is not a Xenomai thread.

Tags

[task-unrestricted](#)

6.44.2.12 `xnthread_set_schedparam()`

```
int xnthread_set_schedparam (
    struct xnthread * thread,
    struct xnsched_class * sched_class,
    const union xnsched_policy_param * sched_param )
```

Change the base scheduling parameters of a thread.

Changes the base scheduling policy and parameters of a thread. If the thread is currently blocked, waiting in priority-pending mode (XNSYNCH\_PRIO) for a synchronization object to be signaled, Cobalt will attempt to reorder the object's wait queue so that it reflects the new sleeper's priority, unless the XNSYNCH\_DREORD flag has been set for the pended object.

## Parameters

<i>thread</i>	The descriptor address of the affected thread. See note.
<i>sched_class</i>	The new scheduling class the thread should be assigned to.
<i>sched_param</i>	The scheduling parameters to set for the thread; <i>sched_param</i> must be valid within the context of <i>sched_class</i> .

It is absolutely required to use this service to change a thread priority, in order to have all the needed housekeeping chores correctly performed. i.e. Do *not* call `xnsched_set_policy()` directly or worse, change the `thread.cprio` field by hand in any case.

## Returns

0 is returned on success. Otherwise, a negative error code indicates the cause of a failure that happened in the scheduling class implementation for *sched\_class*. Invalid parameters passed into *sched\_param* are common causes of error.

## Side effects

- This service does not call the rescheduling procedure but may affect the state of the run queue for the previous and new scheduling classes.
- Assigning the same scheduling class and parameters to a running or ready thread moves it to the end of the run queue, thus causing a manual round-robin, except if a priority boost is undergoing.

## Tags

[task-unregistred](#)

## Note

The changes only apply to the Xenomai scheduling parameters for *thread*. There is no propagation/translation of such changes to the Linux scheduler for the task mated to the Xenomai target thread.

6.44.2.13 `xnthread_set_slice()`

```
int xnthread_set_slice (
    struct xnthread * thread,
    xnticks_t quantum )
```

Set thread time-slicing information.

Update the time-slicing information for a given thread. This service enables or disables round-robin scheduling for the thread, depending on the value of *quantum*. By default, times-slicing is disabled for a new thread initialized by a call to `xnthread_init()`.

Parameters

<i>thread</i>	The descriptor address of the affected thread.
<i>quantum</i>	The time quantum assigned to the thread expressed in nanoseconds. If <i>quantum</i> is different from <code>XN_INFINITE</code> , the time-slice for the thread is set to that value and its current time credit is refilled (i.e. the thread is given a full time-slice to run next). Otherwise, if <i>quantum</i> equals <code>XN_INFINITE</code> , time-slicing is stopped for that thread.

Returns

0 is returned upon success. Otherwise, `-EINVAL` is returned if *quantum* is not `XN_INFINITE` and:

- the base scheduling class of the target thread does not support time-slicing,
- *quantum* is smaller than the master clock gravity for a user thread, which denotes a spurious value.

Tags

[task-unrestricted](#)

6.44.2.14 `xnthread_start()`

```
int xnthread_start (
    struct xnthread * thread,
    const struct xnthread_start_attr * attr )
```

Start a newly created thread.

Starts a (newly) created thread, scheduling it for the first time. This call releases the target thread from the `XNDORMANT` state. This service also sets the initial mode for the new thread.

Parameters

<i>thread</i>	The descriptor address of the started thread which must have been previously initialized by a call to <code>xnthread_init()</code> .
<i>attr</i>	A pointer to an attribute block describing the execution properties of the new thread. Members of this structure are defined as follows:

- mode: The initial thread mode. The following flags can be part of this bitmask:
  - XNLOCK causes the thread to lock the scheduler when it starts. The target thread will have to call the `xnsched_unlock()` service to unlock the scheduler. A non-preemptible thread may still block, in which case, the lock is reasserted when the thread is scheduled back in.
  - XNSUSP makes the thread start in a suspended state. In such a case, the thread will have to be explicitly resumed using the `xnthread_resume()` service for its execution to actually begin.
- entry: The address of the thread's body routine. In other words, it is the thread entry point.
- cookie: A user-defined opaque cookie Cobalt will pass to the emerging thread as the sole argument of its entry point.

Return values

<i>0</i>	if <i>thread</i> could be started ;
<i>-EBUSY</i>	if <i>thread</i> was not dormant or stopped ;

Tags

[task-unrestricted](#), [might-switch](#)

#### 6.44.2.15 xnthread\_suspend()

```
void xnthread_suspend (
    struct xnthread * thread,
    int mask,
    xnticks_t timeout,
    xntmode_t timeout_mode,
    struct xnsynch * wchan )
```

Suspend a thread.

Suspends the execution of a thread according to a given suspensive condition. This thread will not be eligible for scheduling until it all the pending suspensive conditions set by this service are removed by one or more calls to `xnthread_resume()`.

Parameters

<i>thread</i>	The descriptor address of the suspended thread.
<i>mask</i>	The suspension mask specifying the suspensive condition to add to the thread's wait mask. Possible values usable by the caller are:

- XNSUSP. This flag forcibly suspends a thread, regardless of any resource to wait for. A reverse call to `xnthread_resume()` specifying the XNSUSP bit must be issued to remove this condition, which is cumulative with other suspension bits. `wchan` should be NULL when using this suspending mode.
- XNDELAY. This flag denotes a counted delay wait (in ticks) which duration is defined by the value of the timeout parameter.

- **XNPEND**. This flag denotes a wait for a synchronization object to be signaled. The `wchan` argument must point to this object. A timeout value can be passed to bound the wait. This suspending mode should not be used directly by the client interface, but rather through the `xnsynch_sleep_on()` call.

## Parameters

<i>timeout</i>	The timeout which may be used to limit the time the thread pends on a resource. This value is a wait time given in nanoseconds. It can either be relative, absolute monotonic, or absolute adjustable depending on <i>timeout_mode</i> .
----------------	--

Passing `XN_INFINITE` **and** setting *timeout\_mode* to `XN_RELATIVE` specifies an unbounded wait. All other values are used to initialize a watchdog timer. If the current operation mode of the system timer is oneshot and *timeout* elapses before `xnthread_suspend()` has completed, then the target thread will not be suspended, and this routine leads to a null effect.

## Parameters

<i>timeout_mode</i>	The mode of the <i>timeout</i> parameter. It can either be set to <code>XN_RELATIVE</code> , <code>XN_ABSOLUTE</code> , or <code>XN_REALTIME</code> (see also <code>xntimer_start()</code> ).
<i>wchan</i>	The address of a pended resource. This parameter is used internally by the synchronization object implementation code to specify on which object the suspended thread pends. <code>NULL</code> is a legitimate value when this parameter does not apply to the current suspending mode (e.g. <code>XNSUSP</code> ).

## Note

If the target thread has received a Linux-originated signal, then this service immediately exits without suspending the thread, but raises the `XNBREAK` condition in its information mask.

## Tags

[unrestricted](#), [might-switch](#)

6.44.2.16 `xnthread_test_cancel()`

```
void xnthread_test_cancel (
    void ) [inline], [static]
```

Introduce a thread cancellation point.

Terminates the current thread if a cancellation request is pending for it, i.e. if `xnthread_cancel()` was called.

## Tags

[mode-unrestricted](#)

6.44.2.17 `xnthread_unblock()`

```
int xnthread_unblock (
    struct xnthread * thread )
```

Unblock a thread.

Breaks the thread out of any wait it is currently in. This call removes the XNDELAY and XNPEND suspensive conditions previously put by `xnthread_suspend()` on the target thread. If all suspensive conditions are gone, the thread is left in a READY state at which point it becomes eligible anew for scheduling.

Parameters

<code>thread</code>	The descriptor address of the unblocked thread.
---------------------	---

This call neither releases the thread from the XNSUSP, XNRELAX, XNDORMANT or XNHELD suspensive conditions.

When the thread resumes execution, the XNBREAK bit is set in the unblocked thread's information mask. Unblocking a non-blocked thread is perfectly harmless.

Returns

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

Tags

[unrestricted](#), [might-switch](#)

6.44.2.18 `xnthread_wait_period()`

```
int xnthread_wait_period (
    unsigned long * overruns_r )
```

Wait for the next periodic release point.

Make the current thread wait for the next periodic release point in the processor time line.

Parameters

<code>overruns↔ _r</code>	If non-NULL, <code>overruns_r</code> must be a pointer to a memory location which will be written with the count of pending overruns. This value is copied only when <code>xnthread_wait_period()</code> returns <code>-ETIMEDOUT</code> or success; the memory location remains unmodified otherwise. If NULL, this count will never be copied back.
-------------------------------	---



## Returns

0 is returned upon success; if *overruns\_r* is valid, zero is copied to the pointed memory location. Otherwise:

- -EWOULDBLOCK is returned if [xthread\\_set\\_periodic\(\)](#) has not previously been called for the calling thread.
- -EINTR is returned if [xthread\\_unblock\(\)](#) has been called for the waiting thread before the next periodic release point has been reached. In this case, the overrun counter is reset too.
- -ETIMEDOUT is returned if the timer has overrun, which indicates that one or more previous release points have been missed by the calling thread. If *overruns\_r* is valid, the count of pending overruns is copied to the pointed memory location.

## Tags

[primary-only](#), [might-switch](#)

## 6.45 Timer services

The Xenomai timer facility depends on a clock source (xnclock) for scheduling the next activation times.

Collaboration diagram for Timer services:



### Functions

- void [xntimer\\_destroy](#) (struct xntimer \*timer)

*Release a timer object.*

- static xnticks\_t [xntimer\\_interval](#) (struct xntimer \*timer)

*Return the timer interval value.*

- int [xntimer\\_start](#) (struct xntimer \*timer, xnticks\_t value, xnticks\_t interval, xntmode\_t mode)

*Arm a timer.*

- xnticks\_t [xntimer\\_get\\_date](#) (struct xntimer \*timer)

*Return the absolute expiration date.*

- static void [xntimer\\_stop](#) (struct xntimer \*timer)

*Disarm a timer.*

- static xnticks\_t [xntimer\\_get\\_timeout](#) (struct xntimer \*timer)

*Return the relative expiration date.*

- unsigned long long [xntimer\\_get\\_overruns](#) (struct xntimer \*timer, struct xnthread \*waiter, xnticks\_t now)

*Get the count of overruns for the last tick.*

- void [\\_\\_xntimer\\_migrate](#) (struct xntimer \*timer, struct xnsched \*sched)

*Migrate a timer.*

- int [xntimer\\_grab\\_hardware](#) (void)

*Grab the hardware timer on all real-time CPUs.*

- void `xntimer_release_hardware` (void)

*Release hardware timers.*

- void `xntimer_init` (struct `xntimer` \*timer, struct `xncklock` \*clock, void(\*handler)(struct `xntimer` \*timer), struct `xnsched` \*sched, int flags)

*Initialize a timer object.*

- static int `program_htick_shot` (unsigned long delay, struct `clock_event_device` \*cdev)

*Program next host tick as a Xenomai timer event.*

- static void `switch_htick_mode` (enum `clock_event_mode` mode, struct `clock_event_device` \*cdev)

*Tick mode switch emulation callback.*

### 6.45.1 Detailed Description

The Xenomai timer facility depends on a clock source (`xncklock`) for scheduling the next activation times.

The core provides and depends on a monotonic clock source (`nkcklock`) with nanosecond resolution, driving the platform timer hardware exposed by the interrupt pipeline.

### 6.45.2 Function Documentation

#### 6.45.2.1 `__xntimer_migrate()`

```
void __xntimer_migrate (
    struct xntimer * timer,
    struct xnsched * sched )
```

Migrate a timer.

This call migrates a timer to another cpu. In order to avoid pathological cases, it must be called from the CPU to which `timer` is currently attached.

Parameters

<code>timer</code>	The address of the timer object to be migrated.
<code>sched</code>	The address of the destination per-CPU scheduler slot.

Tags

[unrestricted](#), [atomic-entry](#)

6.45.2.2 `program_htick_shot()`

```
static int program_htick_shot (
    unsigned long delay,
    struct clock_event_device * cdev ) [static]
```

Program next host tick as a Xenomai timer event.

Program the next shot for the host tick on the current CPU. Emulation is done using a nucleus timer attached to the master timebase.

Parameters

<i>delay</i>	The time delta from the current date to the next tick, expressed as a count of nanoseconds.
<i>cdev</i>	An pointer to the clock device which notifies us.

Tags

[unrestricted](#)

6.45.2.3 `switch_htick_mode()`

```
void switch_htick_mode (
    enum clock_event_mode mode,
    struct clock_event_device * cdev ) [static]
```

Tick mode switch emulation callback.

Changes the host tick mode for the tick device of the current CPU.

Parameters

<i>mode</i>	The new mode to switch to. The possible values are:
-------------	---

- `CLOCK_EVT_MODE_ONESHOT`, for a switch to oneshot mode.
- `CLOCK_EVT_MODE_PERIODIC`, for a switch to periodic mode. The current implementation for the generic clockevent layer Linux exhibits should never downgrade from a oneshot to a periodic tick mode, so this mode should not be encountered. This said, the associated code is provided, basically for illustration purposes.
- `CLOCK_EVT_MODE_SHUTDOWN`, indicates the removal of the current tick device. Normally, the nucleus only interposes on tick devices which should never be shut down, so this mode should not be encountered.

Parameters

<i>cdev</i>	An opaque pointer to the clock device which notifies us.
-------------	--

Tags

[unrestricted](#)

Note

GENERIC\_CLOCKEVENTS is required from the host kernel.

#### 6.45.2.4 xntimer\_destroy()

```
void xntimer_destroy (  
    struct xntimer * timer )
```

Release a timer object.

Destroys a timer. After it has been destroyed, all resources associated with the timer have been released. The timer is automatically deactivated before deletion if active on entry.

Parameters

<i>timer</i>	The address of a valid timer descriptor.
--------------	--

Tags

[unrestricted](#)

#### 6.45.2.5 xntimer\_get\_date()

```
xnticks_t xntimer_get_date (  
    struct xntimer * timer )
```

Return the absolute expiration date.

Return the next expiration date of a timer as an absolute count of nanoseconds.

## Parameters

<i>timer</i>	The address of a valid timer descriptor.
--------------	--

## Returns

The expiration date in nanoseconds. The special value `XN_INFINITE` is returned if *timer* is currently disabled.

## Tags

[unrestricted](#), [atomic-entry](#)

6.45.2.6 `xntimer_get_overruns()`

```
unsigned long long xntimer_get_overruns (
    struct xntimer * timer,
    struct xnthread * waiter,
    xnticks_t now )
```

Get the count of overruns for the last tick.

This service returns the count of pending overruns for the last tick of a given timer, as measured by the difference between the expected expiry date of the timer and the date *now* passed as argument.

## Parameters

<i>timer</i>	The address of a valid timer descriptor.
<i>waiter</i>	The thread for which the overrun count is being collected.
<i>now</i>	current date (as <code>xnclock_read_raw(xntimer_clock(timer))</code> )

## Returns

the number of overruns of *timer* at date *now*

## Tags

[unrestricted](#), [atomic-entry](#)

6.45.2.7 `xntimer_get_timeout()`

```
xnticks_t xntimer_get_timeout (
    struct xntimer * timer ) [inline], [static]
```

Return the relative expiration date.

This call returns the count of nanoseconds remaining until the timer expires.

Parameters

<i>timer</i>	The address of a valid timer descriptor.
--------------	--

Returns

The count of nanoseconds until expiry. The special value `XN_INFINITE` is returned if *timer* is currently disabled. It might happen that the timer expires when this service runs (even if the associated handler has not been fired yet); in such a case, 1 is returned.

Tags

[unrestricted](#), [atomic-entry](#)

#### 6.45.2.8 `xntimer_grab_hardware()`

```
int xntimer_grab_hardware (
    void )
```

Grab the hardware timer on all real-time CPUs.

[xntimer\\_grab\\_hardware\(\)](#) grabs and tunes the hardware timer for all real-time CPUs.

Host tick emulation is performed for sharing the clock chip between Linux and Xenomai.

Returns

a positive value is returned on success, representing the duration of a Linux periodic tick expressed as a count of nanoseconds; zero should be returned when the Linux kernel does not undergo periodic timing on the given CPU (e.g. oneshot mode). Otherwise:

- `-EBUSY` is returned if the hardware timer has already been grabbed. [xntimer\\_release\\_hardware\(\)](#) must be issued before [xntimer\\_grab\\_hardware\(\)](#) is called again.
- `-ENODEV` is returned if the hardware timer cannot be used. This situation may occur after the kernel disabled the timer due to invalid calibration results; in such a case, such hardware is unusable for any timing duties.

Tags

[secondary-only](#)

6.45.2.9 `xntimer_init()`

```
void xntimer_init (
    struct xntimer * timer,
    struct xnclock * clock,
    void(*)(struct xntimer *timer) handler,
    struct xnsched * sched,
    int flags )
```

Initialize a timer object.

Creates a timer. When created, a timer is left disarmed; it must be started using `xntimer_start()` in order to be activated.

## Parameters

<i>timer</i>	The address of a timer descriptor the nucleus will use to store the object-specific data. This descriptor must always be valid while the object is active therefore it must be allocated in permanent memory.
<i>clock</i>	The clock the timer relates to. Xenomai defines a monotonic system clock, with nanosecond resolution, named <code>nkclock</code> . In addition, external clocks driven by other tick sources may be created dynamically if <code>CONFIG_XENO_OPT_EXTCLOCK</code> is defined.
<i>handler</i>	The routine to call upon expiration of the timer.
<i>sched</i>	An optional pointer to the per-CPU scheduler slot the new timer is affine to. If non-NULL, the timer will fire on the CPU <i>sched</i> is bound to, otherwise it will fire either on the current CPU if real-time, or on the first real-time CPU.
<i>flags</i>	A set of flags describing the timer. A set of clock gravity hints can be passed via the <i>flags</i> argument, used for optimizing the built-in heuristics aimed at latency reduction:

- `XNTIMER_IGRAVITY`, the timer activates a leaf timer handler.
- `XNTIMER_KGRAVITY`, the timer activates a kernel thread.
- `XNTIMER_UGRAVITY`, the timer activates a user-space thread.

There is no limitation on the number of timers which can be created/active concurrently.

## Tags

`unrestricted`

6.45.2.10 `xntimer_interval()`

```
xnticks_t xntimer_interval (
    struct xntimer * timer ) [inline], [static]
```

Return the timer interval value.

Return the timer interval value in nanoseconds.



## Parameters

<i>timer</i>	The address of a valid timer descriptor.
--------------	--

## Returns

The duration of a period in nanoseconds. The special value `XN_INFINITE` is returned if *timer* is currently disabled or one shot.

## Tags

[unrestricted](#), [atomic-entry](#)

6.45.2.11 `xntimer_release_hardware()`

```
void xntimer_release_hardware (
    void )
```

Release hardware timers.

Releases hardware timers previously grabbed by a call to [xntimer\\_grab\\_hardware\(\)](#).

## Tags

[secondary-only](#)

6.45.2.12 `xntimer_start()`

```
int xntimer_start (
    struct xntimer * timer,
    xnticks_t value,
    xnticks_t interval,
    xntmode_t mode )
```

Arm a timer.

Activates a timer so that the associated timeout handler will be fired after each expiration time. A timer can be either periodic or one-shot, depending on the reload value passed to this routine. The given timer must have been previously initialized.

A timer is attached to the clock specified in [xntimer\\_init\(\)](#).

## Parameters

<i>timer</i>	The address of a valid timer descriptor.
<i>value</i>	The date of the initial timer shot, expressed in nanoseconds.
<i>interval</i>	The reload value of the timer. It is a periodic interval value to be used for reprogramming the next timer shot, expressed in nanoseconds. If <i>interval</i> is equal to XN_INFINITE, the timer will not be reloaded after it has expired.
<i>mode</i>	The timer mode. It can be XN_RELATIVE if <i>value</i> shall be interpreted as a relative date, XN_ABSOLUTE for an absolute date based on the monotonic clock of the related time base (as returned by <code>xnclock_read_monotonic()</code> ), or XN_REALTIME if the absolute date is based on the adjustable real-time date for the relevant clock (obtained from <code>xnclock_read_realtime()</code> ).

## Returns

0 is returned upon success, or -ETIMEDOUT if an absolute date in the past has been given. In such an event, the timer is nevertheless armed for the next shot in the timeline if *interval* is different from XN\_INFINITE.

## Tags

[unrestricted](#), [atomic-entry](#)

6.45.2.13 `xntimer_stop()`

```
int xntimer_stop (
    struct xntimer * timer ) [inline], [static]
```

Disarm a timer.

This service deactivates a timer previously armed using `xntimer_start()`. Once disarmed, the timer can be subsequently re-armed using the latter service.

## Parameters

<i>timer</i>	The address of a valid timer descriptor.
--------------	--

## Tags

[unrestricted](#), [atomic-entry](#)

## 6.46 Virtual file services

Virtual files provide a mean to export Xenomai object states to user-space, based on common kernel interfaces.

Collaboration diagram for Virtual file services:



### Data Structures

- struct [xnvfile\\_lock\\_ops](#)  
*Vfile locking operations.*
- struct [xnvfile\\_regular\\_ops](#)  
*Regular vfile operation descriptor.*
- struct [xnvfile\\_regular\\_iterator](#)  
*Regular vfile iterator.*
- struct [xnvfile\\_snapshot\\_ops](#)  
*Snapshot vfile operation descriptor.*
- struct [xnvfile\\_rev\\_tag](#)  
*Snapshot revision tag.*
- struct [xnvfile\\_snapshot](#)  
*Snapshot vfile descriptor.*
- struct [xnvfile\\_snapshot\\_iterator](#)  
*Snapshot-driven vfile iterator.*

### Functions

- int [xnvfile\\_init\\_snapshot](#) (const char \*name, struct [xnvfile\\_snapshot](#) \*vfile, struct [xnvfile\\_directory](#) \*parent)  
  
*Initialize a snapshot-driven vfile.*
- int [xnvfile\\_init\\_regular](#) (const char \*name, struct [xnvfile\\_regular](#) \*vfile, struct [xnvfile\\_directory](#) \*parent)  
  
*Initialize a regular vfile.*
- int [xnvfile\\_init\\_dir](#) (const char \*name, struct [xnvfile\\_directory](#) \*vdir, struct [xnvfile\\_directory](#) \*parent)  
  
*Initialize a virtual directory entry.*
- int [xnvfile\\_init\\_link](#) (const char \*from, const char \*to, struct [xnvfile\\_link](#) \*vlink, struct [xnvfile\\_directory](#) \*parent)  
  
*Initialize a virtual link entry.*

- void `xnvfile_destroy` (struct `xnvfile` \*vfile)

*Removes a virtual file entry.*

- `ssize_t xnvfile_get_blob` (struct `xnvfile_input` \*input, void \*data, `size_t` size)

*Read in a data bulk written to the vfile.*

- `ssize_t xnvfile_get_string` (struct `xnvfile_input` \*input, char \*s, `size_t` maxlen)

*Read in a C-string written to the vfile.*

- `ssize_t xnvfile_get_integer` (struct `xnvfile_input` \*input, long \*valp)

*Evaluate the string written to the vfile as a long integer.*

## Variables

- struct `xnvfile_directory` `cobalt_vfroot`

*Xenomai vfile root directory.*

- struct `xnvfile_directory` `cobalt_vfroot`

*Xenomai vfile root directory.*

### 6.46.1 Detailed Description

Virtual files provide a mean to export Xenomai object states to user-space, based on common kernel interfaces.

This encapsulation is aimed at:

- supporting consistent collection of very large record-based output, without incurring latency peaks for undergoing real-time activities.
- in the future, hiding discrepancies between linux kernel releases, regarding the proper way to export kernel object states to userland, either via the `/proc` interface or by any other mean.

This virtual file implementation offers record-based read support based on `seq_files`, single-buffer write support, directory and link handling, all visible from the `/proc` namespace.

The vfile support exposes four filesystem object types:

- snapshot-driven file (struct `xnvfile_snapshot`). This is commonly used to export real-time object states via the `/proc` filesystem. To minimize the latency involved in protecting the vfile routines from changes applied by real-time code on such objects, a snapshot of the data to output is first taken under proper locking, before the collected data is formatted and sent out in a lockless manner.

Because a large number of records may have to be output, the data collection phase is not strictly atomic as a whole, but only protected at record level. The vfile implementation can be notified of updates to the underlying data set, and restart the collection from scratch until the snapshot is fully consistent.

- regular sequential file (struct `xnvfile_regular`). This is basically an encapsulated sequential file object as available from the host kernel (i.e. `seq_file`), with a few additional features to make it more handy in a Xenomai environment, like implicit locking support and shortened declaration for simplest, single-record output.
- virtual link (struct `xnvfile_link`). This is a symbolic link feature integrated with the vfile semantics. The link target is computed dynamically at creation time from a user-given helper routine.
- virtual directory (struct `xnvfile_directory`). A directory object, which can be used to create a hierarchy for ordering a set of vfile objects.

## 6.46.2 Function Documentation

### 6.46.2.1 `xnvfile_destroy()`

```
void xnvfile_destroy (
    struct xnvfile * vfile )
```

Removes a virtual file entry.

Parameters

<i>vfile</i>	A pointer to the virtual file descriptor to remove.
--------------	---

Tags

[secondary-only](#)

### 6.46.2.2 `xnvfile_get_blob()`

```
ssize_t xnvfile_get_blob (
    struct xnvfile_input * input,
    void * data,
    size_t size )
```

Read in a data bulk written to the vfile.

When writing to a vfile, the associated `store()` handler from the [snapshot-driven vfile](#) or [regular vfile](#) is called, with a single argument describing the input data. `xnvfile_get_blob()` retrieves this data as an untyped binary blob, and copies it back to the caller's buffer.

Parameters

<i>input</i>	A pointer to the input descriptor passed to the <code>store()</code> handler.
<i>data</i>	The address of the destination buffer to copy the input data to.
<i>size</i>	The maximum number of bytes to copy to the destination buffer. If <i>size</i> is larger than the actual data size, the input is truncated to <i>size</i> .

Returns

The number of bytes read and copied to the destination buffer upon success. Otherwise, a negative error code is returned:

- `-EFAULT` indicates an invalid source buffer address.

Tags

[secondary-only](#)

### 6.46.2.3 xnvfile\_get\_integer()

```
ssize_t xnvfile_get_integer (
    struct xnvfile_input * input,
    long * valp )
```

Evaluate the string written to the vfile as a long integer.

When writing to a vfile, the associated store() handler from the [snapshot-driven vfile](#) or [regular vfile](#) is called, with a single argument describing the input data. [xnvfile\\_get\\_integer\(\)](#) retrieves and interprets this data as a long integer, and copies the resulting value back to *valp*.

The long integer can be expressed in decimal, octal or hexadecimal bases depending on the prefix found.

Parameters

<i>input</i>	A pointer to the input descriptor passed to the store() handler.
<i>valp</i>	The address of a long integer variable to receive the value.

Returns

The number of characters read while evaluating the input as a long integer upon success. Otherwise, a negative error code is returned:

- -EINVAL indicates a parse error on the input stream; the written text cannot be evaluated as a long integer.
- -EFAULT indicates an invalid source buffer address.

Tags

[secondary-only](#)

### 6.46.2.4 xnvfile\_get\_string()

```
ssize_t xnvfile_get_string (
    struct xnvfile_input * input,
    char * s,
    size_t maxlen )
```

Read in a C-string written to the vfile.

When writing to a vfile, the associated store() handler from the [snapshot-driven vfile](#) or [regular vfile](#) is called, with a single argument describing the input data. [xnvfile\\_get\\_string\(\)](#) retrieves this data as a null-terminated character string, and copies it back to the caller's buffer.

## Parameters

<i>input</i>	A pointer to the input descriptor passed to the store() handler.
<i>s</i>	The address of the destination string buffer to copy the input data to.
<i>maxlen</i>	The maximum number of bytes to copy to the destination buffer, including the ending null character. If <i>maxlen</i> is larger than the actual string length, the input is truncated to <i>maxlen</i> .

## Returns

The number of characters read upon success. Otherwise, a negative error code is returned:

- -EFAULT indicates an invalid source buffer address.

## Tags

[secondary-only](#)

## 6.46.2.5 xnvfile\_init\_dir()

```
int xnvfile_init_dir (
    const char * name,
    struct xnvfile_directory * vdir,
    struct xnvfile_directory * parent )
```

Initialize a virtual directory entry.

## Parameters

<i>name</i>	The name which should appear in the pseudo-filesystem, identifying the vdir entry.
<i>vdir</i>	A pointer to the virtual directory descriptor to initialize.
<i>parent</i>	A pointer to a virtual directory descriptor standing for the parent directory of the new vdir. If NULL, the /proc root directory will be used. /proc/xenomai is mapped on the globally available <i>cobalt_vfroot</i> vdir.

## Returns

0 is returned on success. Otherwise:

- -ENOMEM is returned if the virtual directory entry cannot be created in the /proc hierarchy.

## Tags

[secondary-only](#)

## 6.46.2.6 xnvfile\_init\_link()

```
int xnvfile_init_link (
    const char * from,
    const char * to,
    struct xnvfile_link * vlink,
    struct xnvfile_directory * parent )
```

Initialize a virtual link entry.

## Parameters

<i>from</i>	The name which should appear in the pseudo-filesystem, identifying the vlink entry.
<i>to</i>	The target file name which should be referred to symbolically by <i>name</i> .
<i>vlink</i>	A pointer to the virtual link descriptor to initialize.
<i>parent</i>	A pointer to a virtual directory descriptor standing for the parent directory of the new vlink. If NULL, the /proc root directory will be used. /proc/xenomai is mapped on the globally available <i>cobalt_vfroot</i> vdir.

## Returns

0 is returned on success. Otherwise:

- -ENOMEM is returned if the virtual link entry cannot be created in the /proc hierarchy.

## Tags

[secondary-only](#)

## 6.46.2.7 xnvfile\_init\_regular()

```
int xnvfile_init_regular (
    const char * name,
    struct xnvfile_regular * vfile,
    struct xnvfile_directory * parent )
```

Initialize a regular vfile.

## Parameters

<i>name</i>	The name which should appear in the pseudo-filesystem, identifying the vfile entry.
<i>vfile</i>	A pointer to a vfile descriptor to initialize from. The following fields in this structure should be filled in prior to call this routine:

- *.privsz* is the size (in bytes) of the private data area to be reserved in the [vfile iterator](#). A NULL



value indicates that no private area should be reserved.

- `entry.lockops` is a pointer to a [locking descriptor](#)", defining the lock and unlock operations for the vfile. This pointer may be left to NULL, in which case no locking will be applied.
- `.ops` is a pointer to an [operation descriptor](#).

#### Parameters

<i>parent</i>	A pointer to a virtual directory descriptor; the vfile entry will be created into this directory. If NULL, the <code>/proc</code> root directory will be used. <code>/proc/xenomai</code> is mapped on the globally available <code>cobalt_vfroot</code> vdir.
---------------	--

#### Returns

0 is returned on success. Otherwise:

- `-ENOMEM` is returned if the virtual file entry cannot be created in the `/proc` hierarchy.

#### Tags

[secondary-only](#)

#### 6.46.2.8 `xnvfile_init_snapshot()`

```
int xnvfile_init_snapshot (
    const char * name,
    struct xnvfile_snapshot * vfile,
    struct xnvfile_directory * parent )
```

Initialize a snapshot-driven vfile.

#### Parameters

<i>name</i>	The name which should appear in the pseudo-filesystem, identifying the vfile entry.
<i>vfile</i>	A pointer to a vfile descriptor to initialize from. The following fields in this structure should be filled in prior to call this routine:

- `.privsz` is the size (in bytes) of the private data area to be reserved in the [vfile iterator](#). A NULL value indicates that no private area should be reserved.
- `.datasz` is the size (in bytes) of a single record to be collected by the [next\(\) handler](#) from the [operation descriptor](#).
- `.tag` is a pointer to a mandatory vfile revision tag structure (struct [xnvfile\\_rev\\_tag](#)). This tag will be monitored for changes by the vfile core while collecting data to output, so that any update detected will cause the current snapshot data to be dropped, and the collection to restart from the beginning.

To this end, any change to the data which may be part of the collected records, should also invoke `xnvfile_touch()` on the associated tag.

- `entry.lockops` is a pointer to a [lock descriptor](#), defining the lock and unlock operations for the vfile. This pointer may be left to NULL, in which case the operations on the nucleus lock (i.e. `nklock`) will be used internally around calls to data collection handlers (see [operation descriptor](#)).
- `.ops` is a pointer to an [operation descriptor](#).

#### Parameters

<i>parent</i>	A pointer to a virtual directory descriptor; the vfile entry will be created into this directory. If NULL, the <code>/proc</code> root directory will be used. <code>/proc/xenomai</code> is mapped on the globally available <code>cobalt_vfroot</code> vdir.
---------------	--

#### Returns

0 is returned on success. Otherwise:

- `-ENOMEM` is returned if the virtual file entry cannot be created in the `/proc` hierarchy.

#### Tags

[secondary-only](#)

### 6.46.3 Variable Documentation

#### 6.46.3.1 `cobalt_vfroot` [1/2]

```
struct xnvfile_directory cobalt_vfroot
```

Xenomai vfile root directory.

This vdir maps the `/proc/xenomai` directory. It can be used to create a hierarchy of Xenomai-related vfiles under this root.

#### 6.46.3.2 `cobalt_vfroot` [2/2]

```
struct xnvfile_directory cobalt_vfroot
```


Xenomai vfile root directory.

This vdir maps the `/proc/xenomai` directory. It can be used to create a hierarchy of Xenomai-related vfiles under this root.

## 6.47 Analogy framework

A RTDM-based interface for implementing DAQ card drivers.

Collaboration diagram for Analogy framework:



### Modules

- [Driver API](#)  
*Programming interface provided to DAQ card drivers.*
- [Analogy user API](#)

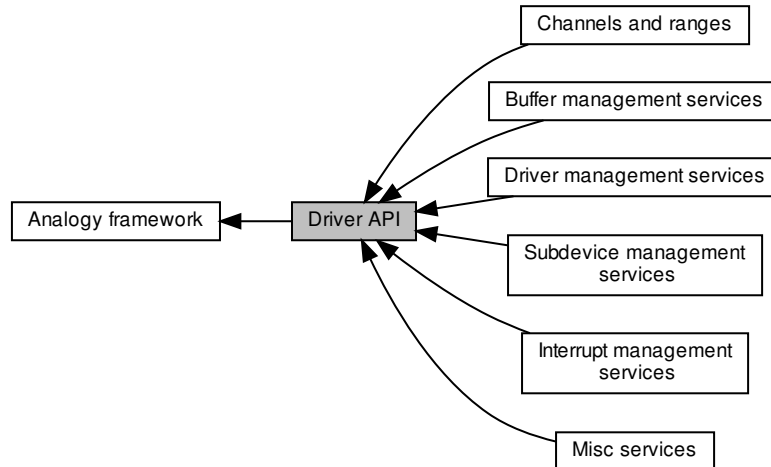
#### 6.47.1 Detailed Description

A RTDM-based interface for implementing DAQ card drivers.

## 6.48 Driver API

Programming interface provided to DAQ card drivers.

Collaboration diagram for Driver API:



### Modules

- [Channels and ranges](#)

*Channels*

- [Driver management services](#)  
*Analogy driver registration / unregistration.*
- [Subdevice management services](#)  
*Subdevice declaration in a driver.*
- [Buffer management services](#)  
*Buffer management services.*
- [Interrupt management services](#)
- [Misc services](#)


### 6.48.1 Detailed Description

Programming interface provided to DAQ card drivers.

## 6.49 Driver management services

Analogy driver registration / unregistration.

Collaboration diagram for Driver management services:



### Functions

- int [a4l\\_register\\_drv](#) (struct [a4l\\_driver](#) \*drv)  
*Register an Analogy driver.*
- int [a4l\\_unregister\\_drv](#) (struct [a4l\\_driver](#) \*drv)  
*Unregister an Analogy driver.*

### 6.49.1 Detailed Description

Analogy driver registration / unregistration.

In a common Linux char driver, the developer has to register a fops structure filled with callbacks for read / write / mmap / ioctl operations.

Analogy drivers do not have to implement read / write / mmap / ioctl functions, these procedures are implemented in the Analogy generic layer. Then, the transfers between user-space and kernel-space are already managed. Analogy drivers work with commands and instructions which are some kind of more dedicated read / write operations. And, instead of registering a fops structure, a Analogy driver must register some [a4l\\_driver](#) structure.

### 6.49.2 Function Documentation

#### 6.49.2.1 a4l\_register\_drv()

```
int a4l_register_drv (  
    struct a4l\_driver * drv )
```

Register an Analogy driver.

After initialising a driver structure, the driver must be made available so as to be attached.

## Parameters

in	<i>drv</i>	Driver descriptor structure
----	------------	-----------------------------

## Returns

0 on success, otherwise negative error code.

References `a4l_driver::board_name`, and `rtdm_available()`.

6.49.2.2 `a4l_unregister_drv()`

```
int a4l_unregister_drv (  
    struct a4l_driver * drv )
```

Unregister an Analogy driver.

This function removes the driver descriptor from the Analogy driver list. The driver cannot be attached anymore.

## Parameters

in	<i>drv</i>	Driver descriptor structure
----	------------	-----------------------------

## Returns


0 on success, otherwise negative error code.

References `a4l_driver::board_name`.

## 6.50 Subdevice management services

Subdevice declaration in a driver.

Collaboration diagram for Subdevice management services:



### Functions

- struct `a4l_subdevice` \* `a4l_alloc_subd` (int sizeof\_priv, void(\*setup)(struct `a4l_subdevice` \*))  
*Allocate a subdevice descriptor.*
- int `a4l_add_subd` (struct `a4l_device` \*dev, struct `a4l_subdevice` \*subd)  
*Add a subdevice to the driver descriptor.*
- struct `a4l_subdevice` \* `a4l_get_subd` (struct `a4l_device` \*dev, int idx)  
*Get a pointer to the subdevice descriptor referenced by its registration index.*

### Subdevices types

Flags to define the subdevice type

- #define `A4L_SUBD_UNUSED` (`A4L_SUBD_MASK_SPECIAL|0x1`)  
*Unused subdevice.*
- #define `A4L_SUBD_AI` (`A4L_SUBD_MASK_READ|0x2`)  
*Analog input subdevice.*
- #define `A4L_SUBD_AO` (`A4L_SUBD_MASK_WRITE|0x4`)  
*Analog output subdevice.*
- #define `A4L_SUBD_DI` (`A4L_SUBD_MASK_READ|0x8`)  
*Digital input subdevice.*
- #define `A4L_SUBD_DO` (`A4L_SUBD_MASK_WRITE|0x10`)  
*Digital output subdevice.*
- #define `A4L_SUBD_DIO` (`A4L_SUBD_MASK_SPECIAL|0x20`)  
*Digital input/output subdevice.*
- #define `A4L_SUBD_COUNTER` (`A4L_SUBD_MASK_SPECIAL|0x40`)  
*Counter subdevice.*
- #define `A4L_SUBD_TIMER` (`A4L_SUBD_MASK_SPECIAL|0x80`)  
*Timer subdevice.*
- #define `A4L_SUBD_MEMORY` (`A4L_SUBD_MASK_SPECIAL|0x100`)  
*Memory, EEPROM, DPRAM.*
- #define `A4L_SUBD_CALIB` (`A4L_SUBD_MASK_SPECIAL|0x200`)

- *Calibration subdevice DACs.*
- #define `A4L_SUBD_PROC` (`A4L_SUBD_MASK_SPECIAL|0x400`)  
*Processor, DSP.*
- #define `A4L_SUBD_SERIAL` (`A4L_SUBD_MASK_SPECIAL|0x800`)  
*Serial IO subdevice.*
- #define `A4L_SUBD_TYPES`  
*Mask which gathers all the types.*

## Subdevice features

Flags to define the subdevice's capabilities

- #define `A4L_SUBD_CMD` `0x1000`  
*The subdevice can handle command (i.e it can perform asynchronous acquisition)*
- #define `A4L_SUBD_MMAP` `0x8000`  
*The subdevice support mmap operations (technically, any driver can do it; however, the developer might want that his driver must be accessed through read / write.*

## Subdevice status

Flags to define the subdevice's status

- #define `A4L_SUBD_BUSY_NR` `0`  
*The subdevice is busy, a synchronous or an asynchronous acquisition is occurring.*
- #define `A4L_SUBD_BUSY` (`1 << A4L_SUBD_BUSY_NR`)  
*The subdevice is busy, a synchronous or an asynchronous acquisition is occurring.*
- #define `A4L_SUBD_CLEAN_NR` `1`  
*The subdevice is about to be cleaned in the middle of the detach procedure.*
- #define `A4L_SUBD_CLEAN` (`1 << A4L_SUBD_CLEAN_NR`)  
*The subdevice is busy, a synchronous or an asynchronous acquisition is occurring.*

### 6.50.1 Detailed Description

Subdevice declaration in a driver.

The subdevice structure is the most complex one in the Analogy driver layer. It contains some description fields to fill and some callbacks to declare.

The description fields are:

- `flags`: to define the subdevice type and its capabilities;
- `chan_desc`: to describe the channels which compose the subdevice;
- `rng_desc`: to declare the usable ranges;

The functions callbakcs are:



- `do_cmd()` and `do_cmdtest()`: to perform asynchronous acquisitions thanks to commands;
- `cancel()`: to abort a working asynchronous acquisition;
- `munge()`: to apply modifications on the data freshly acquired during an asynchronous transfer. Warning: using this feature with can significantly reduce the performances (if the munge operation is complex, it will trigger high CPU charge and if the acquisition device is DMA capable, many cache-misses and cache-replaces will occur (the benefits of the DMA controller will vanish);
- `trigger()`: optionnaly to launch an asynchronous acquisition;
- `insn_read()`, `insn_write()`, `insn_bits()`, `insn_config()`: to perform synchronous acquisition operations.

Once the subdevice is filled, it must be inserted into the driver structure thanks to [a4l\\_add\\_subd\(\)](#).

## 6.50.2 Function Documentation

### 6.50.2.1 a4l\_add\_subd()

```
int a4l_add_subd (
    struct a4l_device * dev,
    struct a4l_subdevice * subd )
```

Add a subdevice to the driver descriptor.

Once the driver descriptor structure is initialized, the function [a4l\\_add\\_subd\(\)](#) must be used so to add some subdevices to the driver.

#### Parameters

in	<i>dev</i>	Device descriptor structure
in	<i>subd</i>	Subdevice descriptor structure

#### Returns

the index with which the subdevice has been registered, in case of error a negative error code is returned.

References `a4l_subdevice::dev`, `a4l_subdevice::idx`, and `a4l_subdevice::list`.

### 6.50.2.2 a4l\_alloc\_subd()

```
struct a4l_subdevice* a4l_alloc_subd (
    int sizeof_priv,
    void(*)(struct a4l_subdevice *) setup )
```

Allocate a subdevice descriptor.

This is a helper function so as to get a suitable subdevice descriptor

## Parameters

in	<i>sizeof_priv</i>	Size of the subdevice's private data
in	<i>setup</i>	Setup function to be called after the allocation

## Returns

the index with which the subdevice has been registered, in case of error a negative error code is returned.

References `rtdm_malloc()`.

6.50.2.3 `a4l_get_subd()`

```
struct a4l_subdevice* a4l_get_subd (
    struct a4l_device * dev,
    int idx )
```

Get a pointer to the subdevice descriptor referenced by its registration index.

This function is scarcely useful as all the drivers callbacks get the related subdevice descriptor as first argument. This function is not optimized, it goes through a linked list to get the proper pointer. So it must not be used in real-time context but at initialization / cleanup time (attach / detach).

## Parameters

in	<i>dev</i>	Device descriptor structure
in	<i>idx</i>	Subdevice index

## Returns

0 on success, otherwise negative error code.

## 6.51 Buffer management services

Buffer management services.

Collaboration diagram for Buffer management services:



### Functions

- int [a4l\\_buf\\_prepare\\_absput](#) (struct [a4l\\_subdevice](#) \*subd, unsigned long count)  
*Update the absolute count of data sent from the device to the buffer since the start of the acquisition and after the next DMA shot.*
- int [a4l\\_buf\\_commit\\_absput](#) (struct [a4l\\_subdevice](#) \*subd, unsigned long count)  
*Set the absolute count of data which was sent from the device to the buffer since the start of the acquisition and until the last DMA shot.*
- int [a4l\\_buf\\_prepare\\_put](#) (struct [a4l\\_subdevice](#) \*subd, unsigned long count)  
*Set the count of data which is to be sent to the buffer at the next DMA shot.*
- int [a4l\\_buf\\_commit\\_put](#) (struct [a4l\\_subdevice](#) \*subd, unsigned long count)  
*Set the count of data sent to the buffer during the last completed DMA shots.*
- int [a4l\\_buf\\_put](#) (struct [a4l\\_subdevice](#) \*subd, void \*bufdata, unsigned long count)  
*Copy some data from the device driver to the buffer.*
- int [a4l\\_buf\\_prepare\\_absget](#) (struct [a4l\\_subdevice](#) \*subd, unsigned long count)  
*Update the absolute count of data sent from the buffer to the device since the start of the acquisition and after the next DMA shot.*
- int [a4l\\_buf\\_commit\\_absget](#) (struct [a4l\\_subdevice](#) \*subd, unsigned long count)  
*Set the absolute count of data which was sent from the buffer to the device since the start of the acquisition and until the last DMA shot.*
- int [a4l\\_buf\\_prepare\\_get](#) (struct [a4l\\_subdevice](#) \*subd, unsigned long count)  
*Set the count of data which is to be sent from the buffer to the device at the next DMA shot.*
- int [a4l\\_buf\\_commit\\_get](#) (struct [a4l\\_subdevice](#) \*subd, unsigned long count)  
*Set the count of data sent from the buffer to the device during the last completed DMA shots.*
- int [a4l\\_buf\\_get](#) (struct [a4l\\_subdevice](#) \*subd, void \*bufdata, unsigned long count)  
*Copy some data from the buffer to the device driver.*
- int [a4l\\_buf\\_evt](#) (struct [a4l\\_subdevice](#) \*subd, unsigned long evts)  
*Signal some event(s) to a user-space program involved in some read / write operation.*
- unsigned long [a4l\\_buf\\_count](#) (struct [a4l\\_subdevice](#) \*subd)  
*Get the data amount available in the Analogy buffer.*
- struct [a4l\\_cmd\\_desc](#) \* [a4l\\_get\\_cmd](#) (struct [a4l\\_subdevice](#) \*subd)  
*Get the current Analogy command descriptor.*
- int [a4l\\_get\\_chan](#) (struct [a4l\\_subdevice](#) \*subd)  
*Get the channel index according to its type.*

### 6.51.1 Detailed Description

Buffer management services.

The buffer is the key component of the Analogy infrastructure. It manages transfers between the user-space and the Analogy drivers thanks to generic functions which are described hereafter. Thanks to the buffer subsystem, the driver developer does not have to care about the way the user program retrieves or sends data.

To write a classical char driver, the developer has to fill a fops structure so as to provide transfer operations to the user program (read, write, ioctl and mmap if need be).

The Analogy infrastructure manages the whole interface with the userspace; the common read, write, mmap, etc. callbacks are generic Analogy functions. These functions manage (and perform, if need be) transfers between the user-space and an asynchronous buffer thanks to lockless mechanisms.

Consequently, the developer has to use the proper buffer functions in order to write / read acquired data into / from the asynchronous buffer.

Here are listed the functions:

- `a4l_buf_prepare_(abs)put()` and `a4l_buf_commit_(abs)put()`
- `a4l_buf_prepare_(abs)get()` and `a4l_buf_commit_(abs)get()`
- [a4l\\_buf\\_put\(\)](#)
- [a4l\\_buf\\_get\(\)](#)
- [a4l\\_buf\\_evt\(\)](#).

The functions count might seem high; however, the developer needs a few of them to write a driver. Having so many functions enables to manage any transfer cases:

- If some DMA controller is available, there is no need to make the driver copy the acquired data into the asynchronous buffer, the DMA controller must directly trigger DMA shots into / from the buffer. In that case, a function `a4l_buf_prepare_*`() must be used so as to set up the DMA transfer and a function `a4l_buf_commit_*`() has to be called to complete the transfer().
- For DMA controllers which need to work with global counter (the transferred data count since the beginning of the acquisition), the functions `a4l_buf_*_abs_*`() have been made available.
- If no DMA controller is available, the driver has to perform the copy between the hardware component and the asynchronous buffer. In such cases, the functions [a4l\\_buf\\_get\(\)](#) and [a4l\\_buf\\_put\(\)](#) are useful.

### 6.51.2 Function Documentation

#### 6.51.2.1 `a4l_buf_commit_absget()`

```
int a4l_buf_commit_absget (
    struct a4l_subdevice * subd,
    unsigned long count )
```

Set the absolute count of data which was sent from the buffer to the device since the start of the acquisition and until the last DMA shot.

The functions `a4l_buf_prepare_(abs)put()`, `a4l_buf_commit_(abs)put()`, `a4l_buf_prepare_(abs)get()` and `a4l_buf_commit_(abs)get()` have been made available for DMA transfers. In such situations, no data copy is needed between the Analogy buffer and the device as some DMA controller is in charge of performing data shots from / to the Analogy buffer. However, some pointers still have to be updated so as to monitor the transfers.

## Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>count</i>	The data count transferred to the device during the last DMA shot plus the data count which have been sent since the beginning of the acquisition

## Returns

0 on success, otherwise negative error code.

References A4L\_SUBD\_BUSY\_NR, a4l\_subdevice::buf, and a4l\_subdevice::status.

## 6.51.2.2 a4l\_buf\_commit\_absput()

```
int a4l_buf_commit_absput (
    struct a4l_subdevice * subd,
    unsigned long count )
```

Set the absolute count of data which was sent from the device to the buffer since the start of the acquisition and until the last DMA shot.

The functions a4l\_buf\_prepare\_(abs)put(), a4l\_buf\_commit\_(abs)put(), a4l\_buf\_prepare\_(abs)get() and a4l\_buf\_commit\_(abs)get() have been made available for DMA transfers. In such situations, no data copy is needed between the Analogy buffer and the device as some DMA controller is in charge of performing data shots from / to the Analogy buffer. However, some pointers still have to be updated so as to monitor the tranfers.

## Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>count</i>	The data count transferred to the buffer during the last DMA shot plus the data count which have been sent / retrieved since the beginning of the acquisition

## Returns

0 on success, otherwise negative error code.

References A4L\_SUBD\_BUSY\_NR, a4l\_subdevice::buf, and a4l\_subdevice::status.

## 6.51.2.3 a4l\_buf\_commit\_get()

```
int a4l_buf_commit_get (
    struct a4l_subdevice * subd,
    unsigned long count )
```

Set the count of data sent from the buffer to the device during the last completed DMA shots.

The functions `a4l_buf_prepare_(abs)put()`, `a4l_buf_commit_(abs)put()`, `a4l_buf_prepare_(abs)get()` and `a4l_buf_commit_(abs)get()` have been made available for DMA transfers. In such situations, no data copy is needed between the Analogy buffer and the device as some DMA controller is in charge of performing data shots from / to the Analogy buffer. However, some pointers still have to be updated so as to monitor the tranfers.

Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>count</i>	The amount of data transferred

Returns

0 on success, otherwise negative error code.

References `A4L_SUBD_BUSY_NR`, `a4l_subdevice::buf`, and `a4l_subdevice::status`.

#### 6.51.2.4 `a4l_buf_commit_put()`

```
int a4l_buf_commit_put (
    struct a4l_subdevice * subd,
    unsigned long count )
```

Set the count of data sent to the buffer during the last completed DMA shots.

The functions `a4l_buf_prepare_(abs)put()`, `a4l_buf_commit_(abs)put()`, `a4l_buf_prepare_(abs)get()` and `a4l_buf_commit_(abs)get()` have been made available for DMA transfers. In such situations, no data copy is needed between the Analogy buffer and the device as some DMA controller is in charge of performing data shots from / to the Analogy buffer. However, some pointers still have to be updated so as to monitor the tranfers.

Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>count</i>	The amount of data transferred

Returns

0 on success, otherwise negative error code.

References `A4L_SUBD_BUSY_NR`, `a4l_subdevice::buf`, and `a4l_subdevice::status`.

## 6.51.2.5 a4l\_buf\_count()

```
unsigned long a4l_buf_count (  
    struct a4l_subdevice * subd )
```

Get the data amount available in the Analogy buffer.

## Parameters

in	<i>subd</i>	Subdevice descriptor structure
----	-------------	--------------------------------

## Returns

the amount of data available in the Analogy buffer.

References A4L\_SUBD\_BUSY\_NR, a4l\_subdevice::buf, and a4l\_subdevice::status.

## 6.51.2.6 a4l\_buf\_evt()

```
int a4l_buf_evt (
    struct a4l_subdevice * subd,
    unsigned long evts )
```

Signal some event(s) to a user-space program involved in some read / write operation.

The function [a4l\\_buf\\_evt\(\)](#) is useful in many cases:

- To wake-up a process waiting for some data to read.
- To wake-up a process waiting for some data to write.
- To notify the user-process an error has occurred during the acquisition.

## Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>evts</i>	Some specific event to notify: <ul style="list-style-type: none"> <li>• A4L_BUF_ERROR to indicate some error has occurred during the transfer</li> <li>• A4L_BUF_EOA to indicate the acquisition is complete (this event is automatically set, it should not be used).</li> </ul>

## Returns

0 on success, otherwise negative error code.

References A4L\_SUBD\_BUSY\_NR, a4l\_subdevice::buf, and a4l\_subdevice::status.

## 6.51.2.7 a4l\_buf\_get()

```
int a4l_buf_get (
    struct a4l_subdevice * subd,
```



```
void * bufdata,
unsigned long count )
```

Copy some data from the buffer to the device driver.

The function `a4l_buf_get()` must copy data coming from the Analogy buffer to some acquisition device. This ring-buffer is an intermediate area between the device driver and the user-space program, which is supposed to provide the data to send to the device.

Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>bufdata</i>	The data buffer to copy into the Analogy buffer
in	<i>count</i>	The amount of data to copy

Returns

0 on success, otherwise negative error code.

References `A4L_SUBD_BUSY_NR`, `a4l_subdevice::buf`, and `a4l_subdevice::status`.

#### 6.51.2.8 `a4l_buf_prepare_absget()`

```
int a4l_buf_prepare_absget (
    struct a4l_subdevice * subd,
    unsigned long count )
```

Update the absolute count of data sent from the buffer to the device since the start of the acquisition and after the next DMA shot.

The functions `a4l_buf_prepare_(abs)put()`, `a4l_buf_commit_(abs)put()`, `a4l_buf_prepare_(abs)get()` and `a4l_buf_commit_(abs)get()` have been made available for DMA transfers. In such situations, no data copy is needed between the Analogy buffer and the device as some DMA controller is in charge of performing data shots from / to the Analogy buffer. However, some pointers still have to be updated so as to monitor the tranfers.

Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>count</i>	The data count to be transferred during the next DMA shot plus the data count which have been copied since the start of the acquisition

Returns

0 on success, otherwise negative error code.

References `A4L_SUBD_BUSY_NR`, `a4l_subdevice::buf`, and `a4l_subdevice::status`.

## 6.51.2.9 a4l\_buf\_prepare\_absput()

```
int a4l_buf_prepare_absput (
    struct a4l_subdevice * subd,
    unsigned long count )
```

Update the absolute count of data sent from the device to the buffer since the start of the acquisition and after the next DMA shot.

The functions `a4l_buf_prepare_(abs)put()`, `a4l_buf_commit_(abs)put()`, `a4l_buf_prepare_(abs)get()` and `a4l_buf_commit_(abs)get()` have been made available for DMA transfers. In such situations, no data copy is needed between the Analogy buffer and the device as some DMA controller is in charge of performing data shots from / to the Analogy buffer. However, some pointers still have to be updated so as to monitor the tranfers.

## Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>count</i>	The data count to be transferred during the next DMA shot plus the data count which have been copied since the start of the acquisition

## Returns

0 on success, otherwise negative error code.

References `A4L_SUBD_BUSY_NR`, `a4l_subdevice::buf`, and `a4l_subdevice::status`.

## 6.51.2.10 a4l\_buf\_prepare\_get()

```
int a4l_buf_prepare_get (
    struct a4l_subdevice * subd,
    unsigned long count )
```

Set the count of data which is to be sent from the buffer to the device at the next DMA shot.

The functions `a4l_buf_prepare_(abs)put()`, `a4l_buf_commit_(abs)put()`, `a4l_buf_prepare_(abs)get()` and `a4l_buf_commit_(abs)get()` have been made available for DMA transfers. In such situations, no data copy is needed between the Analogy buffer and the device as some DMA controller is in charge of performing data shots from / to the Analogy buffer. However, some pointers still have to be updated so as to monitor the tranfers.

## Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>count</i>	The data count to be transferred

## Returns

0 on success, otherwise negative error code.

References A4L\_SUBD\_BUSY\_NR, a4l\_subdevice::buf, and a4l\_subdevice::status.

#### 6.51.2.11 a4l\_buf\_prepare\_put()

```
int a4l_buf_prepare_put (
    struct a4l_subdevice * subd,
    unsigned long count )
```

Set the count of data which is to be sent to the buffer at the next DMA shot.

The functions `a4l_buf_prepare_(abs)put()`, `a4l_buf_commit_(abs)put()`, `a4l_buf_prepare_(abs)get()` and `a4l_buf_commit_(abs)get()` have been made available for DMA transfers. In such situations, no data copy is needed between the Analogy buffer and the device as some DMA controller is in charge of performing data shots from / to the Analogy buffer. However, some pointers still have to be updated so as to monitor the transfers.

Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>count</i>	The data count to be transferred

Returns

0 on success, otherwise negative error code.

References A4L\_SUBD\_BUSY\_NR, a4l\_subdevice::buf, and a4l\_subdevice::status.

#### 6.51.2.12 a4l\_buf\_put()

```
int a4l_buf_put (
    struct a4l_subdevice * subd,
    void * bufdata,
    unsigned long count )
```

Copy some data from the device driver to the buffer.

The function `a4l_buf_put()` must copy data coming from some acquisition device to the Analogy buffer. This ring-buffer is an intermediate area between the device driver and the user-space program, which is supposed to recover the acquired data.

Parameters

in	<i>subd</i>	Subdevice descriptor structure
in	<i>bufdata</i>	The data buffer to copy into the Analogy buffer
in	<i>count</i>	The amount of data to copy

Returns

0 on success, otherwise negative error code.

References A4L\_SUBD\_BUSY\_NR, a4l\_subdevice::buf, and a4l\_subdevice::status.

#### 6.51.2.13 a4l\_get\_chan()

```
int a4l_get_chan (
    struct a4l_subdevice * subd )
```

Get the channel index according to its type.

Parameters

in	<i>subd</i>	Subdevice descriptor structure
----	-------------	--------------------------------

Returns

the channel index.

References A4L\_CHAN\_GLOBAL\_CHANDESC, a4l\_get\_cmd(), a4l\_subdevice::buf, a4l\_subdevice::chan\_desc, a4l\_cmd\_desc::chan\_descs, a4l\_channels\_desc::chans, a4l\_channels\_desc::mode, a4l\_channel::nb\_bits, and a4l\_cmd\_desc::nb\_chan.

#### 6.51.2.14 a4l\_get\_cmd()

```
struct a4l_cmd_desc* a4l_get_cmd (
    struct a4l_subdevice * subd )
```

Get the current Analogy command descriptor.

Parameters

in	<i>subd</i>	Subdevice descriptor structure
----	-------------	--------------------------------


Returns

the command descriptor.

Referenced by a4l\_get\_chan().

## 6.52 Interrupt management services

Collaboration diagram for Interrupt management services:



### Functions

- unsigned int [a4l\\_get\\_irq](#) (struct a4l\_device \*dev)  
*Get the interrupt number in use for a specific device.*
- int [a4l\\_request\\_irq](#) (struct a4l\_device \*dev, unsigned int irq, a4l\_irq\_hdlr\_t handler, unsigned long flags, void \*cookie)  
*Register an interrupt handler for a specific device.*
- int [a4l\\_free\\_irq](#) (struct a4l\_device \*dev, unsigned int irq)  
*Release an interrupt handler for a specific device.*

#### 6.52.1 Detailed Description

#### 6.52.2 Function Documentation

##### 6.52.2.1 a4l\_free\_irq()

```
int a4l_free_irq (
    struct a4l_device * dev,
    unsigned int irq )
```

Release an interrupt handler for a specific device.

Parameters

in	<i>dev</i>	Device descriptor structure
in	<i>irq</i>	Line number of the addressed IRQ

Returns

0 on success, otherwise negative error code.

## 6.52.2.2 a4l\_get\_irq()

```
unsigned int a4l_get_irq (
    struct a4l_device * dev )
```

Get the interrupt number in use for a specific device.

Parameters

in	<i>dev</i>	Device descriptor structure
----	------------	-----------------------------

Returns

the line number used or A4L\_IRQ\_UNUSED if no interrupt is registered.

## 6.52.2.3 a4l\_request\_irq()

```
int a4l_request_irq (
    struct a4l_device * dev,
    unsigned int irq,
    a4l_irq_hdlr_t handler,
    unsigned long flags,
    void * cookie )
```

Register an interrupt handler for a specific device.

Parameters


in	<i>dev</i>	Device descriptor structure
in	<i>irq</i>	Line number of the addressed IRQ
in	<i>handler</i>	Interrupt handler
in	<i>flags</i>	Registration flags: <ul style="list-style-type: none"> <li>• RTDM_IRQTYPE_SHARED: enable IRQ-sharing with other drivers (Warning: real-time drivers and non-real-time drivers cannot share an interrupt line).</li> <li>• RTDM_IRQTYPE_EDGE: mark IRQ as edge-triggered (Warning: this flag is meaningless in RTDM-less context).</li> <li>• A4L_IRQ_DISABLED: keep IRQ disabled when calling the action handler (Warning: this flag is ignored in RTDM-enabled configuration).</li> </ul>
in	<i>cookie</i>	Pointer to be passed to the interrupt handler on invocation

## Returns

0 on success, otherwise negative error code.

## 6.53 Misc services

Collaboration diagram for Misc services:



### Functions

- unsigned long long [a4l\\_get\\_time](#) (void)  
*Get the absolute time in nanoseconds.*

#### 6.53.1 Detailed Description

#### 6.53.2 Function Documentation

##### 6.53.2.1 a4l\_get\_time()

```
unsigned long long a4l_get_time (  
    void )
```

Get the absolute time in nanoseconds.

#### Returns

the absolute time expressed in nanoseconds

References [rtdm\\_clock\\_read\(\)](#).



## 6.54 Clocks and timers

Cobalt/POSIX clock and timer services.

Collaboration diagram for Clocks and timers:



### Functions

- int [clock\\_getres](#) (clockid\_t clock\_id, struct timespec \*tp)
 

*Get the resolution of the specified clock.*
- int [clock\\_gettime](#) (clockid\_t clock\_id, struct timespec \*tp)
 

*Read the specified clock.*
- int [clock\\_settime](#) (clockid\_t clock\_id, const struct timespec \*tp)
 

*Set the specified clock.*
- int [clock\\_nanosleep](#) (clockid\_t clock\_id, int flags, const struct timespec \*rqtp, struct timespec \*rmtp)
 

*Sleep some amount of time.*
- int [nanosleep](#) (const struct timespec \*rqtp, struct timespec \*rmtp)
 

*Sleep some amount of time.*
- int [timer\\_create](#) (clockid\_t clockid, const struct sigevent \*\_\_restrict\_\_ evp, timer\_t \*\_\_restrict\_\_ timerid)
 

*Create a timer*
- int [timer\\_delete](#) (timer\_t timerid)
 

*Delete a timer object.*
- int [timer\\_settime](#) (timer\_t timerid, int flags, const struct itimerspec \*\_\_restrict\_\_ value, struct itimerspec \*\_\_restrict\_\_ ovalue)
 

*Start or stop a timer*
- int [timer\\_gettime](#) (timer\_t timerid, struct itimerspec \*value)
 

*Get timer next expiration date and reload value.*
- int [timer\\_getoverrun](#) (timer\_t timerid)
 

*Get expiration overruns count since the most recent timer expiration signal delivery.*

### 6.54.1 Detailed Description

Cobalt/POSIX clock and timer services.

Cobalt supports three built-in clocks:

`CLOCK_REALTIME` maps to the nucleus system clock, keeping time as the amount of time since the Epoch, with a resolution of one nanosecond.

`CLOCK_MONOTONIC` maps to an architecture-dependent high resolution counter, so is suitable for measuring short time intervals. However, when used for sleeping (with `clock_nanosleep()`), the `CLOCK_MONOTONIC` clock has a resolution of one nanosecond, like the `CLOCK_REALTIME` clock.

`CLOCK_MONOTONIC_RAW` is Linux-specific, and provides monotonic time values from a hardware timer which is not adjusted by NTP. This is strictly equivalent to `CLOCK_MONOTONIC` with Cobalt, which is not NTP adjusted either.

In addition, external clocks can be dynamically registered using the `cobalt_clock_register()` service. These clocks are fully managed by Cobalt extension code, which should advertise each incoming tick by calling `xncklock_tick()` for the relevant clock, from an interrupt context.

Timer objects may be created with the `timer_create()` service using any of the built-in or external clocks. The resolution of these timers is clock-specific. However, built-in clocks all have nanosecond resolution, as specified for `clock_nanosleep()`.

See also

[Specification.](#)

### 6.54.2 Function Documentation

#### 6.54.2.1 `clock_getres()`

```
int clock_getres (
    clockid_t clock_id,
    struct timespec * tp )
```

Get the resolution of the specified clock.

This service returns, at the address *res*, if it is not `NULL`, the resolution of the clock *clock\_id*.

For both `CLOCK_REALTIME` and `CLOCK_MONOTONIC`, this resolution is the duration of one system clock tick. No other clock is supported.

Parameters

<i>clock_id</i>	clock identifier, either <code>CLOCK_REALTIME</code> or <code>CLOCK_MONOTONIC</code> ;
<i>tp</i>	the address where the resolution of the specified clock will be stored on success.

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, <i>clock_id</i> is invalid;</li> </ul>

See also

[Specification.](#)

Tags

[unrestricted](#)

#### 6.54.2.2 clock\_gettime()

```
int clock_gettime (
    clockid_t clock_id,
    struct timespec * tp )
```

Read the specified clock.

This service returns, at the address *tp* the current value of the clock *clock\_id*. If *clock\_id* is:

- CLOCK\_REALTIME, the clock value represents the amount of time since the Epoch, with a precision of one system clock tick;
- CLOCK\_MONOTONIC or CLOCK\_MONOTONIC\_RAW, the clock value is given by an architecture-dependent high resolution counter, with a precision independent from the system clock tick duration.
- CLOCK\_HOST\_REALTIME, the clock value as seen by the host, typically Linux. Resolution and precision depend on the host, but it is guaranteed that both, host and Cobalt, see the same information.

Parameters

<i>clock_id</i>	clock identifier, either CLOCK_REALTIME, CLOCK_MONOTONIC, or CLOCK_HOST_REALTIME;
<i>tp</i>	the address where the value of the specified clock will be stored.

Return values

0	on success;
---	-------------

Return values

-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, <i>clock_id</i> is invalid.</li> </ul>
----	---

See also

[Specification.](#)

Tags

[unrestricted](#)

### 6.54.2.3 clock\_nanosleep()

```
int clock_nanosleep (
    clockid_t clock_id,
    int flags,
    const struct timespec * rntp,
    struct timespec * rntp )
```

Sleep some amount of time.

This service suspends the calling thread until the wakeup time specified by *rntp*, or a signal is delivered to the caller. If the flag `TIMER_ABSTIME` is set in the *flags* argument, the wakeup time is specified as an absolute value of the clock *clock\_id*. If the flag `TIMER_ABSTIME` is not set, the wakeup time is specified as a time interval.

If this service is interrupted by a signal, the flag `TIMER_ABSTIME` is not set, and *rntp* is not `NULL`, the time remaining until the specified wakeup time is returned at the address *rntp*.

The resolution of this service is one system clock tick.

Parameters

<i>clock_id</i>	clock identifier, either <code>CLOCK_REALTIME</code> or <code>CLOCK_MONOTONIC</code> .
<i>flags</i>	one of: <ul style="list-style-type: none"> <li>• 0 meaning that the wakeup time <i>rntp</i> is a time interval;</li> <li>• <code>TIMER_ABSTIME</code>, meaning that the wakeup time is an absolute value of the clock <i>clock_id</i>.</li> </ul>
<i>rntp</i>	address of the wakeup time.
<i>rntp</i>	address where the remaining time before wakeup will be stored if the service is interrupted by a signal.

## Returns

0 on success;

an error number if:

- EPERM, the caller context is invalid;
- ENOTSUP, the specified clock is unsupported;
- EINVAL, the specified wakeup time is invalid;
- EINTR, this service was interrupted by a signal.

## See also

[Specification.](#)

## Tags

[xthread-only](#), [switch-primary](#)

## 6.54.2.4 clock\_settime()

```
int clock_settime (
    clockid_t clock_id,
    const struct timespec * tp )
```

Set the specified clock.

This allow setting the CLOCK\_REALTIME clock.

## Parameters

<i>clock_id</i>	the id of the clock to be set, only CLOCK_REALTIME is supported.
<i>tp</i>	the address of a struct timespec specifying the new date.

## Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, <i>clock_id</i> is not CLOCK_REALTIME;</li> <li>• EINVAL, the date specified by <i>tp</i> is invalid.</li> </ul>

## See also

[Specification.](#)

## Tags

[unrestricted](#)

## 6.54.2.5 nanosleep()

```
int nanosleep (
    const struct timespec * rntp,
    struct timespec * rntp )
```

Sleep some amount of time.

This service suspends the calling thread until the wakeup time specified by *rntp*, or a signal is delivered. The wakeup time is specified as a time interval.

If this service is interrupted by a signal and *rntp* is not *NULL*, the time remaining until the specified wakeup time is returned at the address *rntp*.

The resolution of this service is one system clock tick.

## Parameters

<i>rntp</i>	address of the wakeup time.
<i>rntp</i>	address where the remaining time before wakeup will be stored if the service is interrupted by a signal.

## Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EPERM, the caller context is invalid;</li> <li>• EINVAL, the specified wakeup time is invalid;</li> <li>• EINTR, this service was interrupted by a signal.</li> </ul>

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

## 6.54.2.6 timer\_create()

```
int timer_create (
    clockid_t clockid,
```

```
const struct sigevent *__restrict__ evp,
timer_t *__restrict__ timerid )
```

### Create a timer

This service creates a timer based on the clock *clockid*.

If *evp* is not *NULL*, it describes the notification mechanism used on timer expiration. Only thread-directed notification is supported (*evp->sigev\_notify* set to *SIGEV\_THREAD\_ID*).

If *evp* is *NULL*, the current Cobalt thread will receive the notifications with signal *SIGALRM*.

The recipient thread is delivered notifications when it calls any of the *sigwait()*, *sigtimedwait()* or *sigwait-info()* services.

If this service succeeds, an identifier for the created timer is returned at the address *timerid*. The timer is unarmed until started with the [timer\\_settime\(\)](#) service.

### Parameters

<i>clockid</i>	clock used as a timing base;
<i>evp</i>	description of the asynchronous notification to occur when the timer expires;
<i>timerid</i>	address where the identifier of the created timer will be stored on success.

### Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• <i>EINVAL</i>, the clock <i>clockid</i> is invalid;</li> <li>• <i>EINVAL</i>, the member <i>sigev_notify</i> of the <b>sigevent</b> structure at the address <i>evp</i> is not <i>SIGEV_THREAD_ID</i>;</li> <li>• <i>EINVAL</i>, the member <i>sigev_signo</i> of the <b>sigevent</b> structure is an invalid signal number;</li> <li>• <i>EAGAIN</i>, the maximum number of timers was exceeded, recompile with a larger value.</li> </ul>

See also

[Specification](#).

Tags

[thread-unrestricted](#)

Referenced by [pthread\\_make\\_periodic\\_np\(\)](#).

## 6.54.2.7 timer\_delete()

```
int timer_delete (
    timer_t timerid )
```

Delete a timer object.

This service deletes the timer *timerid*.

Parameters

<i>timerid</i>	identifier of the timer to be removed;
----------------	--

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, <i>timerid</i> is invalid;</li> <li>• EPERM, the timer <i>timerid</i> does not belong to the current process.</li> </ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

## 6.54.2.8 timer\_getoverrun()

```
int timer_getoverrun (
    timer_t timerid )
```

Get expiration overruns count since the most recent timer expiration signal delivery.

This service returns *timerid* expiration overruns count since the most recent timer expiration signal delivery. If this count is more than *DELAYTIMER\_MAX* expirations, *DELAYTIMER\_MAX* is returned.

Parameters

<i>timerid</i>	Timer identifier.
----------------	-------------------



## Returns

the overruns count on success;  
-1 with *errno* set if:

- EINVAL, *timerid* is invalid;
- EPERM, the timer *timerid* does not belong to the current process.

## See also

[Specification.](#)

## Tags

[unrestricted](#)

## 6.54.2.9 timer\_gettime()

```
int timer_gettime (
    timer_t timerid,
    struct itimerspec * value )
```

Get timer next expiration date and reload value.

This service stores, at the address *value*, the expiration date (member *it\_value*) and reload value (member *it\_interval*) of the timer *timerid*. The values are returned as time intervals, and as multiples of the system clock tick duration (see note in section [Clocks and timers services](#) for details on the duration of the system clock tick). If the timer was not started, the returned members *it\_value* and *it\_interval* of *value* are zero.

## Parameters

<i>timerid</i>	timer identifier;
<i>value</i>	address where the timer expiration date and reload value are stored on success.

## Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, <i>timerid</i> is invalid. For <i>timerid</i> to be valid, it must belong to the current process.</li> </ul>

## See also

[Specification.](#)

## Tags

[unrestricted](#)

6.54.2.10 `timer_settime()`

```

timer_settime (
    timer_t timerid,
    int flags,
    const struct itimerspec *__restrict__ value,
    struct itimerspec *__restrict__ ovalue )

```

## Start or stop a timer

This service sets a timer expiration date and reload value of the timer *timerid*. If *ovalue* is not *NULL*, the current expiration date and reload value are stored at the address *ovalue* as with `timer_gettime()`.

If the member *it\_value* of the **itimerspec** structure at *value* is zero, the timer is stopped, otherwise the timer is started. If the member *it\_interval* is not zero, the timer is periodic. The current thread must be a Cobalt thread (created with `pthread_create()`) and will be notified via signal of timer expirations.

When starting the timer, if *flags* is `TIMER_ABSTIME`, the expiration value is interpreted as an absolute date of the clock passed to the `timer_create()` service. Otherwise, the expiration value is interpreted as a time interval.

Expiration date and reload value are rounded to an integer count of nanoseconds.

## Parameters

<i>timerid</i>	identifier of the timer to be started or stopped;
<i>flags</i>	one of 0 or <code>TIMER_ABSTIME</code> ;
<i>value</i>	address where the specified timer expiration date and reload value are read;
<i>ovalue</i>	address where the specified timer previous expiration date and reload value are stored if not <i>NULL</i> .

## Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• <code>EINVAL</code>, the specified timer identifier, expiration date or reload value is invalid. For <i>timerid</i> to be valid, it must belong to the current process.</li> </ul>

See also

[Specification](#).

Tags


[xcontext](#), [switch-primary](#)

Referenced by `pthread_make_periodic_np()`.

## 6.55 Condition variables

Cobalt/POSIX condition variable services.

Collaboration diagram for Condition variables:



### Functions

- int [pthread\\_cond\\_init](#) (pthread\_cond\_t \*cond, const pthread\_condattr\_t \*attr)
  - Initialize a condition variable*
- int [pthread\\_cond\\_destroy](#) (pthread\_cond\_t \*cond)
  - Destroy a condition variable*
- int [pthread\\_cond\\_wait](#) (pthread\_cond\_t \*cond, pthread\_mutex\_t \*mutex)
  - Wait on a condition variable.*
- int [pthread\\_cond\\_timedwait](#) (pthread\_cond\_t \*cond, pthread\_mutex\_t \*mutex, const struct timespec \*abstime)
  - Wait a bounded time on a condition variable.*
- int [pthread\\_cond\\_signal](#) (pthread\_cond\_t \*cond)
  - Signal a condition variable.*
- int [pthread\\_cond\\_broadcast](#) (pthread\_cond\_t \*cond)
  - Broadcast a condition variable.*
- int [pthread\\_condattr\\_init](#) (pthread\_condattr\_t \*attr)
  - Initialize a condition variable attributes object.*
- int [pthread\\_condattr\\_destroy](#) (pthread\_condattr\_t \*attr)
  - Destroy a condition variable attributes object.*
- int [pthread\\_condattr\\_getclock](#) (const pthread\_condattr\_t \*attr, clockid\_t \*clk\_id)
  - Get the clock selection attribute from a condition variable attributes object.*
- int [pthread\\_condattr\\_setclock](#) (pthread\_condattr\_t \*attr, clockid\_t clk\_id)
  - Set the clock selection attribute of a condition variable attributes object.*
- int [pthread\\_condattr\\_getpshared](#) (const pthread\_condattr\_t \*attr, int \*pshared)
  - Get the process-shared attribute from a condition variable attributes object.*
- int [pthread\\_condattr\\_setpshared](#) (pthread\_condattr\_t \*attr, int pshared)
  - Set the process-shared attribute of a condition variable attributes object.*

### 6.55.1 Detailed Description

Cobalt/POSIX condition variable services.

A condition variable is a synchronization object that allows threads to suspend execution until some predicate on shared data is satisfied. The basic operations on conditions are: signal the condition (when the predicate becomes true), and wait for the condition, suspending the thread execution until another thread signals the condition.

A condition variable must always be associated with a mutex, to avoid the race condition where a thread prepares to wait on a condition variable and another thread signals the condition just before the first thread actually waits on it.

Before it can be used, a condition variable has to be initialized with `pthread_cond_init()`. An attribute object, which reference may be passed to this service, allows to select the features of the created condition variable, namely the *clock* used by the `pthread_cond_timedwait()` service (`CLOCK_REALTIME` is used by default), and whether it may be shared between several processes (it may not be shared by default, see `pthread_condattr_setpshared()`).

Note that `pthread_cond_init()` should be used to initialize a condition variable, using the static initializer `PTHREAD_COND_INITIALIZER` will delay the initialization to the first method called on the condition variable and will most likely introduce switches to secondary mode. The documentation (and specifically api-tags) of the condition variable services assumes the condition variable was explicitly initialised with `pthread_cond_init()`.

### 6.55.2 Function Documentation

#### 6.55.2.1 `pthread_cond_broadcast()`

```
int pthread_cond_broadcast (
    pthread_cond_t * cond )
```

Broadcast a condition variable.

This service unblocks all threads blocked on the condition variable *cond*.

Parameters

<i>cond</i>	the condition variable to be signalled.
-------------	---

Returns

0 on succes,

an error number if:

- EINVAL, the condition variable is invalid;
- EPERM, the condition variable is not process-shared and does not belong to the current process.

See also

[Specification.](#)

Tags

[xthread-only](#)

### 6.55.2.2 pthread\_cond\_destroy()

```
int pthread_cond_destroy (
    pthread_cond_t * cond )
```

Destroy a condition variable

This service destroys the condition variable *cond*, if no thread is currently blocked on it. The condition variable becomes invalid for all condition variable services (they all return the EINVAL error) except [pthread\\_cond\\_init\(\)](#).

Parameters

<i>cond</i>	the condition variable to be destroyed.
-------------	---

Returns

0 on success,

an error number if:

- EINVAL, the condition variable *cond* is invalid;
- EPERM, the condition variable is not process-shared and does not belong to the current process;
- EBUSY, some thread is currently using the condition variable.

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

### 6.55.2.3 pthread\_cond\_init()

```
int pthread_cond_init (
    pthread_cond_t * cond,
    const pthread_condattr_t * attr )
```

Initialize a condition variable

This service initializes the condition variable *cond*, using the condition variable attributes object *attr*. If *attr* is *NULL*, default attributes are used (see [pthread\\_condattr\\_init\(\)](#)).

## Parameters

<i>cond</i>	the condition variable to be initialized;
<i>attr</i>	the condition variable attributes object.

## Returns

0 on success,  
an error number if:

- EINVAL, the condition variable attributes object *attr* is invalid or uninitialized;
- EBUSY, the condition variable *cond* was already initialized;
- ENOMEM, insufficient memory available from the system heap to initialize the condition variable, increase CONFIG\_XENO\_OPT\_SYS\_HEAPSZ.
- EAGAIN, no registry slot available, check/raise CONFIG\_XENO\_OPT\_REGISTRY\_NRSL←OTS.

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

## 6.55.2.4 pthread\_cond\_signal()

```
int pthread_cond_signal (
    pthread_cond_t * cond )
```

Signal a condition variable.

This service unblocks one thread blocked on the condition variable *cond*.

If more than one thread is blocked on the specified condition variable, the highest priority thread is unblocked.

## Parameters

<i>cond</i>	the condition variable to be signalled.
-------------	---

## Returns

0 on success,  
an error number if:

- EINVAL, the condition variable is invalid;

- EPERM, the condition variable is not process-shared and does not belong to the current process.

See also

[Specification.](#)

Tags

[xthread-only](#)

### 6.55.2.5 pthread\_cond\_timedwait()

```
int pthread_cond_timedwait (
    pthread_cond_t * cond,
    pthread_mutex_t * mutex,
    const struct timespec * abstime )
```

Wait a bounded time on a condition variable.

This service is equivalent to [pthread\\_cond\\_wait\(\)](#), except that the calling thread remains blocked on the condition variable *cond* only until the timeout specified by *abstime* expires.

The timeout *abstime* is expressed as an absolute value of the *clock* attribute passed to [pthread\\_cond\\_init\(\)](#). By default, *CLOCK\_REALTIME* is used.

Parameters

<i>cond</i>	the condition variable to wait for;
<i>mutex</i>	the mutex associated with <i>cond</i> ;
<i>abstime</i>	the timeout, expressed as an absolute value of the clock attribute passed to <a href="#">pthread_cond_init()</a> .

Returns

0 on success,  
an error number if:

- EPERM, the caller context is invalid;
- EPERM, the specified condition variable is not process-shared and does not belong to the current process;
- EINVAL, the specified condition variable, mutex or timeout is invalid;
- EINVAL, another thread is currently blocked on *cond* using another mutex than *mx*;
- EPERM, the specified mutex is not owned by the caller;
- ETIMEDOUT, the specified timeout expired.

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

### 6.55.2.6 pthread\_cond\_wait()

```
int pthread_cond_wait (
    pthread_cond_t * cond,
    pthread_mutex_t * mutex )
```

Wait on a condition variable.

This service atomically unlocks the mutex *mx*, and block the calling thread until the condition variable *cond* is signalled using [pthread\\_cond\\_signal\(\)](#) or [pthread\\_cond\\_broadcast\(\)](#). When the condition is signaled, this service re-acquire the mutex before returning.

Spurious wakeups occur if a signal is delivered to the blocked thread, so, an application should not assume that the condition changed upon successful return from this service.

Even if the mutex *mx* is recursive and its recursion count is greater than one on entry, it is unlocked before blocking the caller, and the recursion count is restored once the mutex is re-acquired by this service before returning.

Once a thread is blocked on a condition variable, a dynamic binding is formed between the condition variable *cond* and the mutex *mx*; if another thread calls this service specifying *cond* as a condition variable but another mutex than *mx*, this service returns immediately with the EINTR status.

This service is a cancellation point for Cobalt threads (created with the [pthread\\_create\(\)](#) service). When such a thread is cancelled while blocked in a call to this service, the mutex *mx* is re-acquired before the cancellation cleanup handlers are called.

Parameters

<i>cond</i>	the condition variable to wait for;
<i>mutex</i>	the mutex associated with <i>cond</i> .

Returns

0 on success,  
an error number if:

- EPERM, the caller context is invalid;
- EINVAL, the specified condition variable or mutex is invalid;
- EPERM, the specified condition variable is not process-shared and does not belong to the current process;
- EINVAL, another thread is currently blocked on *cond* using another mutex than *mx*;
- EPERM, the specified mutex is not owned by the caller.



See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

#### 6.55.2.7 pthread\_condattr\_destroy()

```
int pthread_condattr_destroy (
    pthread_condattr_t * attr )
```

Destroy a condition variable attributes object.

This service destroys the condition variable attributes object *attr*. The object becomes invalid for all condition variable services (they all return EINVAL) except [pthread\\_condattr\\_init\(\)](#).

Parameters

<i>attr</i>	the initialized mutex attributes object to be destroyed.
-------------	--

Returns

- 0 on success;
- an error number if:
  - EINVAL, the mutex attributes object *attr* is invalid.

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

#### 6.55.2.8 pthread\_condattr\_getclock()

```
int pthread_condattr_getclock (
    const pthread_condattr_t * attr,
    clockid_t * clk_id )
```

Get the clock selection attribute from a condition variable attributes object.

This service stores, at the address *clk\_id*, the value of the *clock* attribute in the condition variable attributes object *attr*.

See [pthread\\_cond\\_timedwait\(\)](#) for a description of the effect of this attribute on a condition variable. The clock ID returned is *CLOCK\_REALTIME* or *CLOCK\_MONOTONIC*.

## Parameters

<i>attr</i>	an initialized condition variable attributes object,
<i>clk</i> ↔ <i>_id</i>	address where the <i>clock</i> attribute value will be stored on success.

## Returns

- 0 on success,  
an error number if:
- EINVAL, the attribute object *attr* is invalid.

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

## 6.55.2.9 pthread\_condattr\_getpshared()

```
int pthread_condattr_getpshared (
    const pthread_condattr_t * attr,
    int * pshared )
```

Get the process-shared attribute from a condition variable attributes object.

This service stores, at the address *pshared*, the value of the *pshared* attribute in the condition variable attributes object *attr*.

The *pshared* attribute may only be one of *PTHREAD\_PROCESS\_PRIVATE* or *PTHREAD\_PROCESS\_SHARED*. See [pthread\\_condattr\\_setpshared\(\)](#) for the meaning of these two constants.

## Parameters

<i>attr</i>	an initialized condition variable attributes object.
<i>pshared</i>	address where the value of the <i>pshared</i> attribute will be stored on success.

## Returns

- 0 on success,  
an error number if:
- EINVAL, the *pshared* address is invalid;
  - EINVAL, the condition variable attributes object *attr* is invalid.

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

#### 6.55.2.10 pthread\_condattr\_init()

```
int pthread_condattr_init (
    pthread_condattr_t * attr )
```

Initialize a condition variable attributes object.

This services initializes the condition variable attributes object *attr* with default values for all attributes. Default value are:

- for the *clock* attribute, *CLOCK\_REALTIME*;
- for the *pshared* attribute *PTHREAD\_PROCESS\_PRIVATE*.

If this service is called specifying a condition variable attributes object that was already initialized, the attributes object is reinitialized.

Parameters

<i>attr</i>	the condition variable attributes object to be initialized.
-------------	---

Returns

0 on success;

an error number if:

- ENOMEM, the condition variable attribute object pointer *attr* is *NULL*.

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

## 6.55.2.11 pthread\_condattr\_setclock()

```
int pthread_condattr_setclock (
    pthread_condattr_t * attr,
    clockid_t clk_id )
```

Set the clock selection attribute of a condition variable attributes object.

This service set the *clock* attribute of the condition variable attributes object *attr*.

See [pthread\\_cond\\_timedwait\(\)](#) for a description of the effect of this attribute on a condition variable.

## Parameters

<i>attr</i>	an initialized condition variable attributes object,
<i>clk</i> ↔ <i>_id</i>	value of the <i>clock</i> attribute, may be <i>CLOCK_REALTIME</i> or <i>CLOCK_MONOTONIC</i> .

## Returns

0 on success,

an error number if:

- EINVAL, the condition variable attributes object *attr* is invalid;
- EINVAL, the value of *clk\_id* is invalid for the *clock* attribute.

## See also

[Specification](#).

## Tags

[thread-unrestricted](#)

## 6.55.2.12 pthread\_condattr\_setpshared()

```
int pthread_condattr_setpshared (
    pthread_condattr_t * attr,
    int pshared )
```

Set the process-shared attribute of a condition variable attributes object.

This service set the *pshared* attribute of the condition variable attributes object *attr*.

## Parameters

<i>attr</i>	an initialized condition variable attributes object.
<i>pshared</i>	value of the <i>pshared</i> attribute, may be one of:
	<ul style="list-style-type: none"> <li>• PTHREAD_PROCESS_PRIVATE, meaning that a condition variable created with the attributes object <i>attr</i> will only be accessible by threads within the same process as the thread that initialized the condition variable;</li> <li>• PTHREAD_PROCESS_SHARED, meaning that a condition variable created with</li> </ul>

## Returns

0 on success,  
an error status if:

- EINVAL, the condition variable attributes object *attr* is invalid;
- EINVAL, the value of *pshared* is invalid.

## See also

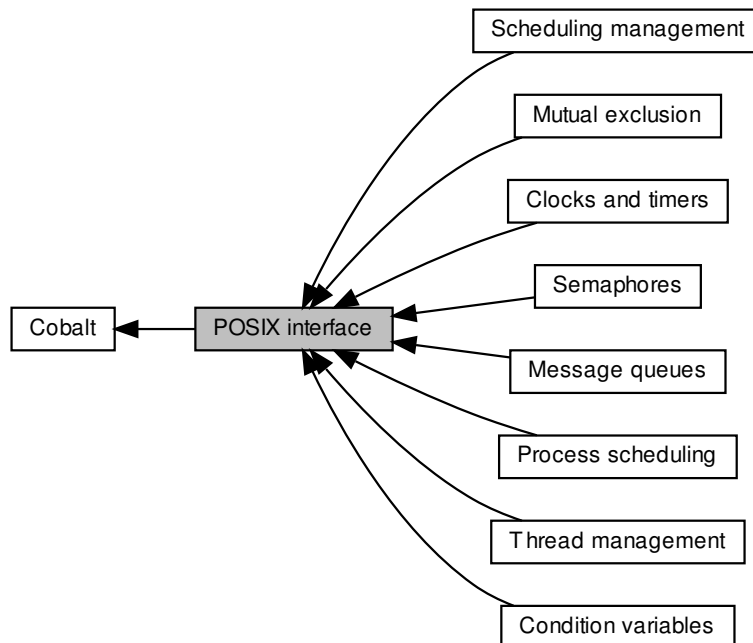
[Specification.](#)

## Tags

[thread-unrestricted](#)

## 6.56 POSIX interface

Collaboration diagram for POSIX interface:



### Modules

- [Clocks and timers](#)  
*Cobalt/POSIX clock and timer services.*
- [Condition variables](#)  
*Cobalt/POSIX condition variable services.*
- [Message queues](#)  
*Cobalt/POSIX message queue services.*
- [Mutual exclusion](#)  
*Cobalt/POSIX mutual exclusion services.*
- [Process scheduling](#)  
*Cobalt/POSIX process scheduling.*
- [Semaphores](#)  
*Cobalt/POSIX semaphore services.*
- [Thread management](#)  
*Cobalt (POSIX) thread management services.*
- [Scheduling management](#)  
*Cobalt scheduling management services.*

### 6.56.1 Detailed Description

The Cobalt/POSIX interface is an implementation of a subset of the [Single Unix specification](#) over the Cobalt core.

The routines from this subset are implemented as wrapper functions as defined by the linker (`-wrap` option, see `man ld(1)`). The linker flags for enabling symbol wrapping can be obtained from the following command: `xeno-config --posix --ldflags`. The full documentation for `xeno-config` can be found at <https://xenomai.org/documentation/xenomai-3/html/man1/xeno-config/index.html>.

When symbol wrapping is enabled:

- calls to POSIX services for which Cobalt provides a (real-time) implementation are redirected to the library implementing the wrapper, by default `libcobalt`. A list of wrapped symbols `libcobalt` overrides can be found in the source tree, in `lib/cobalt/cobalt.wrappers`.

With or without symbol wrapping:

- the wrapper function of a POSIX routine can be explicitly invoked by enclosing the function call with the `__RT()` macro. Since the wrapper symbol is weak, it may be overridden by a 3rd party library, typically to implement its own version of the POSIX call, instead or on top of `libcobalt`'s. e.g. `__RT(sem_init(&sem, 0, 0))` would initialize a real-time semaphore, usually from `libcobalt` unless a stronger `sem_init()` wrapper has been provided by a 3rd party library.
- the `libcobalt` implementation of a POSIX routine can be explicitly invoked by enclosing the function call with the `__COBALT()` macro. e.g. `__COBALT(sem_init(&sem, 0, 0))` would always initialize a Cobalt semaphore (strong symbol).
- the regular `libc` implementation of a POSIX routine can be explicitly invoked by enclosing the function call with the `__STD()` macro. This form basically prevents the symbol wrapping to take place. e.g. `__STD(sem_init(&sem, 0, 0))` would always initialize a regular `*libc` semaphore. This is strictly equivalent to calling the `__real_` form of such routine as documented for `ld(1)`.

Qualifying POSIX calls explicitly as described above may prove useful for invoking real-time services selectively within a large POSIX code base, for which globally enabling symbol wrapping would be unpractical. This may also help in implementing real-time service libraries for which depending on the linker's symbol wrapping mechanism is not suitable.

This said, conforming to the POSIX standard unspoiled by macro tricks for developing an application may be a significant upside as well. YMMV.

## 6.57 Message queues

Cobalt/POSIX message queue services.

Collaboration diagram for Message queues:



### Functions

- `mqd_t mq_open` (const char \*name, int oflags,...)

*Open a message queue*

- int `mq_close` (mqd\_t mqd)

*Close a message queue*

- int `mq_unlink` (const char \*name)

*Unlink a message queue*

- int `mq_getattr` (mqd\_t mqd, struct mq\_attr \*attr)

*Get message queue attributes*

- int `mq_setattr` (mqd\_t mqd, const struct mq\_attr \*\_\_restrict attr, struct mq\_attr \*\_\_restrict oattr)

*Set message queue attributes*

- int `mq_send` (mqd\_t q, const char \*buffer, size\_t len, unsigned prio)

*Send a message to a message queue.*

- int `mq_timedsend` (mqd\_t q, const char \*buffer, size\_t len, unsigned prio, const struct timespec \*timeout)

*Attempt, during a bounded time, to send a message to a message queue.*

- ssize\_t `mq_receive` (mqd\_t q, char \*buffer, size\_t len, unsigned \*prio)

*Receive a message from a message queue.*

- ssize\_t `mq_timedreceive` (mqd\_t q, char \*\_\_restrict buffer, size\_t len, unsigned \*\_\_restrict prio, const struct timespec \*\_\_restrict timeout)

*Attempt, during a bounded time, to receive a message from a message queue.*

- int `mq_notify` (mqd\_t mqd, const struct sigevent \*evp)

*Enable notification on message arrival*



### 6.57.1 Detailed Description

Cobalt/POSIX message queue services.

A message queue allow exchanging data between real-time threads. For a POSIX message queue, maximum message length and maximum number of messages are fixed when it is created with [mq\\_open\(\)](#).

### 6.57.2 Function Documentation

#### 6.57.2.1 mq\_close()

```
int mq_close (
    mqd_t mqd )
```

Close a message queue

This service closes the message queue descriptor *mqd*. The message queue is destroyed only when all open descriptors are closed, and when unlinked with a call to the [mq\\_unlink\(\)](#) service.

Parameters

<i>mqd</i>	message queue descriptor.
------------	---------------------------

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"><li>• EBADF, <i>mqd</i> is an invalid message queue descriptor;</li><li>• EPERM, the caller context is invalid.</li></ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#), [switch-secondary](#)

6.57.2.2 `mq_getattr()`

```
int mq_getattr (
    mqd_t mqd,
    struct mq_attr * attr )
```

Get message queue attributes

This service stores, at the address *attr*, the attributes of the messages queue descriptor *mqd*.

The following attributes are set:

- *mq\_flags*, flags of the message queue descriptor *mqd*;
- *mq\_maxmsg*, maximum number of messages in the message queue;
- *mq\_msgsize*, maximum message size;
- *mq\_curmsgs*, number of messages currently in the queue.

Parameters

<i>mqd</i>	message queue descriptor;
<i>attr</i>	address where the message queue attributes will be stored on success.

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EBADF, <i>mqd</i> is not a valid descriptor.</li> </ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

6.57.2.3 `mq_notify()`

```
int mq_notify (
    mqd_t mqd,
    const struct sigevent * evp )
```

Enable notification on message arrival

If *evp* is not *NULL* and is the address of a **sigevent** structure with the *sigev\_notify* member set to *SIGEV\_SIGNAL*, the current thread will be notified by a signal when a message is sent to the message queue *mqd*, the queue is empty, and no thread is blocked in call to [mq\\_receive\(\)](#) or [mq\\_timedreceive\(\)](#). After the notification, the thread is unregistered.

If *evp* is *NULL* or the *sigev\_notify* member is *SIGEV\_NONE*, the current thread is unregistered.

Only one thread may be registered at a time.

If the current thread is not a Cobalt thread (created with [pthread\\_create\(\)](#)), this service fails.

Parameters

<i>mqd</i>	message queue descriptor;
<i>evp</i>	pointer to an event notification structure.

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• <i>EINVAL</i>, <i>evp</i> is invalid;</li> <li>• <i>EPERM</i>, the caller context is invalid;</li> <li>• <i>EBADF</i>, <i>mqd</i> is not a valid message queue descriptor;</li> <li>• <i>EBUSY</i>, another thread is already registered.</li> </ul>

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

#### 6.57.2.4 mq\_open()

```
mqd_t mq_open (
    const char * name,
    int oflags,
    ... )
```

Open a message queue

This service opens the message queue named *name*.

One of the following values should be set in *oflags*:

- `O_RDONLY`, meaning that the returned queue descriptor may only be used for receiving messages;
- `O_WRONLY`, meaning that the returned queue descriptor may only be used for sending messages;
- `O_RDWR`, meaning that the returned queue descriptor may be used for both sending and receiving messages.

If no message queue named *name* exists, and *oflags* has the `O_CREAT` bit set, the message queue is created by this function, taking two more arguments:

- a *mode* argument, of type `mode_t`, currently ignored;
- an *attr* argument, pointer to an `mq_attr` structure, specifying the attributes of the new message queue.

If *oflags* has the two bits `O_CREAT` and `O_EXCL` set and the message queue already exists, this service fails.

If the `O_NONBLOCK` bit is set in *oflags*, the `mq_send()`, `mq_receive()`, `mq_timedsend()` and `mq_timedreceive()` services return `-1` with *errno* set to `EAGAIN` instead of blocking their caller.

The following arguments of the `mq_attr` structure at the address *attr* are used when creating a message queue:

- *mq\_maxmsg* is the maximum number of messages in the queue (128 by default);
- *mq\_msgsize* is the maximum size of each message (128 by default).

*name* may be any arbitrary string, in which slashes have no particular meaning. However, for portability, using a name which starts with a slash and contains no other slash is recommended.

#### Parameters

<i>name</i>	name of the message queue to open;
<i>oflags</i>	flags.

#### Returns

a message queue descriptor on success;

`-1` with *errno* set if:

- `ENAMETOOLONG`, the length of the *name* argument exceeds 64 characters;
- `EEXIST`, the bits `O_CREAT` and `O_EXCL` were set in *oflags* and the message queue already exists;
- `ENOENT`, the bit `O_CREAT` is not set in *oflags* and the message queue does not exist;
- `ENOSPC`, allocation of system memory failed, or insufficient memory available from the system heap to create the queue, try increasing `CONFIG_XENO_OPT_SYS_HEAPSZ`;
- `EPERM`, attempting to create a message queue from an invalid context;
- `EINVAL`, the *attr* argument is invalid;
- `EMFILE`, too many descriptors are currently open.
- `EAGAIN`, no registry slot available, check/raise `CONFIG_XENO_OPT_REGISTRY_NRSLOTS`.

See also

[Specification.](#)

Tags

[thread-unrestricted](#), [switch-secondary](#)

### 6.57.2.5 mq\_receive()

```
ssize_t mq_receive (
    mqd_t q,
    char * buffer,
    size_t len,
    unsigned * prio )
```

Receive a message from a message queue.

If the message queue *fd* is not empty and if *len* is greater than the *mq\_msgsize* of the message queue, this service copies, at the address *buffer*, the queued message with the highest priority.

If the queue is empty and the flag *O\_NONBLOCK* is not set for the descriptor *fd*, the calling thread is suspended until some message is sent to the queue. If the queue is empty and the flag *O\_NONBLOCK* is set for the descriptor *fd*, this service returns immediately a value of -1 with *errno* set to *EAGAIN*.

Parameters

<i>q</i>	the queue descriptor;
<i>buffer</i>	the address where the received message will be stored on success;
<i>len</i>	<i>buffer</i> length;
<i>prio</i>	address where the priority of the received message will be stored on success.

Returns

the message length, and copy a message at the address *buffer* on success;

-1 with no message unqueued and *errno* set if:

- *EBADF*, *fd* is not a valid descriptor open for reading;
- *EMSGSIZE*, the length *len* is lesser than the message queue *mq\_msgsize* attribute;
- *EAGAIN*, the queue is empty, and the flag *O\_NONBLOCK* is set for the descriptor *fd*;
- *EPERM*, the caller context is invalid;
- *EINTR*, the service was interrupted by a signal.

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

6.57.2.6 `mq_send()`

```
int mq_send (
    mqd_t q,
    const char * buffer,
    size_t len,
    unsigned prio )
```

Send a message to a message queue.

If the message queue *fd* is not full, this service sends the message of length *len* pointed to by the argument *buffer*, with priority *prio*. A message with greater priority is inserted in the queue before a message with lower priority.

If the message queue is full and the flag *O\_NONBLOCK* is not set, the calling thread is suspended until the queue is not full. If the message queue is full and the flag *O\_NONBLOCK* is set, the message is not sent and the service returns immediately a value of -1 with *errno* set to *EAGAIN*.

## Parameters

<i>q</i>	message queue descriptor;
<i>buffer</i>	pointer to the message to be sent;
<i>len</i>	length of the message;
<i>prio</i>	priority of the message.

## Returns

0 and send a message on success;

-1 with no message sent and *errno* set if:

- *EBADF*, *fd* is not a valid message queue descriptor open for writing;
- *EMSGSIZE*, the message length *len* exceeds the *mq\_msgsize* attribute of the message queue;
- *EAGAIN*, the flag *O\_NONBLOCK* is set for the descriptor *fd* and the message queue is full;
- *EPERM*, the caller context is invalid;
- *EINTR*, the service was interrupted by a signal.

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

## 6.57.2.7 mq\_setattr()

```
int mq_setattr (
    mqd_t mqd,
    const struct mq_attr *__restrict__ attr,
    struct mq_attr *__restrict__ oattr )
```

## Set message queue attributes

This service sets the flags of the *mqd* descriptor to the value of the member *mq\_flags* of the **mq\_attr** structure pointed to by *attr*.

The previous value of the message queue attributes are stored at the address *oattr* if it is not *NULL*.

Only setting or clearing the *O\_NONBLOCK* flag has an effect.

## Parameters

<i>mqd</i>	message queue descriptor;
<i>attr</i>	pointer to new attributes (only <i>mq_flags</i> is used);
<i>oattr</i>	if not <i>NULL</i> , address where previous message queue attributes will be stored on success.

## Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>EBADF, <i>mqd</i> is not a valid message queue descriptor.</li> </ul>

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

## 6.57.2.8 mq\_timedreceive()

```
ssize_t mq_timedreceive (
    mqd_t q,
    char *__restrict__ buffer,
    size_t len,
    unsigned *__restrict__ prio,
    const struct timespec *__restrict__ timeout )
```

Attempt, during a bounded time, to receive a message from a message queue.

This service is equivalent to [mq\\_receive\(\)](#), except that if the flag *O\_NONBLOCK* is not set for the descriptor *fd* and the message queue is empty, the calling thread is only suspended until the timeout *abs\_timeout* expires.

## Parameters

<i>q</i>	the queue descriptor;
<i>buffer</i>	the address where the received message will be stored on success;
<i>len</i>	<i>buffer</i> length;
<i>prio</i>	address where the priority of the received message will be stored on success.
<i>timeout</i>	the timeout, expressed as an absolute value of the CLOCK_REALTIME clock.

## Returns

the message length, and copy a message at the address *buffer* on success;

-1 with no message unqueued and *errno* set if:

- EBADF, *fd* is not a valid descriptor open for reading;
- EMSGSIZE, the length *len* is lesser than the message queue *mq\_msgsize* attribute;
- EAGAIN, the queue is empty, and the flag *O\_NONBLOCK* is set for the descriptor *fd*;
- EPERM, the caller context is invalid;
- EINTR, the service was interrupted by a signal;
- ETIMEDOUT, the specified timeout expired.

## See also

[Specification.](#)

## Tags

[xthread-only](#), [switch-primary](#)

## 6.57.2.9 mq\_timedsend()

```
int mq_timedsend (
    mqd_t q,
    const char * buffer,
    size_t len,
    unsigned prio,
    const struct timespec * timeout )
```

Attempt, during a bounded time, to send a message to a message queue.

This service is equivalent to [mq\\_send\(\)](#), except that if the message queue is full and the flag *O\_NONBLOCK* is not set for the descriptor *fd*, the calling thread is only suspended until the timeout specified by *abs\_timeout* expires.

## Parameters

<i>q</i>	message queue descriptor;
<i>buffer</i>	pointer to the message to be sent;
<i>len</i>	length of the message;
<i>prio</i>	priority of the message;
<i>timeout</i>	the timeout, expressed as an absolute value of the CLOCK_REALTIME clock.



## Returns

0 and send a message on success;

-1 with no message sent and *errno* set if:

- EBADF, *fd* is not a valid message queue descriptor open for writing;
- EMSGSIZE, the message length exceeds the *mq\_msgsize* attribute of the message queue;
- EAGAIN, the flag O\_NONBLOCK is set for the descriptor *fd* and the message queue is full;
- EPERM, the caller context is invalid;
- ETIMEDOUT, the specified timeout expired;
- EINTR, the service was interrupted by a signal.

## See also

[Specification.](#)

## Tags

[xthread-only](#), [switch-primary](#)

## 6.57.2.10 mq\_unlink()

```
int mq_unlink (
    const char * name )
```

## Unlink a message queue

This service unlinks the message queue named *name*. The message queue is not destroyed until all queue descriptors obtained with the [mq\\_open\(\)](#) service are closed with the [mq\\_close\(\)](#) service. However, after a call to this service, the unlinked queue may no longer be reached with the [mq\\_open\(\)](#) service.

## Parameters

<i>name</i>	name of the message queue to be unlinked.
-------------	---

## Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EPERM, the caller context is invalid;</li> <li>• ENAMETOOLONG, the length of the <i>name</i> argument exceeds 64 characters;</li> <li>• ENOENT, the message queue does not exist.</li> </ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#), [switch-secondary](#)

## 6.58 Mutual exclusion

Cobalt/POSIX mutual exclusion services.

Collaboration diagram for Mutual exclusion:



### Functions

- int [pthread\\_mutex\\_init](#) (pthread\_mutex\_t \*mutex, const pthread\_mutexattr\_t \*attr)

*Initialize a mutex.*

- static int [\\_\\_attribute\\_\\_](#) ((cold))

*Test if a mutex structure contains a valid autoinitializer.*

- int [pthread\\_mutex\\_destroy](#) (pthread\_mutex\_t \*mutex)

*Destroy a mutex.*

- int [pthread\\_mutex\\_lock](#) (pthread\_mutex\_t \*mutex)

*Lock a mutex.*

- int [pthread\\_mutex\\_timedlock](#) (pthread\_mutex\_t \*mutex, const struct timespec \*to)

*Attempt, during a bounded time, to lock a mutex.*

- int [pthread\\_mutex\\_trylock](#) (pthread\_mutex\_t \*mutex)

*Attempt to lock a mutex.*

- int [pthread\\_mutex\\_unlock](#) (pthread\_mutex\_t \*mutex)

*Unlock a mutex.*

- int [pthread\\_mutex\\_setprioceiling](#) (pthread\_mutex\_t \*\_\_restrict mutex, int prioceiling, int \*\_\_restrict old\_ceiling)

*Set a mutex's priority ceiling.*

- int [pthread\\_mutex\\_getprioceiling](#) (pthread\_mutex\_t \*\_\_restrict mutex, int \*\_\_restrict prioceiling)

*Get a mutex's priority ceiling.*

- int [pthread\\_mutexattr\\_init](#) (pthread\_mutexattr\_t \*attr)

*Initialize a mutex attributes object.*

- int [pthread\\_mutexattr\\_destroy](#) (pthread\_mutexattr\_t \*attr)

*Destroy a mutex attributes object.*

- int [pthread\\_mutexattr\\_gettype](#) (const pthread\_mutexattr\_t \*attr, int \*type)

Get the mutex type attribute from a mutex attributes object.

- int [pthread\\_mutexattr\\_settype](#) (pthread\_mutexattr\_t \*attr, int type)

Set the mutex type attribute of a mutex attributes object.

- int [pthread\\_mutexattr\\_getprotocol](#) (const pthread\_mutexattr\_t \*attr, int \*proto)

Get the protocol attribute from a mutex attributes object.

- int [pthread\\_mutexattr\\_setprotocol](#) (pthread\_mutexattr\_t \*attr, int proto)

Set the protocol attribute of a mutex attributes object.

- int [pthread\\_mutexattr\\_getpshared](#) (const pthread\_mutexattr\_t \*attr, int \*pshared)

Get the process-shared attribute of a mutex attributes object.

- int [pthread\\_mutexattr\\_setpshared](#) (pthread\_mutexattr\_t \*attr, int pshared)

Set the process-shared attribute of a mutex attributes object.

### 6.58.1 Detailed Description

Cobalt/POSIX mutual exclusion services.

A mutex is a MUTual EXclusion device, and is useful for protecting shared data structures from concurrent modifications, and implementing critical sections and monitors.

A mutex has two possible states: unlocked (not owned by any thread), and locked (owned by one thread). A mutex can never be owned by two different threads simultaneously. A thread attempting to lock a mutex that is already locked by another thread is suspended until the owning thread unlocks the mutex first.

Before it can be used, a mutex has to be initialized with [pthread\\_mutex\\_init\(\)](#). An attribute object, which reference may be passed to this service, allows to select the features of the created mutex, namely its *type* (see [pthread\\_mutexattr\\_settype\(\)](#)), the priority *protocol* it uses (see [pthread\\_mutexattr\\_setprotocol\(\)](#)) and whether it may be shared between several processes (see [pthread\\_mutexattr\\_setpshared\(\)](#)).

By default, Cobalt mutexes are of the normal type, use no priority protocol and may not be shared between several processes.

Note that [pthread\\_mutex\\_init\(\)](#) should be used to initialize a mutex, using the static initializer `PTHREAD_MUTEX_INITIALIZER` will delay the initialization to the first method called on the mutex and will most likely introduce switches to secondary mode. The documentation (and specifically api-tags) of the mutex services assumes a mutex was explicitly initialised with [pthread\\_mutex\\_init\(\)](#).

### 6.58.2 Function Documentation

6.58.2.1 `__attribute__()`

```
static int __attribute__ (  
    (cold) ) [static]
```

Test if a mutex structure contains a valid autoinitializer.

Returns

the mutex type on success,  
-1 if not in supported autoinitializer state

6.58.2.2 `pthread_mutex_destroy()`

```
int pthread_mutex_destroy (  
    pthread_mutex_t * mutex )
```

Destroy a mutex.

This service destroys the mutex *mx*, if it is unlocked and not referenced by any condition variable. The mutex becomes invalid for all mutex services (they all return the EINVAL error) except [pthread\\_mutex\\_init\(\)](#).

Parameters

<i>mutex</i>	the mutex to be destroyed.
--------------	----------------------------

Returns

0 on success,  
an error number if:

- EINVAL, the mutex *mx* is invalid;
- EPERM, the mutex is not process-shared and does not belong to the current process;
- EBUSY, the mutex is locked, or used by a condition variable.

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

## 6.58.2.3 pthread\_mutex\_getprioceiling()

```
int pthread_mutex_getprioceiling (
    pthread_mutex_t *__restrict mutex,
    int *__restrict prioceiling )
```

Get a mutex's priority ceiling.

This routine retrieves the priority ceiling value of the specified mutex.

Parameters

<i>mutex</i>	the target mutex.
<i>prioceiling</i>	on success, the current ceiling value is copied to this address.

Returns

- 0 on success;
- an error number if:
  - EINVAL, *mutex* is invalid;
  - EINVAL, *mutex* is not of type PTHREAD\_PRIO\_PROTECT;

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

## 6.58.2.4 pthread\_mutex\_init()

```
int pthread_mutex_init (
    pthread_mutex_t * mutex,
    const pthread_mutexattr_t * attr )
```

Initialize a mutex.

This services initializes the mutex *mx*, using the mutex attributes object *attr*. If *attr* is *NULL*, default attributes are used (see [pthread\\_mutexattr\\_init\(\)](#)).

Parameters

<i>mutex</i>	the mutex to be initialized;
<i>attr</i>	the mutex attributes object.

## Returns

0 on success,  
an error number if:

- EINVAL, the mutex attributes object *attr* is invalid or uninitialized;
- EBUSY, the mutex *mx* was already initialized;
- ENOMEM, insufficient memory available from the system heap to initialize the mutex, increase CONFIG\_XENO\_OPT\_SYS\_HEAPSZ.
- EAGAIN, insufficient memory available to initialize the mutex, increase CONFIG\_XENO\_OPT\_SHARED\_HEAPSZ for a process-shared mutex, or CONFIG\_XENO\_OPT\_PRIVATE\_HEAPSZ for a process-private mutex.
- EAGAIN, no registry slot available, check/raise CONFIG\_XENO\_OPT\_REGISTRY\_NRSLOTS.
- ENOSYS, *attr* mentions priority protection (PTHREAD\_PRIO\_PROTECT), but the C library does not provide pthread\_mutexattr\_get/setprioceiling().

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

## 6.58.2.5 pthread\_mutex\_lock()

```
int pthread_mutex_lock (  
    pthread_mutex_t * mutex )
```

Lock a mutex.

This service attempts to lock the mutex *mx*. If the mutex is free, it becomes locked. If it was locked by another thread than the current one, the current thread is suspended until the mutex is unlocked. If it was already locked by the current mutex, the behaviour of this service depends on the mutex type :

- for mutexes of the *PTHREAD\_MUTEX\_NORMAL* type, this service deadlocks;
- for mutexes of the *PTHREAD\_MUTEX\_ERRORCHECK* type, this service returns the EDEADLK error number;
- for mutexes of the *PTHREAD\_MUTEX\_RECURSIVE* type, this service increments the lock recursion count and returns 0.

## Parameters

<i>mutex</i>	the mutex to be locked.
--------------	-------------------------

## Returns

0 on success  
 an error number if:

- EPERM, the caller is not allowed to perform the operation;
- EINVAL, the mutex *mx* is invalid;
- EPERM, the mutex is not process-shared and does not belong to the current process;
- EDEADLK, the mutex is of the *PTHREAD\_MUTEX\_ERRORCHECK* type and was already locked by the current thread;
- EAGAIN, the mutex is of the *PTHREAD\_MUTEX\_RECURSIVE* type and the maximum number of recursive locks has been exceeded.

See also

[Specification.](#)

## Tags

[xthread-only](#), [switch-primary](#)

## 6.58.2.6 pthread\_mutex\_setprioceiling()

```
int pthread_mutex_setprioceiling (
    pthread_mutex_t *__restrict mutex,
    int prioceiling,
    int *__restrict old_ceiling )
```

Set a mutex's priority ceiling.

This routine acquires the specified mutex, then changes the associated priority ceiling value and releases it. *prioceiling* must be between the values returned by [sched\\_get\\_priority\\_min\(\)](#) and [sched\\_get\\_priority\\_max\(\)](#), inclusive.

The Cobalt implementation applies the priority ceiling protocol using the previous ceiling value during this operation. The new priority ceiling will apply next time the *mutex* transitions from the unlocked to locked state.

## Parameters

<i>mutex</i>	the target mutex.
<i>prioceiling</i>	the new ceiling value.
<i>old_ceiling</i>	on success and if this parameter is non-NULL, the previous ceiling value is copied to this address.

## Returns

0 on success;  
 an error number if:



- EPERM, the caller is not allowed to perform the operation;
- EINVAL, *mutex* is invalid;
- EINVAL, *mutex* is not of type PTHREAD\_PRIO\_PROTECT;
- EINVAL, *prioceiling* is out of range;

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

Note

If the calling thread's priority is higher than the mutex's new priority ceiling, the operation will nevertheless succeed; the Cobalt core never decreases the effective priority of a thread which locks a priority-protected mutex.

#### 6.58.2.7 pthread\_mutex\_timedlock()

```
int pthread_mutex_timedlock (
    pthread_mutex_t * mutex,
    const struct timespec * to )
```

Attempt, during a bounded time, to lock a mutex.

This service is equivalent to [pthread\\_mutex\\_lock\(\)](#), except that if the mutex *mx* is locked by another thread than the current one, this service only suspends the current thread until the timeout specified by *to* expires.

Parameters

<i>mutex</i>	the mutex to be locked;
<i>to</i>	the timeout, expressed as an absolute value of the CLOCK_REALTIME clock.

Returns

0 on success;

an error number if:

- EPERM, the caller is not allowed to perform the operation;
- EINVAL, the mutex *mx* is invalid;
- EPERM, the mutex is not process-shared and does not belong to the current process;
- ETIMEDOUT, the mutex could not be locked and the specified timeout expired;
- EDEADLK, the mutex is of the PTHREAD\_MUTEX\_ERRORCHECK type and the mutex was already locked by the current thread;

- EAGAIN, the mutex is of the `PTHREAD_MUTEX_RECURSIVE` type and the maximum number of recursive locks has been exceeded.

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

### 6.58.2.8 pthread\_mutex\_trylock()

```
int pthread_mutex_trylock (  
    pthread_mutex_t * mutex )
```

Attempt to lock a mutex.

This service is equivalent to [pthread\\_mutex\\_lock\(\)](#), except that if the mutex *mx* is locked by another thread than the current one, this service returns immediately.

Parameters

<i>mutex</i>	the mutex to be locked.
--------------	-------------------------

Returns

0 on success;

an error number if:

- EPERM, the caller is not allowed to perform the operation;
- EINVAL, the mutex is invalid;
- EPERM, the mutex is not process-shared and does not belong to the current process;
- EBUSY, the mutex was locked by another thread than the current one;
- EAGAIN, the mutex is recursive, and the maximum number of recursive locks has been exceeded.

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

## 6.58.2.9 pthread\_mutex\_unlock()

```
int pthread_mutex_unlock (
    pthread_mutex_t * mutex )
```

Unlock a mutex.

This service unlocks the *mutex*. If *mutex* is of the *PTHREAD\_MUTEX\_RECURSIVE* and the locking recursion count is greater than one, the lock recursion count is decremented and the mutex remains locked.

Attempting to unlock a mutex which is not locked or which is locked by another thread than the current one yields the EPERM error, whatever the mutex *type* attribute.

Parameters

<i>mutex</i>	the mutex to be released.
--------------	---------------------------

Returns

0 on success;

an error number if:

- EPERM, the caller is not allowed to perform the operation;
- EINVAL, the mutex *mutex* is invalid;
- EPERM, the mutex was not locked by the current thread.

See also

[Specification](#).

Tags

[xthread-only](#), [switch-primary](#)

## 6.58.2.10 pthread\_mutexattr\_destroy()

```
int pthread_mutexattr_destroy (
    pthread_mutexattr_t * attr )
```

Destroy a mutex attributes object.

This service destroys the mutex attributes object *attr*. The object becomes invalid for all mutex services (they all return EINVAL) except [pthread\\_mutexattr\\_init\(\)](#).

Parameters

<i>attr</i>	the initialized mutex attributes object to be destroyed.
-------------	--

## Returns

- 0 on success;  
an error number if:
- EINVAL, the mutex attributes object *attr* is invalid.

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

## 6.58.2.11 pthread\_mutexattr\_getprotocol()

```
int pthread_mutexattr_getprotocol (
    const pthread_mutexattr_t * attr,
    int * proto )
```

Get the protocol attribute from a mutex attributes object.

This service stores, at the address *proto*, the value of the *protocol* attribute in the mutex attributes object *attr*.

The *protocol* attribute may be one of *PTHREAD\_PRIO\_NONE*, *PTHREAD\_PRIO\_INHERIT* or *PTHREAD\_PRIO\_PROTECT*. See [pthread\\_mutexattr\\_setprotocol\(\)](#) for the meaning of these constants.

## Parameters

<i>attr</i>	an initialized mutex attributes object;
<i>proto</i>	address where the value of the <i>protocol</i> attribute will be stored on success.

## Returns

- 0 on success,  
an error number if:
- EINVAL, the *proto* address is invalid;
  - EINVAL, the mutex attributes object *attr* is invalid.

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

## 6.58.2.12 pthread\_mutexattr\_getpshared()

```
int pthread_mutexattr_getpshared (
    const pthread_mutexattr_t * attr,
    int * pshared )
```

Get the process-shared attribute of a mutex attributes object.

This service stores, at the address *pshared*, the value of the *pshared* attribute in the mutex attributes object *attr*.

The *pshared* attribute may only be one of *PTHREAD\_PROCESS\_PRIVATE* or *PTHREAD\_PROCESS\_SHARED*. See [pthread\\_mutexattr\\_setpshared\(\)](#) for the meaning of these two constants.

Parameters

<i>attr</i>	an initialized mutex attributes object;
<i>pshared</i>	address where the value of the <i>pshared</i> attribute will be stored on success.

Returns

- 0 on success;
- an error number if:
  - EINVAL, the *pshared* address is invalid;
  - EINVAL, the mutex attributes object *attr* is invalid.

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

## 6.58.2.13 pthread\_mutexattr\_gettype()

```
int pthread_mutexattr_gettype (
    const pthread_mutexattr_t * attr,
    int * type )
```

Get the mutex type attribute from a mutex attributes object.

This service stores, at the address *type*, the value of the *type* attribute in the mutex attributes object *attr*.

See [pthread\\_mutex\\_lock\(\)](#) and [pthread\\_mutex\\_unlock\(\)](#) for a description of the values of the *type* attribute and their effect on a mutex.

## Parameters

<i>attr</i>	an initialized mutex attributes object,
<i>type</i>	address where the <i>type</i> attribute value will be stored on success.

## Returns

- 0 on success,  
 an error number if:
- EINVAL, the *type* address is invalid;
  - EINVAL, the mutex attributes object *attr* is invalid.

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

## 6.58.2.14 pthread\_mutexattr\_init()

```
int pthread_mutexattr_init (
    pthread_mutexattr_t * attr )
```

Initialize a mutex attributes object.

This service initializes the mutex attributes object *attr* with default values for all attributes. Default values are :

- for the *type* attribute, *PTHREAD\_MUTEX\_NORMAL*;
- for the *protocol* attribute, *PTHREAD\_PRIO\_NONE*;
- for the *pshared* attribute, *PTHREAD\_PROCESS\_PRIVATE*.

If this service is called specifying a mutex attributes object that was already initialized, the attributes object is reinitialized.

## Parameters

<i>attr</i>	the mutex attributes object to be initialized.
-------------	--

## Returns

- 0 on success;  
 an error number if:
- ENOMEM, the mutex attributes object pointer *attr* is *NULL*.

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

## 6.58.2.15 pthread\_mutexattr\_setprotocol()

```
int pthread_mutexattr_setprotocol (
    pthread_mutexattr_t * attr,
    int proto )
```

Set the protocol attribute of a mutex attributes object.

This service set the *type* attribute of the mutex attributes object *attr*.

## Parameters

<i>attr</i>	an initialized mutex attributes object,
<i>proto</i>	value of the <i>protocol</i> attribute, may be one of: <ul style="list-style-type: none"> <li>• PTHREAD_PRIO_NONE, meaning that a mutex created with the attributes object <i>attr</i> will not follow any priority protocol;</li> <li>• PTHREAD_PRIO_INHERIT, meaning that a mutex created with the attributes object <i>attr</i>, will follow the priority inheritance protocol.</li> <li>• PTHREAD_PRIO_PROTECT, meaning that a mutex created with the attributes object <i>attr</i>, will follow the priority protect protocol.</li> </ul>

## Returns

- 0 on success,  
 an error number if:
- EINVAL, the mutex attributes object *attr* is invalid;
  - ENOTSUP, the value of *proto* is unsupported;
  - EINVAL, the value of *proto* is invalid.

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

#### 6.58.2.16 pthread\_mutexattr\_setpshared()

```
int pthread_mutexattr_setpshared (
    pthread_mutexattr_t * attr,
    int pshared )
```

Set the process-shared attribute of a mutex attributes object.

This service set the *pshared* attribute of the mutex attributes object *attr*.

Parameters

<i>attr</i>	an initialized mutex attributes object.
<i>pshared</i>	value of the <i>pshared</i> attribute, may be one of: <ul style="list-style-type: none"> <li>• PTHREAD_PROCESS_PRIVATE, meaning that a mutex created with the attributes object <i>attr</i> will only be accessible by threads within the same process as the thread that initialized the mutex;</li> <li>• PTHREAD_PROCESS_SHARED, meaning that a mutex created with the attributes object <i>attr</i> will be accessible by any thread that has access to the memory where the mutex is allocated.</li> </ul>

Returns

0 on success,  
an error status if:

- EINVAL, the mutex attributes object *attr* is invalid;
- EINVAL, the value of *pshared* is invalid.

See also

[Specification.](#)

Tags

[thread-unrestricted](#)



## 6.58.2.17 pthread\_mutexattr\_settype()

```
int pthread_mutexattr_settype (
    pthread_mutexattr_t * attr,
    int type )
```

Set the mutex type attribute of a mutex attributes object.

This service set the *type* attribute of the mutex attributes object *attr*.

See [pthread\\_mutex\\_lock\(\)](#) and [pthread\\_mutex\\_unlock\(\)](#) for a description of the values of the *type* attribute and their effect on a mutex.

The `PTHREAD_MUTEX_DEFAULT` default *type* is the same as `PTHREAD_MUTEX_NORMAL`. Note that using a recursive Cobalt mutex with a Cobalt condition variable is safe (see [pthread\\_cond\\_wait\(\)](#) documentation).

## Parameters

<i>attr</i>	an initialized mutex attributes object,
<i>type</i>	value of the <i>type</i> attribute.

## Returns

0 on success,

an error number if:

- EINVAL, the mutex attributes object *attr* is invalid;
- EINVAL, the value of *type* is invalid for the *type* attribute.

## See also

[Specification.](#)


## Tags

[thread-unrestricted](#)

## 6.59 Process scheduling

Cobalt/POSIX process scheduling.

Collaboration diagram for Process scheduling:



### Functions

- int [sched\\_yield](#) (void)
  - Yield the processor.*
- int [sched\\_get\\_priority\\_min](#) (int policy)
  - Get minimum priority of the specified scheduling policy.*
- int [sched\\_get\\_priority\\_min\\_ex](#) (int policy)
  - Get extended minimum priority of the specified scheduling policy.*
- int [sched\\_get\\_priority\\_max](#) (int policy)
  - Get maximum priority of the specified scheduling policy.*
- int [sched\\_setscheduler](#) (pid\_t pid, int policy, const struct sched\_param \*param)
  - Set the scheduling policy and parameters of the specified process.*
- int [sched\\_setscheduler\\_ex](#) (pid\_t pid, int policy, const struct sched\_param\_ex \*param\_ex)
  - Set extended scheduling policy of a process*
- int [sched\\_getscheduler](#) (pid\_t pid)
  - Get the scheduling policy of the specified process.*
- int [sched\\_getscheduler\\_ex](#) (pid\_t pid, int \*policy\_r, struct sched\_param\_ex \*param\_ex)
  - Get extended scheduling policy of a process*
- int [sched\\_get\\_priority\\_max\\_ex](#) (int policy)
  - Get extended maximum priority of the specified scheduling policy.*
- int [sched\\_setconfig\\_np](#) (int cpu, int policy, const union sched\_config \*config, size\_t len)
  - Set CPU-specific scheduler settings for a policy*
- ssize\_t [sched\\_getconfig\\_np](#) (int cpu, int policy, union sched\_config \*config, size\_t \*len\_r)
  - Retrieve CPU-specific scheduler settings for a policy*

### 6.59.1 Detailed Description

Cobalt/POSIX process scheduling.

See also

[Specification.](#)

### 6.59.2 Function Documentation

#### 6.59.2.1 sched\_get\_priority\_max()

```
int sched_get_priority_max (
    int policy )
```

Get maximum priority of the specified scheduling policy.

This service returns the maximum priority of the scheduling policy *policy*.

Parameters

<i>policy</i>	scheduling policy.
---------------	--------------------

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"><li>• EINVAL, <i>policy</i> is invalid.</li></ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#), [switch-secondary](#)

Note

Fetching the maximum priority level of SCHED\_FIFO, SCHED\_RR or any Xenomai-specific policy never leads to a mode switch. Any other value of *policy* may switch the caller to secondary mode.

## 6.59.2.2 sched\_get\_priority\_max\_ex()

```
int sched_get_priority_max_ex (
    int policy )
```

Get extended maximum priority of the specified scheduling policy.

This service returns the maximum priority of the scheduling policy *policy*, reflecting any Cobalt extension to standard classes.

Parameters

<i>policy</i>	scheduling policy.
---------------	--------------------

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, <i>policy</i> is invalid.</li> </ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

## 6.59.2.3 sched\_get\_priority\_min()

```
int sched_get_priority_min (
    int policy )
```

Get minimum priority of the specified scheduling policy.

This service returns the minimum priority of the scheduling policy *policy*.

Parameters

<i>policy</i>	scheduling policy.
---------------	--------------------

Return values

0	on success;
---	-------------

Return values

-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, <i>policy</i> is invalid.</li> </ul>
----	---

See also

[Specification.](#)

Tags

[thread-unrestricted](#), [switch-secondary](#)

Note

Fetching the minimum priority level of SCHED\_FIFO, SCHED\_RR or any Xenomai-specific policy never leads to a mode switch. Any other value of *policy* may switch the caller to secondary mode.

#### 6.59.2.4 sched\_get\_priority\_min\_ex()

```
int sched_get_priority_min_ex (
    int policy )
```

Get extended minimum priority of the specified scheduling policy.

This service returns the minimum priority of the scheduling policy *policy*, reflecting any Cobalt extension to the standard classes.

Parameters

<i>policy</i>	scheduling policy.
---------------	--------------------

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, <i>policy</i> is invalid.</li> </ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

## 6.59.2.5 sched\_getconfig\_np()

```
ssize_t sched_getconfig_np (
    int cpu,
    int policy,
    union sched_config * config,
    size_t * len_r )
```

Retrieve CPU-specific scheduler settings for a policy

A configuration is strictly local to the target *cpu*, and may differ from other processors.

Parameters

<i>cpu</i>	processor to retrieve the configuration of.
<i>policy</i>	scheduling policy to which the configuration data applies. Currently, only SCHED_TP and SCHED_QUOTA are valid input.
<i>config</i>	a pointer to a memory area which receives the configuration settings upon success of this call.

SCHED\_TP specifics

On successful return, *config->quota.tp* contains the TP schedule active on *cpu*.

SCHED\_QUOTA specifics

On entry, *config->quota.get.tgid* must contain the thread group identifier to inquire about.

On successful exit, *config->quota.info* contains the information related to the thread group referenced to by *config->quota.get.tgid*.

Parameters

<i>in, out</i>	<i>len↔ _r</i>	a pointer to a variable for collecting the overall length of the configuration data returned (in bytes). This variable must contain the amount of space available in <i>config</i> when the request is issued.
----------------	--------------------	--

Returns

the number of bytes copied to *config* on success;  
a negative error number if:

- EINVAL, *cpu* is invalid, or *policy* is unsupported by the current kernel configuration, or *len* cannot hold the retrieved configuration data.

- ESRCH, with *policy* equal to SCHED\_QUOTA, if the group identifier required to perform the operation is not valid (i.e. `config->quota.get.tgid` is invalid).
- ENOMEM, lack of memory to perform the operation.
- ENOSPC, *len* is too short.

Tags

[thread-unrestricted](#), [switch-primary](#)

#### 6.59.2.6 sched\_getscheduler()

```
int sched_getscheduler (  
    pid_t pid )
```

Get the scheduling policy of the specified process.

This service retrieves the scheduling policy of the Cobalt process identified by *pid*.

If *pid* does not identify an existing Cobalt thread/process, this service falls back to the regular [sched\\_getscheduler\(\)](#) service.

Parameters

<i>pid</i>	target process/thread;
------------	------------------------

Returns

0 on success;

an error number if:

- ESRCH, *pid* is not found;
- EINVAL, *pid* is negative
- EFAULT, *param\_ex* is an invalid address;

See also

[Specification](#).

Tags

[thread-unrestricted](#)

## 6.59.2.7 sched\_getscheduler\_ex()

```
int sched_getscheduler_ex (
    pid_t pid,
    int * policy_r,
    struct sched_param_ex * param_ex )
```

Get extended scheduling policy of a process

This service is an extended version of the [sched\\_getscheduler\(\)](#) service, which supports Cobalt-specific and/or additional scheduling policies, not available with the host Linux environment. It retrieves the scheduling policy of the Cobalt process/thread identified by *pid*, and the associated scheduling parameters (e.g. the priority).

Parameters

<i>pid</i>	queried process/thread. If zero, the current thread is assumed.
<i>policy_r</i>	a pointer to a variable receiving the current scheduling policy of <i>pid</i> .
<i>param_ex</i>	a pointer to a structure receiving the current scheduling parameters of <i>pid</i> .

Returns

- 0 on success;
- an error number if:
  - ESRCH, *pid* is not a Cobalt thread;
  - EINVAL, *pid* is negative or *param\_ex* is NULL;
  - EFAULT, *param\_ex* is an invalid address;

Tags

[thread-unrestricted](#)

## 6.59.2.8 sched\_setconfig\_np()

```
int sched_setconfig_np (
    int cpu,
    int policy,
    const union sched_config * config,
    size_t len )
```

Set CPU-specific scheduler settings for a policy

A configuration is strictly local to the target *cpu*, and may differ from other processors.

Parameters

<i>cpu</i>	processor to load the configuration of.
<i>policy</i>	scheduling policy to which the configuration data applies. Currently, SCHED_TP and SCHED_QUOTA are valid.
<i>config</i>	a pointer to the configuration data to load on <i>cpu</i> , applicable to <i>policy</i> .



Settings applicable to SCHED\_TP

This call controls the temporal partitions for *cpu*, depending on the operation requested.

- `config.tp.op` specifies the operation to perform:
- `sched_tp_install` installs a new TP schedule on *cpu*, defined by `config.tp.windows[]`. The global time frame is not activated upon return from this request yet; `sched_tp_start` must be issued to activate the temporal scheduling on *CPU*.
- `sched_tp_uninstall` removes the current TP schedule from *cpu*, releasing all the attached resources. If no TP schedule exists on *CPU*, this request has no effect.
- `sched_tp_start` enables the temporal scheduling on *cpu*, starting the global time frame. If no TP schedule exists on *cpu*, this action has no effect.
- `sched_tp_stop` disables the temporal scheduling on *cpu*. The current TP schedule is not uninstalled though, and may be re-started later by a `sched_tp_start` request.

#### Attention

As a consequence of this request, threads assigned to the un-scheduled partitions may be starved from CPU time.

- for a `sched_tp_install` operation, `config.tp.nr_windows` indicates the number of elements present in the `config.tp.windows[]` array. If `config.tp.nr_windows` is zero, the action taken is identical to `sched_tp_uninstall`.
- if `config.tp.nr_windows` is non-zero, `config.tp.windows[]` is a set scheduling time slots for threads assigned to *cpu*. Each window is specified by its offset from the start of the global time frame (`windows[].offset`), its duration (`windows[].duration`), and the partition id it should activate during such period of time (`windows[].ptid`). This field is not considered for other requests than `sched_tp_install`.

Time slots must be strictly contiguous, i.e. `windows[n].offset + windows[n].duration` shall equal `windows[n + 1].offset`. If `windows[].ptid` is in the range `[0..CONFIG_XENO_OPT_SCHED_TP_NRP-ART-1]`, SCHED\_TP threads which belong to the partition being referred to may be given CPU time on *cpu*, from time `windows[].offset` to `windows[].offset + windows[].duration`, provided those threads are in a runnable state.

Time holes between valid time slots may be defined using windows activating the pseudo partition -1. When such window is active in the global time frame, no CPU time is available to SCHED\_TP threads on *cpu*.

#### Note

The `sched_tp_confsz(nr_windows)` macro returns the length of `config.tp` depending on the number of time slots to be defined in `config.tp.windows[]`, as specified by `config.tp.nr_windows`.

Settings applicable to SCHED\_QUOTA

This call manages thread groups running on *cpu*, defining per-group quota for limiting their CPU consumption.

- `config.quota.op` should define the operation to be carried out. Valid operations are:
  - `sched_quota_add` for creating a new thread group on `cpu`. The new group identifier will be written back to `info.tgid` upon success. A new group is given no initial runtime budget when created. `sched_quota_set` should be issued to enable it.
  - `sched_quota_remove` for deleting a thread group on `cpu`. The group identifier should be passed in `config.quota.remove.tgid`.
  - `sched_quota_set` for updating the scheduling parameters of a thread group defined on `cpu`. The group identifier should be passed in `config.quota.set.tgid`, along with the allotted percentage of the quota interval (`config.quota.set.quota`), and the peak percentage allowed (`config.quota.set.quota_peak`).

All three operations fill in the `config.info` structure with the information reflecting the state of the scheduler on `cpu` with respect to `policy`, after the requested changes have been applied.

Parameters

<code>len</code>	overall length of the configuration data (in bytes).
------------------	--

Returns

0 on success;  
an error number if:

- `EINVAL`, `cpu` is invalid, or `policy` is unsupported by the current kernel configuration, `len` is invalid, or `config` contains invalid parameters.
- `ENOMEM`, lack of memory to perform the operation.
- `EBUSY`, with `policy` equal to `SCHED_QUOTA`, if an attempt is made to remove a thread group which still manages threads.
- `ESRCH`, with `policy` equal to `SCHED_QUOTA`, if the group identifier required to perform the operation is not valid.

Tags

[thread-unrestricted](#), [switch-primary](#)

#### 6.59.2.9 `sched_setscheduler()`

```
int sched_setscheduler (
    pid_t pid,
    int policy,
    const struct sched_param * param )
```

Set the scheduling policy and parameters of the specified process.

This service set the scheduling policy of the Cobalt process identified by `pid` to the value `policy`, and its scheduling parameters (i.e. its priority) to the value pointed to by `param`.

If the current Linux thread ID is passed (see `gettid(2)`), this service turns the current regular POSIX thread into a Cobalt thread. If `pid` is neither the identifier of the current thread nor the identifier of an existing Cobalt thread, this service falls back to the regular `sched_setscheduler()` service.

## Parameters

<i>pid</i>	target process/thread;
<i>policy</i>	scheduling policy, one of SCHED_FIFO, SCHED_RR, or SCHED_OTHER;
<i>param</i>	scheduling parameters address.

## Returns

0 on success;

an error number if:

- ESRCH, *pid* is invalid;
- EINVAL, *policy* or *param*->*sched\_priority* is invalid;
- EAGAIN, insufficient memory available from the system heap, increase CONFIG\_XENO\_O↔PT\_SYS\_HEAPSZ;
- EFAULT, *param* is an invalid address;

## See also

[Specification.](#)

## Note

See [sched\\_setscheduler\\_ex\(\)](#).

## Tags

[thread-unrestricted](#), [switch-secondary](#), [switch-primary](#)

Referenced by [sched\\_setscheduler\\_ex\(\)](#).

## 6.59.2.10 sched\_setscheduler\_ex()

```
int sched_setscheduler_ex (
    pid_t pid,
    int policy,
    const struct sched_param_ex * param_ex )
```

## Set extended scheduling policy of a process

This service is an extended version of the [sched\\_setscheduler\(\)](#) service, which supports Cobalt-specific and/or additional scheduling policies, not available with the host Linux environment. It sets the scheduling policy of the Cobalt process/thread identified by *pid* to the value *policy*, and the scheduling parameters (e.g. its priority) to the value pointed to by *par*.

If the current Linux thread ID or zero is passed (see [gettid\(2\)](#)), this service may turn the current regular POSIX thread into a Cobalt thread.

## Parameters

<i>pid</i>	target process/thread. If zero, the current thread is assumed.
<i>policy</i>	scheduling policy, one of SCHED_WEAK, SCHED_FIFO, SCHED_COBALT, SCHED_RR, SCHED_SPORADIC, SCHED_TP, SCHED_QUOTA or SCHED_NORMAL;
<i>param_ex</i>	address of scheduling parameters. As a special exception, a negative sched_priority value is interpreted as if SCHED_WEAK was given in <i>policy</i> , using the absolute value of this parameter as the weak priority level.

When CONFIG\_XENO\_OPT\_SCHED\_WEAK is enabled, SCHED\_WEAK exhibits priority levels in the [0..99] range (inclusive). Otherwise, sched\_priority must be zero for the SCHED\_WEAK policy.

## Returns

0 on success;

an error number if:

- ESRCH, *pid* is not found;
- EINVAL, *pid* is negative, *param\_ex* is NULL, any of *policy* or *param\_ex->sched\_priority* is invalid;
- EAGAIN, insufficient memory available from the system heap, increase CONFIG\_XENO\_OPT\_SYS\_HEAPSZ;
- EFAULT, *param\_ex* is an invalid address;

## Note

See [sched\\_setscheduler\(\)](#).

## Tags

[thread-unrestricted](#), [switch-secondary](#), [switch-primary](#)

References [sched\\_setscheduler\(\)](#).

## 6.59.2.11 sched\_yield()

```
int sched_yield (
    void )
```

Yield the processor.

This function move the current thread at the end of its priority group.

Return values

0	
---	--

See also

[Specification.](#)

Tags


[thread-unrestricted](#), [switch-primary](#)

References XNRELAX, and XNWEAK.

## 6.60 Semaphores

Cobalt/POSIX semaphore services.

Collaboration diagram for Semaphores:



### Functions

- int [sem\\_init](#) (sem\_t \*sem, int pshared, unsigned int value)
  - Initialize an unnamed semaphore.*
- int [sem\\_destroy](#) (sem\_t \*sem)
  - Destroy an unnamed semaphore*
- int [sem\\_post](#) (sem\_t \*sem)
  - Post a semaphore*
- int [sem\\_trywait](#) (sem\_t \*sem)
  - Attempt to decrement a semaphore*
- int [sem\\_wait](#) (sem\_t \*sem)
  - Decrement a semaphore*
- int [sem\\_timedwait](#) (sem\_t \*sem, const struct timespec \*abs\_timeout)
  - Attempt to decrement a semaphore with a time limit*
- int [sem\\_close](#) (sem\_t \*sem)
  - Close a named semaphore*
- int [sem\\_unlink](#) (const char \*name)
  - Unlink a named semaphore*

### 6.60.1 Detailed Description

Cobalt/POSIX semaphore services.

Semaphores are counters for resources shared between threads. The basic operations on semaphores are: increment the counter atomically, and wait until the counter is non-null and decrement it atomically.

Semaphores have a maximum value past which they cannot be incremented. The macro `SEM_VALUE_MAX` is defined to be this maximum value.

## 6.60.2 Function Documentation

### 6.60.2.1 `sem_close()`

```
int sem_close (
    sem_t * sem )
```

Close a named semaphore

This service closes the semaphore *sem*. The semaphore is destroyed only when unlinked with a call to the [sem\\_unlink\(\)](#) service and when each call to `sem_open()` matches a call to this service.

When a semaphore is destroyed, the memory it used is returned to the system heap, so that further references to this semaphore are not guaranteed to fail, as is the case for unnamed semaphores.

This service fails if *sem* is an unnamed semaphore.

Parameters

<i>sem</i>	the semaphore to be closed.
------------	-----------------------------

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, the semaphore <i>sem</i> is invalid or is an unnamed semaphore.</li> </ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#), [switch-secondary](#)

### 6.60.2.2 `sem_destroy()`

```
int sem_destroy (
    sem_t * sem )
```

Destroy an unnamed semaphore

This service destroys the semaphore *sem*. Threads currently blocked on *sem* are unblocked and the service they called return -1 with *errno* set to EINVAL. The semaphore is then considered invalid by all semaphore services (they all fail with *errno* set to EINVAL) except [sem\\_init\(\)](#).

This service fails if *sem* is a named semaphore.

## Parameters

<i>sem</i>	the semaphore to be destroyed.
------------	--------------------------------

## Return values

<i>always</i>	0 on success. If SEM_WARNDEL was mentioned in <code>sem_init_np()</code> , the semaphore is deleted as requested and a strictly positive value is returned to warn the caller if threads were pending on it, otherwise zero is returned. If SEM_NOBUSYDEL was mentioned in <code>sem_init_np()</code> , <code>sem_destroy()</code> may succeed only if no thread is waiting on the semaphore to delete, otherwise -EBUSY is returned.
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, the semaphore <i>sem</i> is invalid or a named semaphore;</li> <li>• EPERM, the semaphore <i>sem</i> is not process-shared and does not belong to the current process.</li> <li>• EBUSY, a thread is currently waiting on the semaphore <i>sem</i> with SEM_NOBUSYDEL set.</li> </ul>

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

6.60.2.3 `sem_init()`

```
int sem_init (
    sem_t * sem,
    int pshared,
    unsigned int value )
```

Initialize an unnamed semaphore.

This service initializes the semaphore *sm*, with the value *value*.

This service fails if *sm* is already initialized or is a named semaphore.

## Parameters

<i>sem</i>	the semaphore to be initialized;
<i>pshared</i>	if zero, means that the new semaphore may only be used by threads in the same process as the thread calling <code>sem_init()</code> ; if non zero, means that the new semaphore may be used by any thread that has access to the memory where the semaphore is allocated.
<i>value</i>	the semaphore initial value.



## Return values

<i>0</i>	on success,
<i>-1</i>	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EBUSY, the semaphore <i>sm</i> was already initialized;</li> <li>• EAGAIN, insufficient memory available to initialize the semaphore, increase CONFIG_XENO_OPT_SHARED_HEAPSZ for a process-shared semaphore, or CONFIG_XENO_OPT_PRIVATE_HEAPSZ for a process-private semaphore.</li> <li>• EAGAIN, no registry slot available, check/raise CONFIG_XENO_OPT_REGISTRY_NRSLOTS.</li> <li>• EINVAL, the <i>value</i> argument exceeds SEM_VALUE_MAX.</li> </ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#)

## 6.60.2.4 sem\_post()

```
int sem_post (
    sem_t * sem )
```

Post a semaphore

This service posts the semaphore *sem*.

If no thread is currently blocked on this semaphore, its count is incremented unless "pulse" mode is enabled for it (see sem\_init\_np(), SEM\_PULSE). If a thread is blocked on the semaphore, the thread heading the wait queue is unblocked.

Parameters

<i>sem</i>	the semaphore to be signaled.
------------	-------------------------------

Return values

<i>0</i>	on success;
----------	-------------

## Return values

-1	with errno set if: <ul style="list-style-type: none"> <li>• EINVAL, the specified semaphore is invalid or uninitialized;</li> <li>• EPERM, the semaphore <i>sm</i> is not process-shared and does not belong to the current process;</li> <li>• EAGAIN, the semaphore count is <i>SEM_VALUE_MAX</i>.</li> </ul>
----	---

See also

[Specification.](#)

Tags

[unrestricted](#)

## 6.60.2.5 sem\_timedwait()

```
int sem_timedwait (
    sem_t * sem,
    const struct timespec * abs_timeout )
```

Attempt to decrement a semaphore with a time limit

This service is equivalent to [sem\\_wait\(\)](#), except that the caller is only blocked until the timeout *abs\_timeout* expires.

Parameters

<i>sem</i>	the semaphore to be decremented;
<i>abs_timeout</i>	the timeout, expressed as an absolute value of the relevant clock for the semaphore, either CLOCK_MONOTONIC if SEM_RAWCLOCK was mentioned via sem_init_np(), or CLOCK_REALTIME otherwise.

Return values

0	on success;
---	-------------

## Return values

-1	with <i>errno</i> set if: <ul style="list-style-type: none"><li>• EPERM, the caller context is invalid;</li><li>• EINVAL, the semaphore is invalid or uninitialized;</li><li>• EINVAL, the specified timeout is invalid;</li><li>• EPERM, the semaphore <i>sm</i> is not process-shared and does not belong to the current process;</li><li>• EINTR, the caller was interrupted by a signal while blocked in this service;</li><li>• ETIMEDOUT, the semaphore could not be decremented and the specified timeout expired.</li></ul>
----	---

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

6.60.2.6 `sem_trywait()`

```
int sem_trywait (  
    sem_t * sem )
```

### Attempt to decrement a semaphore

This service is equivalent to [sem\\_wait\(\)](#), except that it returns immediately if the semaphore *sem* is currently depleted, and that it is not a cancellation point.

Parameters

<i>sem</i>	the semaphore to be decremented.
------------	----------------------------------

Return values

0	on success;
---	-------------

## Return values

-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EINVAL, the specified semaphore is invalid or uninitialized;</li> <li>• EPERM, the semaphore <i>sem</i> is not process-shared and does not belong to the current process;</li> <li>• EAGAIN, the specified semaphore is currently fully depleted.</li> <li>•</li> </ul>
----	--

## See also

[Specification.](#)

## Tags

[xthread-only](#)

## 6.60.2.7 sem\_unlink()

```
int sem_unlink (
    const char * name )
```

## Unlink a named semaphore

This service unlinks the semaphore named *name*. This semaphore is not destroyed until all references obtained with `sem_open()` are closed by calling `sem_close()`. However, the unlinked semaphore may no longer be reached with the `sem_open()` service.

When a semaphore is destroyed, the memory it used is returned to the system heap, so that further references to this semaphore are not guaranteed to fail, as is the case for unnamed semaphores.

## Parameters

<i>name</i>	the name of the semaphore to be unlinked.
-------------	---

## Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• ENAMETOOLONG, the length of the <i>name</i> argument exceeds 64 characters;</li> <li>• ENOENT, the named semaphore does not exist.</li> </ul>

See also

[Specification.](#)

Tags

[thread-unrestricted](#), [switch-secondary](#)

### 6.60.2.8 sem\_wait()

```
int sem_wait (
    sem_t * sem )
```

Decrement a semaphore

This service decrements the semaphore *sem* if it is currently if its value is greater than 0. If the semaphore's value is currently zero, the calling thread is suspended until the semaphore is posted, or a signal is delivered to the calling thread.

This service is a cancellation point for Cobalt threads (created with the [pthread\\_create\(\)](#) service). When such a thread is cancelled while blocked in a call to this service, the semaphore state is left unchanged before the cancellation cleanup handlers are called.

Parameters

<i>sem</i>	the semaphore to be decremented.
------------	----------------------------------

Return values

0	on success;
-1	with <i>errno</i> set if: <ul style="list-style-type: none"> <li>• EPERM, the caller context is invalid;</li> <li>• EINVAL, the semaphore is invalid or uninitialized;</li> <li>• EPERM, the semaphore <i>sem</i> is not process-shared and does not belong to the current process;</li> <li>• EINTR, the caller was interrupted by a signal while blocked in this service.</li> </ul>

See also

[Specification.](#)

Tags

[xthread-only](#), [switch-primary](#)

## 6.61 Thread management

Cobalt (POSIX) thread management services.

Collaboration diagram for Thread management:



### Functions

- int [pthread\\_create](#) (pthread\_t \*ptid\_r, const pthread\_attr\_t \*attr, void \*(\*start)(void \*), void \*arg)

*Create a new thread*

- int [pthread\\_setmode\\_np](#) (int clrmask, int setmask, int \*mode\_r)

*Set the mode of the current thread.*

- int [pthread\\_setname\\_np](#) (pthread\_t thread, const char \*name)

*Set a thread name.*

- int [pthread\\_kill](#) (pthread\_t thread, int sig)

*Send a signal to a thread.*

- int [pthread\\_join](#) (pthread\_t thread, void \*\*retval)

*Wait for termination of a specified thread.*

### 6.61.1 Detailed Description

Cobalt (POSIX) thread management services.

See also

[Specification.](#)

### 6.61.2 Function Documentation

## 6.61.2.1 pthread\_create()

```
int pthread_create (
    pthread_t * ptid_r,
    const pthread_attr_t * attr,
    void *(*)(void *) start,
    void * arg )
```

Create a new thread

This service creates a thread managed by the Cobalt core in a dual kernel configuration.

Attributes of the new thread depend on the *attr* argument. If *attr* is NULL, default values for these attributes are used.

Returning from the *start* routine has the same effect as calling `pthread_exit()` with the return value.

Parameters

<i>ptid_r</i>	address where the identifier of the new thread will be stored on success;
<i>attr</i>	thread attributes;
<i>start</i>	thread start routine;
<i>arg</i>	opaque user-supplied argument passed to <i>start</i> ;

Returns

0 on success;

an error number if:

- EINVAL, *attr* is invalid;
- EAGAIN, insufficient memory available from the system heap to create a new thread, increase CONFIG\_XENO\_OPT\_SYS\_HEAPSZ;
- EINVAL, thread attribute *inheritsched* is set to PTHREAD\_INHERIT\_SCHED and the calling thread does not belong to the Cobalt interface;

See also

[Specification.](#)

Note

When creating a Cobalt thread for the first time, libcobalt installs an internal handler for the SIGSHADOW signal. If you had previously installed a handler for such signal before that point, such handler will be exclusively called for any SIGSHADOW occurrence Xenomai did not send.

If, however, an application-defined handler for SIGSHADOW is installed afterwards, overriding the libcobalt handler, the new handler is required to call `cobalt_sigshadow_handler()` on entry. This routine returns a non-zero value for every occurrence of SIGSHADOW issued by the Cobalt core. If zero instead, the application-defined handler should process the signal.

```
int cobalt_sigshadow_handler(int sig, siginfo_t *si, void *ctxt);
```

You should register your handler with `sigaction(2)`, setting the SA\_SIGINFO flag.

Tags

[thread-unrestricted](#), [switch-secondary](#)

### 6.61.2.2 pthread\_join()

```
int pthread_join (
    pthread_t thread,
    void ** retval )
```

Wait for termination of a specified thread.

If *thread* is running and joinable, this service blocks the caller until *thread* terminates or detaches. When *thread* terminates, the caller is unblocked and its return value is stored at the address *value\_ptr*.

On the other hand, if *thread* has already finished execution, its return value collected earlier is stored at the address *value\_ptr* and this service returns immediately.

This service is a cancelation point for Cobalt threads: if the calling thread is canceled while blocked in a call to this service, the cancelation request is honored and *thread* remains joinable.

Multiple simultaneous calls to [pthread\\_join\(\)](#) specifying the same running target thread block all the callers until the target thread terminates.

Parameters

<i>thread</i>	identifier of the thread to wait for;
<i>retval</i>	address where the target thread return value will be stored on success.

Returns

- 0 on success;  
an error number if:
- ESRCH, *thread* is invalid;
  - EDEADLK, attempting to join the calling thread;
  - EINVAL, *thread* is detached;
  - EPERM, the caller context is invalid.

See also

[Specification](#).

Tags

[xthread-only](#), [switch-secondary](#), [switch-primary](#)



## 6.61.2.3 pthread\_kill()

```
int pthread_kill (
    pthread_t thread,
    int sig )
```

Send a signal to a thread.

This service send the signal *sig* to the Cobalt thread *thread* (created with [pthread\\_create\(\)](#)). If *sig* is zero, this service check for existence of the thread *thread*, but no signal is sent.

Parameters

<i>thread</i>	thread identifier;
<i>sig</i>	signal number.

Returns

0 on success;  
an error number if:

- EINVAL, *sig* is an invalid signal number;
- EAGAIN, the maximum number of pending signals has been exceeded;
- ESRCH, *thread* is an invalid thread identifier.

See also

[Specification](#).

Tags

[thread-unrestricted](#), [switch-primary](#)

## 6.61.2.4 pthread\_setmode\_np()

```
int pthread_setmode_np (
    int clrmask,
    int setmask,
    int * mode_r )
```

Set the mode of the current thread.

This service sets the mode of the calling thread, which affects its behavior under particular circumstances. *clrmask* and *setmask* are two masks of mode bits which are respectively cleared and set by [pthread\\_setmode\\_np\(\)](#):

- `PTHREAD_LOCK_SCHED`, when set, locks the scheduler, which prevents the current thread from being switched out until the scheduler is unlocked. Unless `PTHREAD_DISABLE_LOCKBREAK` is also set, the thread may still block, dropping the lock temporarily, in which case, the lock will be reacquired automatically when the thread resumes execution. When `PTHREAD_LOCK_SCHED` is cleared, the current thread drops the scheduler lock, and the rescheduling procedure is initiated.
- When set, `PTHREAD_WARNSW` enables debugging notifications for the current thread. A `SIGDEBUG` (Linux-originated) signal is sent when the following atypical or abnormal behavior is detected:
  - the current thread switches to secondary mode. Such notification comes in handy for detecting spurious relaxes, with one of the following reason codes:
    - \* `SIGDEBUG_MIGRATE_SYSCALL`, if the thread issued a regular Linux system call.
    - \* `SIGDEBUG_MIGRATE_SIGNAL`, if the thread had to leave real-time mode for handling a Linux signal.
    - \* `SIGDEBUG_MIGRATE_FAULT`, if the thread had to leave real-time mode for handling a processor fault/exception.
  - the current thread is sleeping on a Cobalt mutex currently owned by a thread running in secondary mode, which reveals a priority inversion. In such an event, the reason code passed to the signal handler will be `SIGDEBUG_MIGRATE_PRIOINV`.
  - the current thread is about to sleep while holding a Cobalt mutex, and `CONFIG_XENO_OPT_DEBUG_MUTEX_SLEEP` is enabled in the kernel configuration. In such an event, the reason code passed to the signal handler will be `SIGDEBUG_MUTEX_SLEEP`. Blocking for acquiring a mutex does not trigger such signal though.
  - the current thread has enabled `PTHREAD_DISABLE_LOCKBREAK` and `PTHREAD_LOCK_SCHED`, then attempts to block on a Cobalt service, which would cause a lock break. In such an event, the reason code passed to the signal handler will be `SIGDEBUG_LOCK_BREAK`.
- `PTHREAD_DISABLE_LOCKBREAK` disallows breaking the scheduler lock. Normally, the scheduler lock is dropped implicitly when the current owner blocks, then reacquired automatically when the owner resumes execution. If `PTHREAD_DISABLE_LOCKBREAK` is set, the scheduler lock owner would return with `EINTR` immediately from any blocking call instead (see `PTHREAD_WARNSW` notifications).
- `PTHREAD_CONFORMING` can be passed in `setmask` to switch the current Cobalt thread to its preferred runtime mode. The only meaningful use of this switch is to force a real-time thread back to primary mode eagerly. Other usages have no effect.

This service is a non-portable extension of the Cobalt interface.

#### Parameters

<code>clrmask</code>	set of bits to be cleared.
<code>setmask</code>	set of bits to be set.
<code>mode_r</code>	If non-NULL, <code>mode_r</code> must be a pointer to a memory location which will be written upon success with the previous set of active mode bits. If NULL, the previous set of active mode bits will not be returned.

#### Returns

0 on success, otherwise:

- `EINVAL`, some bit in `clrmask` or `setmask` is invalid.

## Note

Setting *clrmask* and *setmask* to zero leads to a nop, only returning the previous mode if *mode\_r* is a valid address.

## Attention

Issuing `PTHREAD_CONFORMING` is most likely useless or even introduces pure overhead in regular applications, since the Cobalt core performs the necessary mode switches, only when required.

## Tags

[xthread-only](#), [switch-primary](#)

## 6.61.2.5 pthread\_setname\_np()

```
int pthread_setname_np (
    pthread_t thread,
    const char * name )
```

Set a thread name.

This service set to *name*, the name of *thread*. This name is used for displaying information in `/proc/xenomai/sched`.

This service is a non-portable extension of the Cobalt interface.

## Parameters

<i>thread</i>	target thread;
<i>name</i>	name of the thread.

## Returns

- 0 on success;  
an error number if:
- ESRCH, *thread* is invalid.


## Tags

[xthread-only](#)

## 6.62 Scheduling management

Cobalt scheduling management services.

Collaboration diagram for Scheduling management:



### Functions

- int [pthread\\_setschedparam](#) (pthread\_t thread, int policy, const struct sched\_param \*param)

*Set the scheduling policy and parameters of the specified thread.*

- int [pthread\\_setschedparam\\_ex](#) (pthread\_t thread, int policy, const struct sched\_param\_ex \*param\_ex)

*Set extended scheduling policy of thread*

- int [pthread\\_getschedparam](#) (pthread\_t thread, int \*\_\_restrict\_\_ policy, struct sched\_param \*\_\_restrict\_\_ param)

*Get the scheduling policy and parameters of the specified thread.*

- int [pthread\\_getschedparam\\_ex](#) (pthread\_t thread, int \*\_\_restrict\_\_ policy\_r, struct sched\_param\_ex \*\_\_restrict\_\_ param\_ex)

*Get extended scheduling policy of thread*

- int [pthread\\_yield](#) (void)

*Yield the processor.*

### 6.62.1 Detailed Description

Cobalt scheduling management services.

### 6.62.2 Function Documentation

## 6.62.2.1 pthread\_getschedparam()

```
int pthread_getschedparam (
    pthread_t thread,
    int *__restrict__ policy,
    struct sched_param *__restrict__ param )
```

Get the scheduling policy and parameters of the specified thread.

This service returns, at the addresses *policy* and *par*, the current scheduling policy and scheduling parameters (i.e. priority) of the Cobalt thread *tid*. If *thread* is not the identifier of a Cobalt thread, this service fallback to the regular POSIX [pthread\\_getschedparam\(\)](#) service.

## Parameters

<i>thread</i>	target thread;
<i>policy</i>	address where the scheduling policy of <i>tid</i> is stored on success;
<i>param</i>	address where the scheduling parameters of <i>tid</i> is stored on success.

## Returns

- 0 on success;
- an error number if:
  - ESRCH, *tid* is invalid.

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

Referenced by [pthread\\_getschedparam\\_ex\(\)](#).

## 6.62.2.2 pthread\_getschedparam\_ex()

```
int pthread_getschedparam_ex (
    pthread_t thread,
    int *__restrict__ policy_r,
    struct sched_param_ex *__restrict__ param_ex )
```

Get extended scheduling policy of thread

This service is an extended version of the regular [pthread\\_getschedparam\(\)](#) service, which also supports Cobalt-specific policies, not available with the host Linux environment.

## Parameters

<i>thread</i>	target thread;
<i>policy_r</i>	address where the scheduling policy of <i>thread</i> is stored on success;
<i>param_ex</i>	address where the scheduling parameters of <i>thread</i> are stored on success.

## Returns

- 0 on success;  
 an error number if:
- ESRCH, *thread* is invalid.

## See also

[Specification.](#)

## Tags

[thread-unrestricted](#)

References `pthread_getschedparam()`.

6.62.2.3 `pthread_setschedparam()`

```
int pthread_setschedparam (
    pthread_t thread,
    int policy,
    const struct sched_param * param )
```

Set the scheduling policy and parameters of the specified thread.

This service set the scheduling policy of the Cobalt thread identified by *pid* to the value *policy*, and its scheduling parameters (i.e. its priority) to the value pointed to by *param*.

If `pthread_self()` is passed, this service turns the current thread into a Cobalt thread. If *thread* is not the identifier of a Cobalt thread, this service falls back to the regular `pthread_setschedparam()` service.

## Parameters

<i>thread</i>	target Cobalt thread;
<i>policy</i>	scheduling policy, one of SCHED_FIFO, SCHED_RR, or SCHED_OTHER;
<i>param</i>	address of scheduling parameters.

## Returns

0 on success;

an error number if:

- ESRCH, *pid* is invalid;
- EINVAL, *policy* or *param->sched\_priority* is invalid;
- EAGAIN, insufficient memory available from the system heap, increase CONFIG\_XENO\_OPT\_SYS\_HEAPSZ;
- EFAULT, *param* is an invalid address;

## See also

[Specification.](#)

## Note

See [pthread\\_create\(\)](#), [pthread\\_setschedparam\\_ex\(\)](#).

## Tags

[thread-unrestricted](#), [switch-secondary](#), [switch-primary](#)

Referenced by [pthread\\_setschedparam\\_ex\(\)](#).

## 6.62.2.4 pthread\_setschedparam\_ex()

```
int pthread_setschedparam_ex (
    pthread_t thread,
    int policy,
    const struct sched_param_ex * param_ex )
```

## Set extended scheduling policy of thread

This service is an extended version of the regular [pthread\\_setschedparam\(\)](#) service, which supports Cobalt-specific scheduling policies, not available with the host Linux environment.

This service set the scheduling policy of the Cobalt thread *thread* to the value *policy*, and its scheduling parameters (e.g. its priority) to the value pointed to by *param\_ex*.

If *thread* does not match the identifier of a Cobalt thread, this action falls back to the regular [pthread\\_setschedparam\(\)](#) service.

## Parameters

<i>thread</i>	target Cobalt thread;
<i>policy</i>	scheduling policy, one of SCHED_WEAK, SCHED_FIFO, SCHED_COBALT, SCHED_RR, SCHED_SPORADIC, SCHED_TP, SCHED_QUOTA or SCHED_NORMAL;
<i>param_ex</i>	scheduling parameters address. As a special exception, a negative sched_priority value is interpreted as if SCHED_WEAK was given in <i>policy</i> , using the absolute value of this parameter as the weak priority level.

When `CONFIG_XENO_OPT_SCHED_WEAK` is enabled, `SCHED_WEAK` exhibits priority levels in the `[0..99]` range (inclusive). Otherwise, `sched_priority` must be zero for the `SCHED_WEAK` policy.

#### Returns

0 on success;

an error number if:

- `ESRCH`, *thread* is invalid;
- `EINVAL`, *policy* or *param\_ex->sched\_priority* is invalid;
- `EAGAIN`, insufficient memory available from the system heap, increase `CONFIG_XENO_OPT_SYS_HEAPSZ`;
- `EFAULT`, *param\_ex* is an invalid address;
- `EPERM`, the calling process does not have superuser permissions.

See also

[Specification.](#)

Note

See [pthread\\_create\(\)](#), [pthread\\_setschedparam\(\)](#).

Tags

[thread-unrestricted](#), [switch-secondary](#), [switch-primary](#)

References [pthread\\_setschedparam\(\)](#).

#### 6.62.2.5 pthread\_yield()

```
int pthread_yield (
    void )
```

Yield the processor.

This function move the current thread at the end of its priority group.

Return values

0	
---	--



See also

[Specification.](#)

Tags

[thread-unrestricted](#), [switch-primary](#)

## 6.63 Smokey API

A simple infrastructure for writing and running smoke tests.

A simple infrastructure for writing and running smoke tests.

Smokey is based on the Copperplate API, therefore is available over the single and dual kernel Xenomai configurations indifferently.

The API provides a set of services for declaring any number of test plugins, embodied into a test program. Each plugin usually implements a single smoke test, checking a particular feature of interest. Each plugin present in the running executable is automatically detected by the Smokey init routine. In addition, the Smokey API parses all arguments and options passed on the command line to the executable, running pre-defined actions which are therefore automatically recognized by all programs linked against the Smokey library.

### Writing smoke tests with Smokey

A smoke test is composed of a routine which implements the test code, and a set of runtime settings/attributes for running such code. The routine prototype shall be:

```
int run_<test_name>(struct smokey_test *t, int argc, char *const argv[])
```

The test routine should return a zero value for success, or any negated POSIX error code for indicating the failure to the test driver (e.g. -EINVAL if some value is found to be wrong).

With *t* referring to the Smokey test descriptor, and *argc*, *argv* the argument count and vector expunged from all the inner options which may have been previously interpreted by the Smokey API and inner layers (such as Copperplate).

The Smokey API provides the services to declare a complete test (named **foo** in this example) as follows:

```
#include <smokey/smokey.h>

smokey_test_plugin(foo, // test name
                  SMOKEY_ARGLIST( // argument list
                                SMOKEY_INT(some_integer),
                                SMOKEY_STRING(some_string),
                                SMOKEY_BOOL(some_boolean),
                                ),
                  // description
                  "A dummy Smokey-based test plugin\n"
                  "\taccepting three optional arguments:\n"
                  "\tsome_integer=<value>\n"
                  "\tsome_string=<string>\n"
                  "\tsome_bool[=0/1]\n"
);

static int run_foo(struct smokey_test *t, int argc, char *const argv[])
{
    int i_arg = 0, nargs;
    char *s_arg = NULL;
    bool b_arg = false;

    nargs = smokey_parse_args(t, argc, argv);

    if (SMOKEY_ARG_ISSET(foo, some_integer))
        i_arg = SMOKEY_ARG_INT(foo, some_integer);
    if (SMOKEY_ARG_ISSET(foo, some_string))
        s_arg = SMOKEY_ARG_STRING(foo, some_string);
    if (SMOKEY_ARG_ISSET(foo, some_boolean))
        b_arg = SMOKEY_ARG_BOOL(foo, some_boolean);

    return run_some_hypothetical_smoke_test_code(i_arg, s_arg, b_arg);
}
```

As illustrated, a smoke test is at least composed of a test plugin descriptor (i.e. `smokey_test_plugin()`), and a run handler named after the test.

#### Test arguments

Smokey recognizes three argument declarators, namely: `SMOKEY_INT(name)` for a C (signed) integer, `SMOKEY_BOOL(name)` for a boolean value and `SMOKEY_STRING(name)` for a character string.

Each argument can be passed to the test code as a name=value pair, where *name* should match one of the declarators. Before the test-specific arguments can be accessed, a call to `smokey_parse_args()` must be issued by the test code, passing the parameters received in the run handler. This routine returns the number of arguments found on the command line matching the an entry in `SMOKEY_ARGLIST()`.

Once `smokey_parse_args()` has returned with a non-zero value, each argument can be checked individually for presence. If a valid argument was matched on the command line, `SMOKEY_ARG_ISS->ET(test_name, arg_name)` returns non-zero. In the latter case, its value can be retrieved by a similar call to `SMOKEY_ARG_INT(test_name, arg_name)`, `SMOKEY_ARG_STRING(test_name, arg_name)` or `SMOKEY_ARG_BOOL(test_name, arg_name)`.

In the above example, passing "some\_integer=3" on the command line of any program implementing such Smokey-based test would cause the variable `i_arg` to receive "3" as a value.

#### Pre-defined Smokey options

Any program linked against the Smokey API implicitly recognizes the following options:

- `-list[=<id[,id...]>]` dumps the list of tests implemented in the program to stdout. This list may be restricted to the tests matching the optional regular expression (see `-run`). The information given includes the description strings provided in the plugin declarators (`smokey_test_plugin()`). The position and symbolic name of each test is also issued, which may be used in id specifications with the `-run` option (see below).

#### Note

Test positions may vary depending on changes to the host program like adding or removing other tests, the symbolic name however is stable and identifies each test uniquely.

- `-run[=<id[,id...]>]` selects the tests to be run, determining the active test list among the overall set of tests detected in the host program. The test driver code (e.g. implementing a test harness program on top of Smokey) may then iterate over the `smokey_test_list` for accessing each active test individually, in the enumeration order specified by the user (Use `for_each_smokey_test()` for that).

If no argument is passed to `-run`, Smokey assumes that all tests detected in the current program should be picked, filling `smokey_test_list` with tests by increasing position order.

Otherwise, `id` may be a test position, a symbolic name, or a range thereof delimited by a dash character. A symbolic name may be matched using a `glob(3)` type regular expression.

`id` specification may be:

- 0-9, picks tests #0 to #9

- -3, picks tests #0 to #3
  - 5-, picks tests #5 to the highest possible test position
  - 2-0, picks tests #2 to #0, in decreasing order
  - foo, picks test foo only
  - 0,1,foo- picks tests #0, #1, and any test from foo up to the last test defined
  - fo\* picks any test with a name starting by "fo"
- `–exclude=<id[,id...]>` excludes the given tests from the test list. The format of the argument is identical to the one accepted by the `–run` option.
  - `–keep-going` sets the boolean flag `smokey_keep_going` to a non-zero value, indicating to the test driver that receiving a failure code from a smoke test should not abort the test loop. This flag is not otherwise interpreted by the Smokey API.
  - `–verbose[=level]` sets the integer `smokey_verbose_mode` to a non-zero value, which should be interpreted by all parties as the desired verbosity level (defaults to 1).
  - `–vm` gives a hint to the test code, about running in a virtual environment, such as KVM. When passed, the boolean `smokey_on_vm` is set. Each test may act upon this setting, such as skipping time-dependent checks that may fail due to any slowdown induced by the virtualization.

#### Writing a test driver based on the Smokey API

A test driver provides the `main()` entry point, which should iterate over the test list (`smokey_test_list`) prepared by the Smokey API, for running each test individually. The `for_each_smokey_test()` helper is available for iterating over the active test list.

When this entry point is called, all the initialization chores, including the test detection and the active test selection have been performed by the Smokey API already.

#### Issuing information notices

The printf-like `smokey_note()` routine is available for issuing notices to the output device (currently `stdout`), unless `–silent` was detected on the command line. `smokey_note()` outputs a terminating newline character. Notes are enabled for any verbosity level greater than zero.

#### Issuing trace messages

The printf-like `smokey_trace()` routine is available for issuing progress messages to the output device (currently `stdout`), unless `–silent` was detected on the command line. `smokey_trace()` outputs a terminating newline character. Traces are enabled for any verbosity level greater than one.

Therefore, a possible implementation of a test driver could be as basic as:

```
#include <stdio.h>
#include <error.h>
#include <smokey/smokey.h>

int main(int argc, char *const argv[])
{
    struct smokey_test *t;
    int ret;


    if (pvlist_empty(&smokey_test_list))
        return 0;

    for_each_smokey_test(t) {
        ret = t->run(t, argc, argv);
        if (ret) {
            if (smokey_keep_going)
                continue;
            error(1, -ret, "test %s failed", t->name);
        }
        smokey_note("%s OK", t->name);
    }

    return 0;
}
```

## 6.64 Asynchronous acquisition API

Collaboration diagram for Asynchronous acquisition API:



### Data Structures

- struct [a4l\\_cmd\\_desc](#)  
*Structure describing the asynchronous instruction.*

### Functions

- int [a4l\\_snd\\_command](#) ([a4l\\_desc\\_t](#) \*dsc, [a4l\\_cmd\\_t](#) \*cmd)  
*Send a command to an Analogy device.*
- int [a4l\\_snd\\_cancel](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd)  
*Cancel an asynchronous acquisition.*
- int [a4l\\_set\\_bufsize](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned long size)  
*Change the size of the asynchronous buffer.*
- int [a4l\\_get\\_bufsize](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned long \*size)  
*Get the size of the asynchronous buffer.*
- int [a4l\\_mark\\_bufrw](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned long cur, unsigned long \*new)  
*Update the asynchronous buffer state.*
- int [a4l\\_poll](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned long ms\_timeout)  
*Get the available data count.*
- int [a4l\\_mmap](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned long size, void \*\*ptr)  
*Map the asynchronous ring-buffer into a user-space.*

### ANALOGY\_CMD\_XXX

#### Common command flags definitions

- #define [A4L\\_CMD\\_SIMUL](#) 0x1  
*Do not execute the command, just check it.*
- #define [A4L\\_CMD\\_BULK](#) 0x2  
*Perform data recovery / transmission in bulk mode.*
- #define [A4L\\_CMD\\_WRITE](#) 0x4  
*Perform a command which will write data to the device.*

## TRIG\_XXX

## Command triggers flags definitions

- #define TRIG\_NONE 0x00000001  
*Never trigger.*
- #define TRIG\_NOW 0x00000002  
*Trigger now + N ns.*
- #define TRIG\_FOLLOW 0x00000004  
*Trigger on next lower level trig.*
- #define TRIG\_TIME 0x00000008  
*Trigger at time N ns.*
- #define TRIG\_TIMER 0x00000010  
*Trigger at rate N ns.*
- #define TRIG\_COUNT 0x00000020  
*Trigger when count reaches N.*
- #define TRIG\_EXT 0x00000040  
*Trigger on external signal N.*
- #define TRIG\_INT 0x00000080  
*Trigger on analogy-internal signal N.*
- #define TRIG\_OTHER 0x00000100  
*Driver defined trigger.*
- #define TRIG\_WAKE\_EOS 0x0020  
*Wake up on end-of-scan.*
- #define TRIG\_ROUND\_MASK 0x00030000  
*Trigger not implemented yet.*
- #define TRIG\_ROUND\_NEAREST 0x00000000  
*Trigger not implemented yet.*
- #define TRIG\_ROUND\_DOWN 0x00010000  
*Trigger not implemented yet.*
- #define TRIG\_ROUND\_UP 0x00020000  
*Trigger not implemented yet.*
- #define TRIG\_ROUND\_UP\_NEXT 0x00030000  
*Trigger not implemented yet.*

## Channel macros

Specific precompilation macros and constants useful for the channels descriptors tab located in the command structure

- #define CHAN(a) ((a) & 0xffff)  
*Channel indication macro.*
- #define RNG(a) (((a) & 0xff) << 16)  
*Range definition macro.*
- #define AREF(a) (((a) & 0x03) << 24)  
*Reference definition macro.*
- #define FLAGS(a) ((a) & CR\_FLAGS\_MASK)  
*Flags definition macro.*

- #define `PACK(a, b, c)` (`a` | `RNG(b)` | `AREF(c)`)  
*Channel + range + reference definition macro.*
- #define `PACK_FLAGS(a, b, c, d)` (`PACK(a, b, c)` | `FLAGS(d)`)  
*Channel + range + reference + flags definition macro.*
- #define `AREF_GROUND` 0x00  
*Analog reference is analog ground.*
- #define `AREF_COMMON` 0x01  
*Analog reference is analog common.*
- #define `AREF_DIFF` 0x02  
*Analog reference is differential.*
- #define `AREF_OTHER` 0x03  
*Analog reference is undefined.*

### 6.64.1 Detailed Description

### 6.64.2 Function Documentation

#### 6.64.2.1 `a4l_get_bufsize()`

```
int a4l_get_bufsize (
    a4l_desc_t * dsc,
    unsigned int idx_subd,
    unsigned long * size )
```

Get the size of the asynchronous buffer.

During asynchronous acquisition, a ring-buffer enables the transfers from / to user-space. Functions like `a4l_read()` or `a4l_write()` recovers / sends data through this intermediate buffer. Please note, there is one ring-buffer per subdevice capable of asynchronous acquisition. By default, each buffer size is set to 64 KB.

Parameters

in	<code>dsc</code>	Device descriptor filled by <code>a4l_open()</code> (and optionally <code>a4l_fill_desc()</code> )
in	<code>idx_subd</code>	Index of the concerned subdevice
out	<code>size</code>	Buffer size

Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong

References `a4l_descriptor::fd`.



## 6.64.2.2 a4l\_mark\_bufrw()

```
int a4l_mark_bufrw (
    a4l_desc_t * dsc,
    unsigned int idx_subd,
    unsigned long cur,
    unsigned long * new )
```

Update the asynchronous buffer state.

When the mapping of the asynchronous ring-buffer (thanks to [a4l\\_mmap\(\)](#) is disabled, common read / write syscalls have to be used. In input case, [a4l\\_read\(\)](#) must be used for:

- the retrieval of the acquired data.
- the notification to the Analogy layer that the acquired data have been consumed, then the area in the ring-buffer which was containing becomes available. In output case, [a4l\\_write\(\)](#) must be called to:
- send some data to the Analogy layer.
- signal the Analogy layer that a chunk of data in the ring-buffer must be used by the driver.

In mmap configuration, these features are provided by unique function named [a4l\\_mark\\_bufrw\(\)](#). In input case, [a4l\\_mark\\_bufrw\(\)](#) can :

- recover the count of data newly available in the ring-buffer.
- notify the Analogy layer how many bytes have been consumed. In output case, [a4l\\_mark\\_bufrw\(\)](#) can:
- recover the count of data available for writing.
- notify Analogy that some bytes have been written.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
in	<i>idx_subd</i>	Index of the concerned subdevice
in	<i>cur</i>	Amount of consumed data
out	<i>new</i>	Amount of available data

## Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong; the descriptor and the new pointer should be checked; check also the kernel log ("dmesg")
- -EFAULT is returned if a user <-> kernel transfer went wrong

References [a4l\\_descriptor::fd](#).

## 6.64.2.3 a4l\_mmap()

```
int a4l_mmap (
    a4l_desc_t * dsc,
    unsigned int idx_subd,
    unsigned long size,
    void ** ptr )
```

Map the asynchronous ring-buffer into a user-space.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
in	<i>idx_subd</i>	Index of the concerned subdevice
in	<i>size</i>	Size of the buffer to map
out	<i>ptr</i>	Address of the pointer containing the assigned address on return

## Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong, the descriptor and the pointer should be checked; check also the kernel log
- -EPERM is returned if the function is called in an RT context or if the buffer to resize is mapped in user-space (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -EBUSY is returned if the buffer is already mapped in user-space

References [a4l\\_descriptor::fd](#).

## 6.64.2.4 a4l\_poll()

```
int a4l_poll (
    a4l_desc_t * dsc,
    unsigned int idx_subd,
    unsigned long ms_timeout )
```

Get the available data count.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
in	<i>idx_subd</i>	Index of the concerned subdevice
in	<i>ms_timeout</i>	The number of milliseconds to wait for some data to be available. Passing A4L_INFINITE causes the caller to block indefinitely until some data is available. Passing A4L_NONBLOCK causes the function to return immediately without waiting for any available data

## Returns

the available data count. Otherwise:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -EINTR is returned if calling task has been unblocked by a signal

References `a4l_descriptor::fd`.

6.64.2.5 `a4l_set_bufsize()`

```
int a4l_set_bufsize (
    a4l_desc_t * dsc,
    unsigned int idx_subd,
    unsigned long size )
```

Change the size of the asynchronous buffer.

During asynchronous acquisition, a ring-buffer enables the transfers from / to user-space. Functions like `a4l_read()` or `a4l_write()` recovers / sends data through this intermediate buffer. The function `a4l_set_bufsize()` can change the size of the ring-buffer. Please note, there is one ring-buffer per subdevice capable of asynchronous acquisition. By default, each buffer size is set to 64 KB.

## Parameters

in	<code>dsc</code>	Device descriptor filled by <code>a4l_open()</code> (and optionally <code>a4l_fill_desc()</code> )
in	<code>idx_subd</code>	Index of the concerned subdevice
in	<code>size</code>	New buffer size, the maximal tolerated value is 16MB ( <code>A4L_BUF_MAXSIZE</code> )

## Returns

0 on success. Otherwise:

- -EINVAL is returned if the analogy descriptor is not correct or if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EPERM is returned if the function is called in an RT context or if the buffer to resize is mapped in user-space (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -EBUSY is returned if the selected subdevice is already processing an asynchronous operation
- -ENOMEM is returned if the system is out of memory

References `a4l_sys_bufcfg()`, and `a4l_descriptor::fd`.

6.64.2.6 `a4l_snd_cancel()`

```
int a4l_snd_cancel (
    a4l_desc_t * dsc,
    unsigned int idx_subd )
```

Cancel an asynchronous acquisition.

The function `a4l_snd_cancel()` is devoted to stop an asynchronous acquisition configured thanks to an Analogy command.

Parameters

in	<i>dsc</i>	Device descriptor filled by <code>a4l_open()</code> (and optionally <code>a4l_fill_desc()</code> )
in	<i>idx_subd</i>	Subdevice index

Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EIO is returned if the selected subdevice does not support asynchronous operation

References `a4l_descriptor::fd`.

6.64.2.7 `a4l_snd_command()`

```
int a4l_snd_command (
    a4l_desc_t * dsc,
    a4l_cmd_t * cmd )
```

Send a command to an Analogy device.

The function `a4l_snd_command()` triggers asynchronous acquisition.

Parameters

in	<i>dsc</i>	Device descriptor filled by <code>a4l_open()</code> (and optionally <code>a4l_fill_desc()</code> )
in	<i>cmd</i>	Command structure

Returns

0 on success. Otherwise:


- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -ENOMEM is returned if the system is out of memory

- -EFAULT is returned if a user <-> kernel transfer went wrong
- -EIO is returned if the selected subdevice cannot handle command
- -EBUSY is returned if the selected subdevice is already processing an asynchronous operation

References `a4l_descriptor::fd`.

## 6.65 Asynchronous acquisition API

Collaboration diagram for Asynchronous acquisition API:



### Functions

- int [a4l\\_async\\_read](#) ([a4l\\_desc\\_t](#) \*dsc, void \*buf, size\_t nbyte, unsigned long ms\_timeout)  
*Perform asynchronous read operation on the analog input subdevice.*
- int [a4l\\_async\\_write](#) ([a4l\\_desc\\_t](#) \*dsc, void \*buf, size\_t nbyte, unsigned long ms\_timeout)  
*Perform asynchronous write operation on the analog input subdevice.*

#### 6.65.1 Detailed Description

#### 6.65.2 Function Documentation

##### 6.65.2.1 a4l\_async\_read()

```

int a4l_async_read (
    a4l_desc_t * dsc,
    void * buf,
    size_t nbyte,
    unsigned long ms_timeout )
  
```

Perform asynchronous read operation on the analog input subdevice.

The function [a4l\\_async\\_read\(\)](#) is only useful for acquisition configured through an Analogy command.

#### Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
out	<i>buf</i>	Input buffer
in	<i>nbyte</i>	Number of bytes to read
in	<i>ms_timeout</i>	The number of milliseconds to wait for some data to be available. Passing A4L_INFINITE causes the caller to block indefinitely until some data is available. Passing A4L_NONBLOCK causes the function to return immediately without waiting for any available data

## Returns

Number of bytes read, otherwise negative error code:

- -EINVAL is returned if some argument is missing or wrong, the descriptor should be checked; check also the kernel log
- -ENOENT is returned if the device's reading subdevice is idle (no command was sent)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -EINTR is returned if calling task has been unblocked by a signal

## 6.65.2.2 a4l\_async\_write()

```
int a4l_async_write (
    a4l_desc_t * dsc,
    void * buf,
    size_t nbyte,
    unsigned long ms_timeout )
```

Perform asynchronous write operation on the analog input subdevice.

The function [a4l\\_async\\_write\(\)](#) is only useful for acquisition configured through an Analogy command.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
in	<i>buf</i>	Output buffer
in	<i>nbyte</i>	Number of bytes to write
in	<i>ms_timeout</i>	The number of milliseconds to wait for some free area to be available. Passing A4L_INFINITE causes the caller to block indefinitely until some data is available. Passing A4L_NONBLOCK causes the function to return immediately without waiting any available space to write data.


## Returns

Number of bytes written, otherwise negative error code:

- -EINVAL is returned if some argument is missing or wrong, the descriptor should be checked; check also the kernel log
- -ENOENT is returned if the device's reading subdevice is idle (no command was sent)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -EINTR is returned if calling task has been unblocked by a signal

## 6.66 Software calibration API

Collaboration diagram for Software calibration API:



### Functions

- int [a4l\\_read\\_calibration\\_file](#) (char \*name, struct a4l\_calibration\_data \*data)  
*Read the analogy generated calibration file.*
- int [a4l\\_get\\_softcal\\_converter](#) (struct a4l\_polynomial \*converter, int subd, int chan, int range, struct a4l\_calibration\_data \*data)  
*Get the polynomial that will be use for the software calibration.*
- int [a4l\\_rawtodcal](#) (a4l\_chinfo\_t \*chan, double \*dst, void \*src, int cnt, struct a4l\_polynomial \*converter)  
*Convert raw data (from the driver) to calibrated double units.*
- int [a4l\\_dcaltoraw](#) (a4l\_chinfo\_t \*chan, void \*dst, double \*src, int cnt, struct a4l\_polynomial \*converter)  
*Convert double values to raw calibrated data using polynomials.*

### 6.66.1 Detailed Description

### 6.66.2 Function Documentation

#### 6.66.2.1 a4l\_dcaltoraw()

```

int a4l_dcaltoraw (
    a4l_chinfo_t * chan,
    void * dst,
    double * src,
    int cnt,
    struct a4l_polynomial * converter )
  
```

Convert double values to raw calibrated data using polynomials.

#### Parameters

in	<i>chan</i>	Channel descriptor
out	<i>dst</i>	Ouput buffer
in	<i>src</i>	Input buffer
in	<i>cnt</i>	Count of conversion to perform
in	<i>converter</i>	Conversion polynomial



## Returns

the count of conversion performed, otherwise a negative error code:

- -EINVAL is returned if some argument is missing or wrong; *chan*, *rng* and the pointers should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_dcaltoraw\(\)](#)

## 6.66.2.2 a4l\_get\_softcal\_converter()

```
int a4l_get_softcal_converter (
    struct a4l_polynomial * converter,
    int subd,
    int chan,
    int range,
    struct a4l_calibration_data * data )
```

Get the polynomial that will be use for the software calibration.

## Parameters

out	<i>converter</i>	Polynomial to be used on the software calibration
in	<i>subd</i>	Subdevice index
in	<i>chan</i>	Channel
in	<i>range</i>	Range
in	<i>data</i>	Calibration data read from the calibration file

## Returns

-1 on error

## 6.66.2.3 a4l\_rawtodcal()

```
int a4l_rawtodcal (
    a4l_chinfo_t * chan,
    double * dst,
    void * src,
    int cnt,
    struct a4l_polynomial * converter )
```

Convert raw data (from the driver) to calibrated double units.

## Parameters

in	<i>chan</i>	Channel descriptor
out	<i>dst</i>	Ouput buffer
in	<i>src</i>	Input buffer
in	<i>cnt</i>	Count of conversion to perform
in	<i>converter</i>	Conversion polynomial

## Returns

the count of conversion performed, otherwise a negative error code:

- -EINVAL is returned if some argument is missing or wrong; chan, rng and the pointers should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_rawtodcal\(\)](#)

## 6.66.2.4 a4l\_read\_calibration\_file()

```
int a4l_read_calibration_file (
    char * name,
    struct a4l_calibration_data * data )
```


Read the analogy generated calibration file.

## Parameters

in	<i>name</i>	Name of the calibration file
out	<i>data</i>	Pointer to the calibration file contents

## 6.67 Descriptor Syscall API

Collaboration diagram for Descriptor Syscall API:



### Data Structures

- struct [a4l\\_descriptor](#)  
*Structure containing device-information useful to users.*

### Functions

- int [a4l\\_sys\\_desc](#) (int fd, [a4l\\_desc\\_t](#) \*dsc, int pass)  
*Get a descriptor on an attached device.*

### ANALOGY\_XXX\_DESC

Constants used as argument so as to define the description depth to recover

- #define [A4L\\_BSC\\_DESC](#) 0x0  
*BSC stands for basic descriptor (device data)*
- #define [A4L\\_CPLX\\_DESC](#) 0x1  
*CPLX stands for complex descriptor (subdevice + channel + range data)*

#### 6.67.1 Detailed Description

#### 6.67.2 Function Documentation

## 6.67.2.1 a4l\_sys\_desc()

```
int a4l_sys_desc (
    int fd,
    a4l_desc_t * dsc,
    int pass )
```

Get a descriptor on an attached device.

Once the device has been attached, the function `a4l_get_desc()` retrieves various information on the device (subdevices, channels, ranges, etc.). The function `a4l_get_desc()` can be called twice:

- The first time, almost all the fields, except `sldata`, are set (`board_name`, `nb_subd`, `idx_read_subd`, `idx_write_subd`, `magic`, `sbsize`); the last field, `sldata`, is supposed to be a pointer on a buffer, which size is defined by the field `sbsize`.
- The second time, the buffer pointed by `sldata` is filled with data about the subdevices, the channels and the ranges.

Between the two calls, an allocation must be performed in order to recover a buffer large enough to contain all the data. These data are set up according a root-leaf organization (device -> subdevice -> channel -> range). They cannot be accessed directly; specific functions are available so as to retrieve them:

- [a4l\\_get\\_subdinfo\(\)](#) to get some subdevice's characteristics.
- [a4l\\_get\\_chaninfo\(\)](#) to get some channel's characteristics.
- [a4l\\_get\\_rnginfo\(\)](#) to get some range's characteristics.

## Parameters

in	<i>fd</i>	Driver file descriptor
out	<i>dsc</i>	Device descriptor
in	<i>pass</i>	Description level to retrieve: <ul style="list-style-type: none"> <li>• <code>A4L_BSC_DESC</code> to get the basic descriptor (notably the size of the data buffer to allocate).</li> <li>• <code>A4L_CPLX_DESC</code> to get the complex descriptor, the data buffer is filled with characteristics about the subdevices, the channels and the ranges.</li> </ul>

## Returns

0 on success. Otherwise:


- `-EINVAL` is returned if some argument is missing or wrong; the `pass` argument should be checked; check also the kernel log ("dmesg")
- `-EFAULT` is returned if a user <-> kernel transfer went wrong
- `-ENODEV` is returned if the descriptor is incoherent (the device may be unattached)

References `A4L_BSC_DESC`, and `a4l_descriptor::magic`.

Referenced by `a4l_open()`.

## 6.68 Descriptor API

Collaboration diagram for Descriptor API:



### Functions

- int `a4l_open` (`a4l_desc_t *dsc`, const char \*fname)  
*Open an Analogy device and basically fill the descriptor.*
- int `a4l_close` (`a4l_desc_t *dsc`)  
*Close the Analogy device related with the descriptor.*
- int `a4l_fill_desc` (`a4l_desc_t *dsc`)  
*Fill the descriptor with subdevices, channels and ranges data.*
- int `a4l_get_subdinfo` (`a4l_desc_t *dsc`, unsigned int subd, `a4l_sbinfo_t **info`)  
*Get an information structure on a specified subdevice.*
- int `a4l_get_chinfo` (`a4l_desc_t *dsc`, unsigned int subd, unsigned int chan, `a4l_chinfo_t **info`)  
*Get an information structure on a specified channel.*
- int `a4l_get_rnginfo` (`a4l_desc_t *dsc`, unsigned int subd, unsigned int chan, unsigned int rng, `a4l_rnginfo_t **info`)  
*Get an information structure on a specified range.*

### 6.68.1 Detailed Description

This is the API interface used to fill and use Analogy device descriptor structure

### 6.68.2 Function Documentation

#### 6.68.2.1 `a4l_close()`

```
int a4l_close (
    a4l_desc_t * dsc )
```

Close the Analogy device related with the descriptor.

The file descriptor is associated with a context. The context is one of the enabler of asynchronous transfers. So, by closing the file descriptor, the programmer must keep in mind that the currently occurring asynchronous transfer will be cancelled.

## Parameters

in	<i>dsc</i>	Device descriptor
----	------------	-------------------

## Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong; the the *dsc* pointer should be checked; check also the kernel log ("dmesg")

References `a4l_sys_close()`, and `a4l_descriptor::fd`.

6.68.2.2 `a4l_fill_desc()`

```
int a4l_fill_desc (
    a4l_desc_t * dsc )
```

Fill the descriptor with subdevices, channels and ranges data.

## Parameters

in	<i>dsc</i>	Device descriptor partly filled by <code>a4l_open()</code> .
----	------------	--

## Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong; the the *dsc* pointer should be checked; check also the kernel log ("dmesg")
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -ENODEV is returned if the descriptor is incoherent (the device may be unattached)

References `a4l_descriptor::fd`, and `a4l_descriptor::magic`.

6.68.2.3 `a4l_get_chinfo()`

```
int a4l_get_chinfo (
    a4l_desc_t * dsc,
    unsigned int subd,
    unsigned int chan,
    a4l_chinfo_t ** info )
```

Get an information structure on a specified channel.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> and <a href="#">a4l_fill_desc()</a>
in	<i>subd</i>	Subdevice index
in	<i>chan</i>	Channel index
out	<i>info</i>	Channel information structure

## Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong; *subd*, *chan* and the *dsc* pointer should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_get\\_chinfo\(\)](#)

References [a4l\\_descriptor::magic](#).

## 6.68.2.4 a4l\_get\_rnginfo()

```
int a4l_get_rnginfo (
    a4l_desc_t * dsc,
    unsigned int subd,
    unsigned int chan,
    unsigned int rng,
    a4l_rnginfo_t ** info )
```

Get an information structure on a specified range.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> and <a href="#">a4l_fill_desc()</a>
in	<i>subd</i>	Subdevice index
in	<i>chan</i>	Channel index
in	<i>rng</i>	Range index
out	<i>info</i>	Range information structure

## Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong; *subd*, *chan*, *rng* and the *dsc* pointer should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_get\\_rnginfo\(\)](#)

References [a4l\\_descriptor::magic](#).

6.68.2.5 `a4l_get_subdinfo()`

```
int a4l_get_subdinfo (
    a4l_desc_t * dsc,
    unsigned int subd,
    a4l_sbinfo_t ** info )
```

Get an information structure on a specified subdevice.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> and <a href="#">a4l_fill_desc()</a>
in	<i>subd</i>	Subdevice index
out	<i>info</i>	Subdevice information structure

## Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong; subd and the dsc pointer should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_get\\_subdinfo\(\)](#).

References [a4l\\_descriptor::magic](#).

Referenced by [a4l\\_sync\\_dio\(\)](#).

6.68.2.6 `a4l_open()`

```
int a4l_open (
    a4l_desc_t * dsc,
    const char * fname )
```

Open an Analogy device and basically fill the descriptor.

## Parameters

out	<i>dsc</i>	Device descriptor
in	<i>fname</i>	Device name

## Returns

0 on success. Otherwise:


- -EINVAL is returned if some argument is missing or wrong; the fname and the dsc pointer should be checked; check also the kernel log ("dmesg")
- -EFAULT is returned if a user <-> kernel transfer went wrong

References [A4L\\_BSC\\_DESC](#), [a4l\\_sys\\_close\(\)](#), [a4l\\_sys\\_desc\(\)](#), [a4l\\_sys\\_open\(\)](#), and [a4l\\_descriptor::fd](#).



## 6.69 Math API

Collaboration diagram for Math API:



### Functions

- int [a4l\\_math\\_polyfit](#) (unsigned r\_dim, double \*r, double orig, const unsigned dim, double \*x, double \*y)  
*Calculate the polynomial fit.*
- void [a4l\\_math\\_mean](#) (double \*pmean, double \*val, unsigned nr)  
*Calculate the arithmetic mean of an array of values.*
- void [a4l\\_math\\_stddev](#) (double \*pstddev, double mean, double \*val, unsigned nr)  
*Calculate the standard deviation of an array of values.*
- void [a4l\\_math\\_stddev\\_of\\_mean](#) (double \*pstddevm, double mean, double \*val, unsigned nr)  
*Calculate the standard deviation of the mean.*

#### 6.69.1 Detailed Description

#### 6.69.2 Function Documentation

##### 6.69.2.1 a4l\_math\_mean()

```

void a4l_math_mean (
    double * pmean,
    double * val,
    unsigned nr )
  
```

Calculate the arithmetic mean of an array of values.

#### Parameters

out	<i>pmean</i>	Pointer to the resulting value
in	<i>val</i>	Array of input values
in	<i>nr</i>	Number of array elements

## 6.69.2.2 a4l\_math\_polyfit()

```
int a4l_math_polyfit (
    unsigned r_dim,
    double * r,
    double orig,
    const unsigned dim,
    double * x,
    double * y )
```

Calculate the polynomial fit.

Parameters

in	<i>r_dim</i>	Number of elements of the resulting polynomial
out	<i>r</i>	Polynomial
in	<i>orig</i>	
in	<i>dim</i>	Number of elements in the polynomials
in	<i>x</i>	Polynomial
in	<i>y</i>	Polynomial

Operation:

We are looking for Res such that  $A \cdot \text{Res} = Y$ , with A the Vandermonde matrix made from the X vector.

Using the least square method, this means finding Res such that:  $A^t \cdot A \cdot \text{Res} = A^t \cdot Y$

If we write  $A = Q \cdot R$  with  $Q^t \cdot Q = 1$ , and R non singular, this can be reduced to:  $R \cdot \text{Res} = Q^t \cdot Y$

mat\_qr() gives us R and  $Q^t \cdot Y$  from A and Y

We can then obtain Res by back substitution using mat\_upper\_triangular\_backsub() with R upper triangular.

## 6.69.2.3 a4l\_math\_stddev()

```
void a4l_math_stddev (
    double * pstddev,
    double mean,
    double * val,
    unsigned nr )
```

Calculate the standard deviation of an array of values.

Parameters

out	<i>pstddev</i>	Pointer to the resulting value
in	<i>mean</i>	Mean value
in	<i>val</i>	Array of input values
in	<i>nr</i>	Number of array elements

Referenced by `a4l_math_stddev_of_mean()`.

#### 6.69.2.4 `a4l_math_stddev_of_mean()`

```
void a4l_math_stddev_of_mean (  
    double * pstddevm,  
    double mean,  
    double * val,  
    unsigned nr )
```

Calculate the standard deviation of the mean.

Parameters

out	<i>pstddevm</i>	Pointer to the resulting value
in	<i>mean</i>	Mean value
in	<i>val</i>	Array of input values
in	<i>nr</i>	Number of array elements

References `a4l_math_stddev()`.

## 6.70 Range / conversion API

Collaboration diagram for Range / conversion API:



### Functions

- int [a4l\\_sizeof\\_chan](#) (a4l\_chinfo\_t \*chan)  
*Get the size in memory of an acquired element.*
- int [a4l\\_sizeof\\_subd](#) (a4l\_sbinfo\_t \*subd)  
*Get the size in memory of a digital acquired element.*
- int [a4l\\_find\\_range](#) (a4l\_desc\_t \*dsc, unsigned int idx\_subd, unsigned int idx\_chan, unsigned long unit, double min, double max, a4l\_rnginfo\_t \*\*rng)  
*Find the most suitable range.*
- int [a4l\\_rawtoul](#) (a4l\_chinfo\_t \*chan, unsigned long \*dst, void \*src, int cnt)  
*Unpack raw data (from the driver) into unsigned long values.*
- int [a4l\\_rawtof](#) (a4l\_chinfo\_t \*chan, a4l\_rnginfo\_t \*rng, float \*dst, void \*src, int cnt)  
*Convert raw data (from the driver) to float-typed samples.*
- int [a4l\\_rawtod](#) (a4l\_chinfo\_t \*chan, a4l\_rnginfo\_t \*rng, double \*dst, void \*src, int cnt)  
*Convert raw data (from the driver) to double-typed samples.*
- int [a4l\\_ultoraw](#) (a4l\_chinfo\_t \*chan, void \*dst, unsigned long \*src, int cnt)  
*Pack unsigned long values into raw data (for the driver)*
- int [a4l\\_ftoraw](#) (a4l\_chinfo\_t \*chan, a4l\_rnginfo\_t \*rng, void \*dst, float \*src, int cnt)  
*Convert float-typed samples to raw data (for the driver)*
- int [a4l\\_dtoraw](#) (a4l\_chinfo\_t \*chan, a4l\_rnginfo\_t \*rng, void \*dst, double \*src, int cnt)  
*Convert double-typed samples to raw data (for the driver)*

### 6.70.1 Detailed Description

### 6.70.2 Function Documentation

#### 6.70.2.1 a4l\_dtoraw()

```

int a4l_dtoraw (
    a4l_chinfo_t * chan,
    a4l_rnginfo_t * rng,
    void * dst,
    double * src,
    int cnt )
  
```

Convert double-typed samples to raw data (for the driver)

## Parameters

in	<i>chan</i>	Channel descriptor
in	<i>rng</i>	Range descriptor
out	<i>dst</i>	Ouput buffer
in	<i>src</i>	Input buffer
in	<i>cnt</i>	Count of conversion to perform

## Returns

the count of conversion performed, otherwise a negative error code:

- -EINVAL is returned if some argument is missing or wrong; *chan*, *rng* and the pointers should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_dtoraw\(\)](#)

## 6.70.2.2 a4l\_find\_range()

```
int a4l_find_range (
    a4l_desc_t * dsc,
    unsigned int idx_subd,
    unsigned int idx_chan,
    unsigned long unit,
    double min,
    double max,
    a4l_rnginfo_t ** rng )
```

Find the most suitable range.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> and <a href="#">a4l_fill_desc()</a>
in	<i>idx_subd</i>	Index of the concerned subdevice
in	<i>idx_chan</i>	Index of the concerned channel
in	<i>unit</i>	Unit type used in the range
in	<i>min</i>	Minimal limit value
in	<i>max</i>	Maximal limit value
out	<i>rng</i>	Found range

## Returns

The index of the most suitable range on success. Otherwise:

- -ENOENT is returned if a suitable range is not found.
- -EINVAL is returned if some argument is missing or wrong; *idx\_subd*, *idx\_chan* and the *dsc* pointer should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_find\\_range\(\)](#)

References `a4l_descriptor::magic`.

### 6.70.2.3 `a4l_ftoraw()`

```
int a4l_ftoraw (
    a4l_chinfo_t * chan,
    a4l_rnginfo_t * rng,
    void * dst,
    float * src,
    int cnt )
```

Convert float-typed samples to raw data (for the driver)

Parameters

in	<i>chan</i>	Channel descriptor
in	<i>rng</i>	Range descriptor
out	<i>dst</i>	Ouput buffer
in	<i>src</i>	Input buffer
in	<i>cnt</i>	Count of conversion to perform

Returns

the count of conversion performed, otherwise a negative error code:

- `-EINVAL` is returned if some argument is missing or wrong; `chan`, `rng` and the pointers should be checked; check also the kernel log ("`dmesg`"); WARNING: `a4l_fill_desc()` should be called before using `a4l_ftoraw()`

### 6.70.2.4 `a4l_rawtod()`

```
int a4l_rawtod (
    a4l_chinfo_t * chan,
    a4l_rnginfo_t * rng,
    double * dst,
    void * src,
    int cnt )
```

Convert raw data (from the driver) to double-typed samples.

Parameters

in	<i>chan</i>	Channel descriptor
in	<i>rng</i>	Range descriptor
out	<i>dst</i>	Ouput buffer
in	<i>src</i>	Input buffer
in	<i>cnt</i>	Count of conversion to perform

## Returns

the count of conversion performed, otherwise a negative error code:

- -EINVAL is returned if some argument is missing or wrong; *chan*, *rng* and the pointers should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_rawtod\(\)](#)

## 6.70.2.5 a4l\_rawtof()

```
int a4l_rawtof (
    a4l_chinfo_t * chan,
    a4l_rnginfo_t * rng,
    float * dst,
    void * src,
    int cnt )
```

Convert raw data (from the driver) to float-typed samples.

## Parameters

in	<i>chan</i>	Channel descriptor
in	<i>rng</i>	Range descriptor
out	<i>dst</i>	Ouput buffer
in	<i>src</i>	Input buffer
in	<i>cnt</i>	Count of conversion to perform

## Returns

the count of conversion performed, otherwise a negative error code:

- -EINVAL is returned if some argument is missing or wrong; *chan*, *rng* and the pointers should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_rawtod\(\)](#)

## 6.70.2.6 a4l\_rawtoul()

```
int a4l_rawtoul (
    a4l_chinfo_t * chan,
    unsigned long * dst,
    void * src,
    int cnt )
```

Unpack raw data (from the driver) into unsigned long values.

This function takes as input driver-specific data and scatters each element into an entry of an unsigned long table. It is a convenience routine which performs no conversion, just copy.

## Parameters

in	<i>chan</i>	Channel descriptor
out	<i>dst</i>	Ouput buffer
in	<i>src</i>	Input buffer
in	<i>cnt</i>	Count of transfer to copy

## Returns

the count of copy performed, otherwise a negative error code:

- -EINVAL is returned if some argument is missing or wrong; *chan*, *dst* and *src* pointers should be checked; check also the kernel log ("dmesg"); WARNING: [a4l\\_fill\\_desc\(\)](#) should be called before using [a4l\\_ultoraw\(\)](#)

## 6.70.2.7 a4l\_sizeof\_chan()

```
int a4l_sizeof_chan (
    a4l_chinfo_t * chan )
```

Get the size in memory of an acquired element.

According to the board, the channels have various acquisition widths. With values like 8, 16 or 32, there is no problem finding out the size in memory (1, 2, 4); however with widths like 12 or 24, this function might be helpful to guess the size needed in RAM for a single acquired element.

## Parameters

in	<i>chan</i>	Channel descriptor
----	-------------	--------------------

## Returns

the size in memory of an acquired element, otherwise a negative error code:

- -EINVAL is returned if the argument *chan* is NULL

## 6.70.2.8 a4l\_sizeof\_subd()

```
int a4l_sizeof_subd (
    a4l_sbinfo_t * subd )
```

Get the size in memory of a digital acquired element.



This function is only useful for DIO subdevices. Digital subdevices are a specific kind of subdevice on which channels are regarded as bits composing the subdevice's bitfield. During a DIO acquisition, all bits are sampled. Therefore, `a4l_sizeof_chan()` is useless in this case and we have to use `a4l_sizeof_subd()`. With bitfields which sizes are 8, 16 or 32, there is no problem finding out the size in memory (1, 2, 4); however with widths like 12 or 24, this function might be helpful to guess the size needed in RAM for a single acquired element.

Parameters

in	<i>subd</i>	Subdevice descriptor
----	-------------	----------------------

Returns

the size in memory of an acquired element, otherwise a negative error code:

- -EINVAL is returned if the argument *chan* is NULL or if the subdevice is not a digital subdevice

References A4L\_SUBD\_DI, A4L\_SUBD\_DIO, A4L\_SUBD\_DO, and A4L\_SUBD\_TYPES.

Referenced by `a4l_sync_dio()`.

#### 6.70.2.9 `a4l_ultoraw()`

```
int a4l_ultoraw (
    a4l_chinfo_t * chan,
    void * dst,
    unsigned long * src,
    int cnt )
```

Pack unsigned long values into raw data (for the driver)

This function takes as input a table of unsigned long values and gather them according to the channel width. It is a convenience routine which performs no conversion, just formatting.

Parameters

in	<i>chan</i>	Channel descriptor
out	<i>dst</i>	Ouput buffer
in	<i>src</i>	Input buffer
in	<i>cnt</i>	Count of transfer to copy


Returns

the count of copy performed, otherwise a negative error code:

- -EINVAL is returned if some argument is missing or wrong; *chan*, *dst* and *src* pointers should be checked; check also the kernel log ("dmesg"); WARNING: `a4l_fill_desc()` should be called before using `a4l_ultoraw()`

## 6.71 Level 1 API

Collaboration diagram for Level 1 API:




### Modules

- [Asynchronous acquisition API](#)
- [Descriptor API](#)
- [Synchronous acquisition API](#)

#### 6.71.1 Detailed Description

## 6.72 Synchronous acquisition API

Collaboration diagram for Synchronous acquisition API:



### Data Structures

- struct [a4l\\_instruction](#)  
*Structure describing the synchronous instruction.*
- struct [a4l\\_instruction\\_list](#)  
*Structure describing the list of synchronous instructions.*

### Macros

- #define [A4L\\_INSN\\_WAIT\\_MAX](#) 100000  
*Maximal wait duration.*

### Functions

- int [a4l\\_snd\\_insnlist](#) ([a4l\\_desc\\_t](#) \*dsc, [a4l\\_insnlst\\_t](#) \*arg)  
*Perform a list of synchronous acquisition misc operations.*
- int [a4l\\_snd\\_insn](#) ([a4l\\_desc\\_t](#) \*dsc, [a4l\\_insn\\_t](#) \*arg)  
*Perform a synchronous acquisition misc operation.*

### Instruction type

Flags to define the type of instruction

- #define [A4L\\_INSN\\_READ](#) (0 | [A4L\\_INSN\\_MASK\\_READ](#))  
*Read instruction.*
- #define [A4L\\_INSN\\_WRITE](#) (1 | [A4L\\_INSN\\_MASK\\_WRITE](#))  
*Write instruction.*
- #define [A4L\\_INSN\\_BITS](#)  
*"Bits" instruction*
- #define [A4L\\_INSN\\_CONFIG](#)  
*Configuration instruction.*
- #define [A4L\\_INSN\\_GTOD](#)  
*Get time instruction.*
- #define [A4L\\_INSN\\_WAIT](#)  
*Wait instruction.*
- #define [A4L\\_INSN\\_INTTRIG](#)  
*Trigger instruction (to start asynchronous acquisition)*

## Configuration instruction type

Values to define the type of configuration instruction

- `#define A4L_INSN_CONFIG_DIO_INPUT 0`
- `#define A4L_INSN_CONFIG_DIO_OUTPUT 1`
- `#define A4L_INSN_CONFIG_DIO_OPENDRAIN 2`
- `#define A4L_INSN_CONFIG_ANALOG_TRIG 16`
- `#define A4L_INSN_CONFIG_ALT_SOURCE 20`
- `#define A4L_INSN_CONFIG_DIGITAL_TRIG 21`
- `#define A4L_INSN_CONFIG_BLOCK_SIZE 22`
- `#define A4L_INSN_CONFIG_TIMER_1 23`
- `#define A4L_INSN_CONFIG_FILTER 24`
- `#define A4L_INSN_CONFIG_CHANGE_NOTIFY 25`
- `#define A4L_INSN_CONFIG_SERIAL_CLOCK 26`
- `#define A4L_INSN_CONFIG_BIDIRECTIONAL_DATA 27`
- `#define A4L_INSN_CONFIG_DIO_QUERY 28`
- `#define A4L_INSN_CONFIG_PWM_OUTPUT 29`
- `#define A4L_INSN_CONFIG_GET_PWM_OUTPUT 30`
- `#define A4L_INSN_CONFIG_ARM 31`
- `#define A4L_INSN_CONFIG_DISARM 32`
- `#define A4L_INSN_CONFIG_GET_COUNTER_STATUS 33`
- `#define A4L_INSN_CONFIG_RESET 34`
- `#define A4L_INSN_CONFIG_GPCT_SINGLE_PULSE_GENERATOR 1001 /* Use CTR as single pulsegenerator */`
- `#define A4L_INSN_CONFIG_GPCT_PULSE_TRAIN_GENERATOR 1002 /* Use CTR as pulse-traingenerator */`
- `#define A4L_INSN_CONFIG_GPCT_QUADRATURE_ENCODER 1003 /* Use the counter as encoder */`
- `#define A4L_INSN_CONFIG_SET_GATE_SRC 2001 /* Set gate source */`
- `#define A4L_INSN_CONFIG_GET_GATE_SRC 2002 /* Get gate source */`
- `#define A4L_INSN_CONFIG_SET_CLOCK_SRC 2003 /* Set master clock source */`
- `#define A4L_INSN_CONFIG_GET_CLOCK_SRC 2004 /* Get master clock source */`
- `#define A4L_INSN_CONFIG_SET_OTHER_SRC 2005 /* Set other source */`
- `#define A4L_INSN_CONFIG_SET_COUNTER_MODE 4097`
- `#define A4L_INSN_CONFIG_SET_ROUTING 4099`
- `#define A4L_INSN_CONFIG_GET_ROUTING 4109`

## Counter status bits

Status bits for `INSN_CONFIG_GET_COUNTER_STATUS`

- `#define A4L_COUNTER_ARMED 0x1`
- `#define A4L_COUNTER_COUNTING 0x2`
- `#define A4L_COUNTER_TERMINAL_COUNT 0x4`

## IO direction

Values to define the IO polarity

- `#define A4L_INPUT 0`
- `#define A4L_OUTPUT 1`
- `#define A4L_OPENDRAIN 2`

## Events types

Values to define the Analogy events. They might used to send some specific events through the instruction interface.

- `#define A4L_EV_START 0x00040000`
- `#define A4L_EV_SCAN_BEGIN 0x00080000`
- `#define A4L_EV_CONVERT 0x00100000`
- `#define A4L_EV_SCAN_END 0x00200000`
- `#define A4L_EV_STOP 0x00400000`

### 6.72.1 Detailed Description

### 6.72.2 Function Documentation

#### 6.72.2.1 `a4l_snd_insn()`

```
int a4l_snd_insn (
    a4l_desc_t * dsc,
    a4l_insn_t * arg )
```

Perform a synchronous acquisition misc operation.

The function `a4l_snd_insn()` triggers a synchronous acquisition.

Parameters

in	<i>dsc</i>	Device descriptor filled by <code>a4l_open()</code> (and optionally <code>a4l_fill_desc()</code> )
in	<i>arg</i>	Instruction structure

Returns

0 on success. Otherwise:

- `-EINVAL` is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- `-EFAULT` is returned if a user <-> kernel transfer went wrong
- `-ENOMEM` is returned if the system is out of memory

References `a4l_descriptor::fd`.

Referenced by `a4l_sync_dio()`.

## 6.72.2.2 a4l\_snd\_insnlist()

```
int a4l_snd_insnlist (
    a4l_desc_t * dsc,
    a4l_insnlst_t * arg )
```

Perform a list of synchronous acquisition misc operations.

The function [a4l\\_snd\\_insnlist\(\)](#) is able to send many synchronous instructions on a various set of sub-devices, channels, etc.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
in	<i>arg</i>	Instructions list structure

## Returns

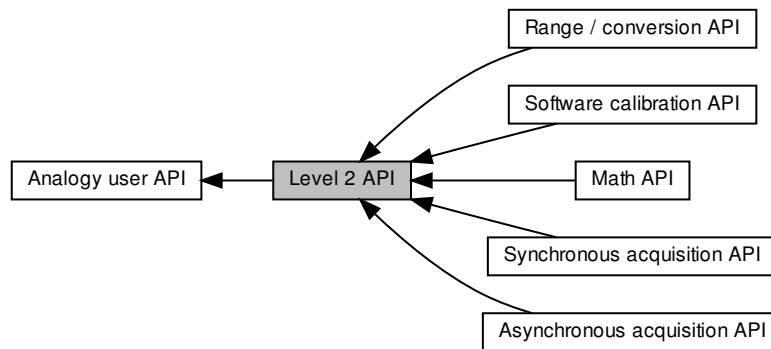
0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -ENOMEM is returned if the system is out of memory

References [a4l\\_descriptor::fd](#).

## 6.73 Level 2 API

Collaboration diagram for Level 2 API:




### Modules

- [Asynchronous acquisition API](#)
- [Software calibration API](#)
- [Math API](#)
- [Range / conversion API](#)
- [Synchronous acquisition API](#)

### 6.73.1 Detailed Description

## 6.74 Synchronous acquisition API

Collaboration diagram for Synchronous acquisition API:



### Functions

- int [a4l\\_sync\\_write](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned int chan\_desc, unsigned int ns\_delay, void \*buf, size\_t nbyte)  
*Perform a synchronous acquisition write operation.*
- int [a4l\\_sync\\_read](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned int chan\_desc, unsigned int ns\_delay, void \*buf, size\_t nbyte)  
*Perform a synchronous acquisition read operation.*
- int [a4l\\_sync\\_dio](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, void \*mask, void \*buf)  
*Perform a synchronous acquisition digital acquisition.*
- int [a4l\\_config\\_subd](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned int type,...)  
*Configure a subdevice.*

#### 6.74.1 Detailed Description

#### 6.74.2 Function Documentation

##### 6.74.2.1 a4l\_config\_subd()

```
int a4l_config_subd (
    a4l\_desc\_t * dsc,
    unsigned int idx_subd,
    unsigned int type,
    ... )
```

Configure a subdevice.

[a4l\\_config\\_subd\(\)](#) takes a variable count of arguments. According to the configuration type, some additional argument is necessary:

- A4L\_INSN\_CONFIG\_DIO\_INPUT: the channel index (unsigned int)
- A4L\_INSN\_CONFIG\_DIO\_OUTPUT: the channel index (unsigned int)
- A4L\_INSN\_CONFIG\_DIO\_QUERY: the returned DIO polarity (unsigned int \*)



## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
in	<i>idx_subd</i>	Index of the concerned subdevice
in	<i>type</i>	Configuration parameter

## Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -ENOSYS is returned if the configuration parameter is not supported

References `A4L_INSN_CONFIG`, and `a4l_instruction::type`.

6.74.2.2 `a4l_sync_dio()`

```
int a4l_sync_dio (
    a4l_desc_t * dsc,
    unsigned int idx_subd,
    void * mask,
    void * buf )
```

Perform a synchronous acquisition digital acquisition.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
in	<i>idx_subd</i>	Index of the concerned subdevice
in	<i>mask</i>	Write mask which indicates which bit(s) must be modified
in, out	<i>buf</i>	Input / output buffer

## Returns

Number of bytes read, otherwise negative error code:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -ENOMEM is returned if the system is out of memory
- -ENOSYS is returned if the driver does not provide any handler "instruction bits"

References `a4l_get_subdinfo()`, `A4L_INSN_BITS`, `a4l_sizeof_subd()`, `a4l_snd_insn()`, `a4l_instruction->data_size`, and `a4l_instruction::type`.

## 6.74.2.3 a4l\_sync\_read()

```
int a4l_sync_read (
    a4l_desc_t * dsc,
    unsigned int idx_subd,
    unsigned int chan_desc,
    unsigned int ns_delay,
    void * buf,
    size_t nbyte )
```

Perform a synchronous acquisition read operation.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
in	<i>idx_subd</i>	Index of the concerned subdevice
in	<i>chan_desc</i>	Channel descriptor (channel, range and reference)
in	<i>ns_delay</i>	Optional delay (in nanoseconds) to wait between the setting of the input channel and sample(s) acquisition(s).
in	<i>buf</i>	Input buffer
in	<i>nbyte</i>	Number of bytes to read

## Returns

Number of bytes read, otherwise negative error code:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -ENOMEM is returned if the system is out of memory

References A4L\_INSN\_READ, A4L\_INSN\_WAIT, and a4l\_instruction::type.

## 6.74.2.4 a4l\_sync\_write()

```
int a4l_sync_write (
    a4l_desc_t * dsc,
    unsigned int idx_subd,
    unsigned int chan_desc,
    unsigned int ns_delay,
    void * buf,
    size_t nbyte )
```

Perform a synchronous acquisition write operation.

## Parameters

in	<i>dsc</i>	Device descriptor filled by <a href="#">a4l_open()</a> (and optionally <a href="#">a4l_fill_desc()</a> )
in	<i>idx_subd</i>	Index of the concerned subdevice
in	<i>chan_desc</i>	Channel descriptor (channel, range and reference) <span style="float: right;">Generated by Doxygen</span>
in	<i>ns_delay</i>	Optional delay (in nanoseconds) to wait between the setting of the input channel and sample(s) acquisition(s).
in	<i>buf</i>	Output buffer

## Returns


Number of bytes written, otherwise negative error code:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -ENOMEM is returned if the system is out of memory

References A4L\_INSN\_WAIT, A4L\_INSN\_WRITE, and a4l\_instruction::type.

## 6.75 Analogy user API

Collaboration diagram for Analogy user API:



### Modules


- [Level 1 API](#)
- [Level 2 API](#)
- [Level 0 API](#)

### 6.75.1 Detailed Description

This is the API interface of Analogy library

## 6.76 Level 0 API

Collaboration diagram for Level 0 API:



### Modules

- [Descriptor Syscall API](#)
- [Basic Syscall API](#)
- [Attach / detach Syscall API](#)

#### 6.76.1 Detailed Description

System call interface to core Analogy services

This interface should not be used directly by applications.

## 6.77 Basic Syscall API

Collaboration diagram for Basic Syscall API:



### Functions

- int [a4l\\_sys\\_open](#) (const char \*fname)  
*Open an Analogy device.*
- int [a4l\\_sys\\_close](#) (int fd)  
*Close an Analogy device.*
- int [a4l\\_sys\\_read](#) (int fd, void \*buf, size\_t nbyte)  
*Read from an Analogy device.*
- int [a4l\\_sys\\_write](#) (int fd, void \*buf, size\_t nbyte)  
*Write to an Analogy device.*

### 6.77.1 Detailed Description

### 6.77.2 Function Documentation

#### 6.77.2.1 [a4l\\_sys\\_close\(\)](#)

```
int a4l_sys_close (  
    int fd )
```

Close an Analogy device.

Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">a4l_sys_open()</a>
----	-----------	---

Returns

0 on success, otherwise a negative error code.

Referenced by [a4l\\_close\(\)](#), and [a4l\\_open\(\)](#).

## 6.77.2.2 a4l\_sys\_open()

```
int a4l_sys_open (
    const char * fname )
```

Open an Analogy device.

Parameters

in	<i>fname</i>	Device name
----	--------------	-------------

Returns

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

Referenced by a4l\_open().

## 6.77.2.3 a4l\_sys\_read()

```
int a4l_sys_read (
    int fd,
    void * buf,
    size_t nbyte )
```

Read from an Analogy device.

The function a4l\_read() is only useful for acquisition configured through an Analogy command.

Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">a4l_sys_open()</a>
out	<i>buf</i>	Input buffer
in	<i>nbyte</i>	Number of bytes to read

Returns

Number of bytes read. Otherwise:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -ENOENT is returned if the device's reading subdevice is idle (no command was sent)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -EINTR is returned if calling task has been unblocked by a signal

## 6.77.2.4 a4l\_sys\_write()

```
int a4l_sys_write (
    int fd,
    void * buf,
    size_t nbyte )
```

Write to an Analogy device.

The function `a4l_write()` is only useful for acquisition configured through an Analogy command.

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">a4l_sys_open()</a>
in	<i>buf</i>	Output buffer
in	<i>nbyte</i>	Number of bytes to write

## Returns

Number of bytes written. Otherwise:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -ENOENT is returned if the device's writing subdevice is idle (no command was sent)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -EINTR is returned if calling task has been unblocked by a signal



## 6.78 Attach / detach Syscall API

Collaboration diagram for Attach / detach Syscall API:



### Functions

- int [a4l\\_sys\\_attach](#) (int fd, a4l\_lnkdesc\_t \*arg)  
*Attach an Analogy device to a driver.*
- int [a4l\\_sys\\_detach](#) (int fd)  
*Detach an Analogy device from a driver.*
- int [a4l\\_sys\\_bufcfg](#) (int fd, unsigned int idx\_subd, unsigned long size)  
*Configure the buffer size.*

### 6.78.1 Detailed Description

### 6.78.2 Function Documentation

#### 6.78.2.1 a4l\_sys\_attach()

```
int a4l_sys_attach (
    int fd,
    a4l_lnkdesc_t * arg )
```

Attach an Analogy device to a driver.

#### Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">a4l_sys_open()</a>
in	<i>arg</i>	Link descriptor argument

#### Returns

0 on success. Otherwise:

- -ENOMEM is returned if the system is out of memory

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -ENODEV is returned in case of internal error (Please, type "dmesg" for more info)
- -ENXIO is returned in case of internal error (Please, type "dmesg" for more info)

### 6.78.2.2 a4l\_sys\_bufcfg()

```
int a4l_sys_bufcfg (
    int fd,
    unsigned int idx_subd,
    unsigned long size )
```

Configure the buffer size.

This function can configure the buffer size of the file descriptor currently in use. If the subdevice index is set to A4L\_BUF\_DEFMAGIC, it can also define the default buffser size at open time.

Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">a4l_sys_open()</a>
in	<i>idx_subd</i>	Index of the concerned subdevice
in	<i>size</i>	Buffer size to be set

Returns

0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EPERM is returned if the function is called in an RT context or if the buffer to resize is mapped in user-space (Please, type "dmesg" for more info)
- -EFAULT is returned if a user <-> kernel transfer went wrong
- -EBUSY is returned if the selected subdevice is already processing an asynchronous operation
- -ENOMEM is returned if the system is out of memory

Referenced by [a4l\\_set\\_bufsize\(\)](#).

### 6.78.2.3 a4l\_sys\_detach()

```
int a4l_sys_detach (
    int fd )
```

Detach an Analogy device from a driver.

## Parameters

in	<i>fd</i>	File descriptor as returned by <a href="#">a4l_sys_open()</a>
----	-----------	---

## Returns


0 on success. Otherwise:

- -EINVAL is returned if some argument is missing or wrong (Please, type "dmesg" for more info)
- -EBUSY is returned if the device to be detached is in use
- -EPERM is returned if the devive to be detached still has some buffer mapped in user-space
- -ENODEV is returned in case of internal error (Please, type "dmesg" for more info)
- -ENXIO is returned in case of internal error (Please, type "dmesg" for more info)

## 6.79 Alarm services

General-purpose watchdog timers.

Collaboration diagram for Alarm services:



### Data Structures

- struct [RT\\_ALARM\\_INFO](#)  
*Alarm status descriptor.*

### Functions

- int [rt\\_alarm\\_start](#) (RT\_ALARM \*alarm, RTIME value, RTIME interval)  
*Start an alarm.*
- int [rt\\_alarm\\_stop](#) (RT\_ALARM \*alarm)  
*Stop an alarm.*
- int [rt\\_alarm\\_inquire](#) (RT\_ALARM \*alarm, [RT\\_ALARM\\_INFO](#) \*info)  
*Query alarm status.*
- int [rt\\_alarm\\_create](#) (RT\_ALARM \*alarm, const char \*name, void(\*handler)(void \*arg), void \*arg)  
*Create an alarm object.*
- int [rt\\_alarm\\_delete](#) (RT\_ALARM \*alarm)  
*Delete an alarm.*

#### 6.79.1 Detailed Description

General-purpose watchdog timers.

Alarms are general-purpose watchdog timers. Alchemy tasks may create any number of alarms and use them to run a user-defined handler, after a specified initial delay has elapsed. Alarms can be either one shot or periodic; in the latter case, the real-time system automatically reprograms the alarm for the next shot according to a user-defined interval value.

## 6.79.2 Function Documentation

### 6.79.2.1 `rt_alarm_create()`

```
int rt_alarm_create (
    RT_ALARM * alarm,
    const char * name,
    void(*)(void *arg) handler,
    void * arg )
```

Create an alarm object.

This routine creates an object triggering an alarm routine at a specified time in the future. Alarms can be periodic or oneshot, depending on the reload interval value passed to [rt\\_alarm\\_start\(\)](#).

#### Parameters

<i>alarm</i>	The address of an alarm descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the alarm. When non-NULL and non-empty, a copy of this string is used for indexing the created alarm into the object registry.
<i>handler</i>	The address of the routine to call when the alarm expires. This routine is passed the <i>arg</i> value.
<i>arg</i>	A user-defined opaque argument passed to the <i>handler</i> .

#### Returns

Zero is returned upon success. Otherwise:

- -ENOMEM is returned if the system fails to get memory from the local pool in order to create the alarm.
- -EEXIST is returned if the *name* is conflicting with an already registered alarm.
- -EPERM is returned if this service was called from an asynchronous context.

#### Tags

[mode-unrestricted](#), [switch-secondary](#)

#### Note

Alarms are process-private objects and thus cannot be shared by multiple processes, even if they belong to the same Xenomai session.

6.79.2.2 `rt_alarm_delete()`

```
int rt_alarm_delete (
    RT_ALARM * alarm )
```

Delete an alarm.

This routine deletes an alarm object previously created by a call to [rt\\_alarm\\_create\(\)](#).

Parameters

<i>alarm</i>	The alarm descriptor.
--------------	-----------------------

Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *alarm* is not a valid alarm descriptor.
- -EPERM is returned if this service was called from an asynchronous context.

Tags

[mode-unrestricted](#), [switch-secondary](#)

6.79.2.3 `rt_alarm_inquire()`

```
int rt_alarm_inquire (
    RT_ALARM * alarm,
    RT_ALARM_INFO * info )
```

Query alarm status.

This routine returns the status information about the specified *alarm*.

Parameters

<i>alarm</i>	The alarm descriptor.
<i>info</i>	A pointer to the <a href="#">return</a> buffer" to copy the information to.

Returns

Zero is returned and status information is written to the structure pointed at by *info* upon success. Otherwise:

- -EINVAL is returned if *alarm* is not a valid alarm descriptor.

Tags

[unrestricted](#), [switch-primary](#)

#### 6.79.2.4 `rt_alarm_start()`

```
int rt_alarm_start (
    RT_ALARM * alarm,
    RTIME value,
    RTIME interval )
```

Start an alarm.

This routine programs the trigger date of an alarm object. An alarm can be either periodic or oneshot, depending on the *interval* value.

Alarm handlers are always called on behalf of Xenomai's internal timer event routine. Therefore, Xenomai routines which can be called from such handlers are restricted to the set of services available on behalf of an asynchronous context.

This service overrides any previous setup of the expiry date and reload interval for the alarm.

Parameters

<i>alarm</i>	The alarm descriptor.
<i>value</i>	The relative date of the first expiry, expressed in clock ticks (see note).
<i>interval</i>	The reload value of the alarm. It is a periodic interval value to be used for reprogramming the next alarm shot, expressed in clock ticks (see note). If <i>interval</i> is equal to <code>TM_INFINITE</code> , the alarm will not be reloaded after it has expired.

Returns

Zero is returned upon success. Otherwise:

- `-EINVAL` is returned if *alarm* is not a valid alarm descriptor.
- `-EPERM` is returned if this service was called from an invalid context.

Tags

[xthread-only](#), [switch-primary](#)

Note

Each of the initial *value* and *interval* is interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.79.2.5 `rt_alarm_stop()`

```
int rt_alarm_stop (
    RT_ALARM * alarm )
```

Stop an alarm.

This routine disables an alarm object, preventing any further expiry until it is re-enabled via [rt\\_alarm\\_start\(\)](#).

Parameters

<i>alarm</i>	The alarm descriptor.
--------------	-----------------------

Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *alarm* is not a valid alarm descriptor.

Tags


[unrestricted](#), [switch-primary](#)



## 6.80 Buffer services

Lightweight FIFO IPC mechanism.

Collaboration diagram for Buffer services:



### Data Structures

- struct `RT_BUFFER_INFO`  
*Buffer status descriptor.*

### Macros

- `#define B_PRIO 0x1` /\* Pend by task priority order. \*/  
*Creation flags.*

### Functions

- int `rt_buffer_create` (RT\_BUFFER \*bf, const char \*name, size\_t bufsz, int mode)  
*Create an IPC buffer.*
- int `rt_buffer_delete` (RT\_BUFFER \*bf)  
*Delete an IPC buffer.*
- ssize\_t `rt_buffer_write_timed` (RT\_BUFFER \*bf, const void \*ptr, size\_t size, const struct timespec \*abs\_timeout)  
*Write to an IPC buffer.*
- static ssize\_t `rt_buffer_write_until` (RT\_BUFFER \*bf, const void \*ptr, size\_t size, RTIME timeout)  
*Write to an IPC buffer (with absolute scalar timeout).*
- static ssize\_t `rt_buffer_write` (RT\_BUFFER \*bf, const void \*ptr, size\_t size, RTIME timeout)  
*Write to an IPC buffer (with relative scalar timeout).*
- ssize\_t `rt_buffer_read_timed` (RT\_BUFFER \*bf, void \*ptr, size\_t size, const struct timespec \*abs\_timeout)  
*Read from an IPC buffer.*
- static ssize\_t `rt_buffer_read_until` (RT\_BUFFER \*bf, void \*ptr, size\_t size, RTIME timeout)

*Read from an IPC buffer (with absolute scalar timeout).*

- static ssize\_t [rt\\_buffer\\_read](#) (RT\_BUFFERER \*bf, void \*ptr, size\_t size, RTIME timeout)

*Read from an IPC buffer (with relative scalar timeout).*

- int [rt\\_buffer\\_clear](#) (RT\_BUFFERER \*bf)

*Clear an IPC buffer.*

- int [rt\\_buffer\\_inquire](#) (RT\_BUFFERER \*bf, RT\_BUFFERER\_INFO \*info)

*Query buffer status.*

- int [rt\\_buffer\\_bind](#) (RT\_BUFFERER \*bf, const char \*name, RTIME timeout)

*Bind to an IPC buffer.*

- int [rt\\_buffer\\_unbind](#) (RT\_BUFFERER \*bf)

*Unbind from an IPC buffer.*

### 6.80.1 Detailed Description

Lightweight FIFO IPC mechanism.

A buffer is a lightweight IPC mechanism, implementing a fast, one-way producer-consumer data path. All messages written are buffered in a single memory area in strict FIFO order, until read either in blocking or non-blocking mode.

Message are always atomically handled on the write side (i.e. no interleave, no short writes), whilst only complete messages are normally returned to the read side. However, short reads may happen under a well-defined situation (see note in [rt\\_buffer\\_read\(\)](#)), albeit they can be fully avoided by proper use of the buffer.

### 6.80.2 Macro Definition Documentation

#### 6.80.2.1 B\_PRIO

```
#define B_PRIO 0x1 /* Pend by task priority order. */
```

Creation flags.

Referenced by [rt\\_buffer\\_create\(\)](#).

### 6.80.3 Function Documentation

6.80.3.1 `rt_buffer_bind()`

```
int rt_buffer_bind (
    RT_BUFFER * bf,
    const char * name,
    RTIME timeout )
```

Bind to an IPC buffer.

This routine creates a new descriptor to refer to an existing IPC buffer identified by its symbolic name. If the object does not exist on entry, the caller may block until a buffer of the given name is created.

## Parameters

<i>bf</i>	The address of a buffer descriptor filled in by the operation. Contents of this memory is undefined upon failure.
<i>name</i>	A valid NULL-terminated name which identifies the buffer to bind to. This string should match the object name argument passed to <code>rt_buffer_create()</code> .
<i>timeout</i>	The number of clock ticks to wait for the registration to occur (see note). Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the object is registered. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if the object is not registered on entry.

## Returns

Zero is returned upon success. Otherwise:

- `-EINTR` is returned if `rt_task_unblock()` was called for the current task before the retrieval has completed.
- `-EWOULDBLOCK` is returned if *timeout* is equal to `TM_NONBLOCK` and the searched object is not registered on entry.
- `-ETIMEDOUT` is returned if the object cannot be retrieved within the specified amount of time.
- `-EPERM` is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

The *timeout* value is interpreted as a multiple of the Alchemy clock resolution (see `--alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.80.3.2 `rt_buffer_clear()`

```
int rt_buffer_clear (
    RT_BUFFER * bf )
```

Clear an IPC buffer.

This routine empties a buffer from any data.

## Parameters

<i>bf</i>	The buffer descriptor.
-----------	------------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *bf* is not a valid buffer descriptor.

## Tags

[unrestricted](#), [switch-primary](#)

6.80.3.3 `rt_buffer_create()`

```
int rt_buffer_create (
    RT_BUFFER * bf,
    const char * name,
    size_t bufsz,
    int mode )
```

Create an IPC buffer.

This routine creates an IPC object that allows tasks to send and receive data asynchronously via a memory buffer. Data may be of an arbitrary length, albeit this IPC is best suited for small to medium-sized messages, since data always have to be copied to the buffer during transit. Large messages may be more efficiently handled by message queues (RT\_QUEUE).

## Parameters

<i>bf</i>	The address of a buffer descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the buffer. When non-NULL and non-empty, a copy of this string is used for indexing the created buffer into the object registry.
<i>bufsz</i>	The size of the buffer space available to hold data. The required memory is obtained from the main heap.
<i>mode</i>	The buffer creation mode. The following flags can be OR'ed into this bitmask, each of them affecting the new buffer:

- B\_FIFO makes tasks pend in FIFO order for reading data from the buffer.
- B\_PRIO makes tasks pend in priority order for reading data from the buffer.

This parameter also applies to tasks blocked on the buffer's write side (see [rt\\_buffer\\_write\(\)](#)).

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *mode* is invalid or *bufsz* is zero.
- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the buffer.
- -EEXIST is returned if the *name* is conflicting with an already registered buffer.
- -EPERM is returned if this service was called from an invalid context, e.g. interrupt or non-Xenomai thread.

## Tags

[xthread-only](#), [mode-unrestricted](#), [switch-secondary](#)

## Note

Buffers can be shared by multiple processes which belong to the same Xenomai session.

References B\_PRIO.

6.80.3.4 `rt_buffer_delete()`

```
int rt_buffer_delete (
    RT_BUFFER * bf )
```

Delete an IPC buffer.

This routine deletes a buffer object previously created by a call to [rt\\_buffer\\_create\(\)](#).

## Parameters

<i>bf</i>	The buffer descriptor.
-----------	------------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *bf* is not a valid buffer descriptor.
- -EPERM is returned if this service was called from an asynchronous context.

## Tags

[mode-unrestricted](#), [switch-secondary](#)

6.80.3.5 `rt_buffer_inquire()`

```
int rt_buffer_inquire (
    RT_BUFFER * bf,
    RT_BUFFER_INFO * info )
```

Query buffer status.

This routine returns the status information about the specified buffer.

Parameters

<i>bf</i>	The buffer descriptor.
<i>info</i>	A pointer to the <a href="#">return</a> buffer" to copy the information to.

Returns

Zero is returned and status information is written to the structure pointed at by *info* upon success. Otherwise:

- -EINVAL is returned if *bf* is not a valid buffer descriptor.

Tags

[unrestricted](#), [switch-primary](#)

6.80.3.6 `rt_buffer_read()`

```
ssize_t rt_buffer_read (
    RT_BUFFER * bf,
    void * ptr,
    size_t len,
    RTIME timeout ) [inline], [static]
```

Read from an IPC buffer (with relative scalar timeout).

This routine is a variant of [rt\\_buffer\\_read\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

Parameters

<i>bf</i>	The buffer descriptor.
<i>ptr</i>	A pointer to a memory area which will be written upon success with the received data.
<i>len</i>	The length in bytes of the memory area pointed to by <i>ptr</i> .
<i>timeout</i>	A delay expressed in clock ticks. Passing TM_INFINITE causes the caller to block indefinitely until enough data is available. Passing TM_NONBLOCK causes the service to return immediately without blocking in case not enough data is available.

Tags

[xthread-nowait](#), [switch-primary](#)

### 6.80.3.7 `rt_buffer_read_timed()`

```
ssize_t rt_buffer_read_timed (
    RT_BUFFER * bf,
    void * ptr,
    size_t len,
    const struct timespec * abs_timeout )
```

Read from an IPC buffer.

This routine reads the next message from the specified buffer. If no message is available on entry, the caller is allowed to block until enough data is written to the buffer, or a timeout elapses.

Parameters

<i>bf</i>	The buffer descriptor.
<i>ptr</i>	A pointer to a memory area which will be written upon success with the received data.
<i>len</i>	The length in bytes of the memory area pointed to by <i>ptr</i> . Under normal circumstances, <code>rt_buffer_read_timed()</code> only returns entire messages as specified by the <i>len</i> argument, or an error value. However, short reads are allowed when a potential deadlock situation is detected (see note below).
<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for a message to be available from the buffer (see note). Passing NULL causes the caller to block indefinitely until enough data is available. Passing { .tv_sec = 0, .tv_nsec = 0 } causes the service to return immediately without blocking in case not enough data is available.

Returns

The number of bytes read from the buffer is returned upon success. Otherwise:

- -ETIMEDOUT is returned if *abs\_timeout* is reached before a complete message arrives.
- -EWOULDBLOCK is returned if *abs\_timeout* is { .tv\_sec = 0, .tv\_nsec = 0 } and not enough data is immediately available on entry to form a complete message.
- -EINTR is returned if `rt_task_unblock()` was called for the current task before enough data became available to form a complete message.
- -EINVAL is returned if *bf* is not a valid buffer descriptor, or *len* is greater than the actual buffer length.
- -EIDRM is returned if *bf* is deleted while the caller was waiting for data. In such event, *bf* is no more valid upon return of this service.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

## Note

A short read (i.e. fewer bytes returned than requested by *len*) may happen whenever a pathological use of the buffer is encountered. This condition only arises when the system detects that one or more writers are waiting for sending data, while a reader would have to wait for receiving a complete message at the same time. For instance, consider the following sequence, involving a 1024-byte buffer (*bf*) and two threads:

```
writer thread > rt_write_buffer(&bf, ptr, 1, TM_INFINITE); (one byte to read, 1023 bytes available for sending)
writer thread > rt_write_buffer(&bf, ptr, 1024, TM_INFINITE); (writer blocks - no space for another 1024-byte message)
reader thread > rt_read_buffer(&bf, ptr, 1024, TM_INFINITE); (short read - a truncated (1-byte) message is returned)
```

In order to prevent both threads to wait for each other indefinitely, a short read is allowed, which may be completed by a subsequent call to `rt_buffer_read()` or `rt_buffer_read_until()`. If that case arises, thread priorities, buffer and/or message lengths should likely be fixed, in order to eliminate such condition.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.80.3.8 `rt_buffer_read_until()`

```
ssize_t rt_buffer_read_until (
    RT_BUFFER * bf,
    void * ptr,
    size_t len,
    RTIME abs_timeout ) [inline], [static]
```

Read from an IPC buffer (with absolute scalar timeout).

This routine is a variant of `rt_buffer_read_timed()` accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>bf</i>	The buffer descriptor.
<i>ptr</i>	A pointer to a memory area which will be written upon success with the received data.
<i>len</i>	The length in bytes of the memory area pointed to by <i>ptr</i> .
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until enough data is available. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without blocking in case not enough data is available.



Tags

[xthread-nowait](#), [switch-primary](#)

### 6.80.3.9 `rt_buffer_unbind()`

```
int rt_buffer_unbind (
    RT_BUFFER * bf )
```

Unbind from an IPC buffer.

Parameters

<i>bf</i>	The buffer descriptor.
-----------	------------------------

This routine releases a previous binding to an IPC buffer. After this call has returned, the descriptor is no more valid for referencing this object.

Tags

[thread-unrestricted](#)

### 6.80.3.10 `rt_buffer_write()`

```
ssize_t rt_buffer_write (
    RT_BUFFER * bf,
    const void * ptr,
    size_t len,
    RTIME timeout ) [inline], [static]
```

Write to an IPC buffer (with relative scalar timeout).

This routine is a variant of [rt\\_buffer\\_write\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

Parameters

<i>bf</i>	The buffer descriptor.
<i>ptr</i>	The address of the message data to be written to the buffer.
<i>len</i>	The length in bytes of the message data.
<i>timeout</i>	A delay expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until enough buffer space is available. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without blocking in case of buffer space shortage.

Tags

[xthread-nowait](#), [switch-primary](#)

### 6.80.3.11 `rt_buffer_write_timed()`

```
ssize_t rt_buffer_write_timed (
    RT_BUFFER * bf,
    const void * ptr,
    size_t len,
    const struct timespec * abs_timeout )
```

Write to an IPC buffer.

This routine writes a message to the specified buffer. If not enough buffer space is available on entry to hold the message, the caller is allowed to block until enough room is freed, or a timeout elapses, whichever comes first.

Parameters

<i>bf</i>	The buffer descriptor.
<i>ptr</i>	The address of the message data to be written to the buffer.
<i>len</i>	The length in bytes of the message data. Zero is a valid value, in which case the buffer is left untouched, and zero is returned to the caller.
<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for enough buffer space to be available to hold the message (see note). Passing NULL causes the caller to block indefinitely until enough buffer space is available. Passing { .tv_sec = 0, .tv_nsec = 0 } causes the service to return immediately without blocking in case of buffer space shortage.

Returns

The number of bytes written to the buffer is returned upon success. Otherwise:

- -ETIMEDOUT is returned if the absolute *abs\_timeout* date is reached before enough buffer space is available to hold the message.
- -EWOULDBLOCK is returned if *abs\_timeout* is { .tv\_sec = 0, .tv\_nsec = 0 } and no buffer space is immediately available on entry to hold the message.
- -EINTR is returned if `rt_task_unblock()` was called for the current task before enough buffer space became available to hold the message.
- -EINVAL is returned if *bf* is not a valid buffer descriptor, or *len* is greater than the actual buffer length.
- -EIDRM is returned if *bf* is deleted while the caller was waiting for buffer space. In such event, *bf* is no more valid upon return of this service.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

Tags

[xthread-nowait](#), [switch-primary](#)

## Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.80.3.12 `rt_buffer_write_until()`

```
ssize_t rt_buffer_write_until (
    RT_BUFFER * bf,
    const void * ptr,
    size_t len,
    RTIME abs_timeout ) [inline], [static]
```

Write to an IPC buffer (with absolute scalar timeout).

This routine is a variant of [rt\\_buffer\\_write\\_timed\(\)](#) accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>bf</i>	The buffer descriptor.
<i>ptr</i>	The address of the message data to be written to the buffer.
<i>len</i>	The length in bytes of the message data.
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until enough buffer space is available. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without blocking in case of buffer space shortage.


## Tags

[xthread-nowait](#), [switch-primary](#)

## 6.81 Condition variable services

POSIXish condition variable mechanism.

Collaboration diagram for Condition variable services:



### Data Structures

- struct [RT\\_COND\\_INFO](#)  
*Condition variable status descriptor.*

### Functions

- int [rt\\_cond\\_create](#) (RT\_COND \*cond, const char \*name)  
*Create a condition variable.*
- int [rt\\_cond\\_delete](#) (RT\_COND \*cond)  
*Delete a condition variable.*
- int [rt\\_cond\\_signal](#) (RT\_COND \*cond)  
*Signal a condition variable.*
- int [rt\\_cond\\_broadcast](#) (RT\_COND \*cond)  
*Broadcast a condition variable*
- int [rt\\_cond\\_wait\\_timed](#) (RT\_COND \*cond, RT\_MUTEX \*mutex, const struct timespec \*abs\_↔ timeout)  
*Wait on a condition variable.*
- static int [rt\\_cond\\_wait\\_until](#) (RT\_COND \*cond, RT\_MUTEX \*mutex, RTIME timeout)  
*Wait on a condition variable (with absolute scalar timeout).*
- static int [rt\\_cond\\_wait](#) (RT\_COND \*cond, RT\_MUTEX \*mutex, RTIME timeout)  
*Wait on a condition variable (with relative scalar timeout).*
- int [rt\\_cond\\_inquire](#) (RT\_COND \*cond, [RT\\_COND\\_INFO](#) \*info)  
*Query condition variable status.*
- int [rt\\_cond\\_bind](#) (RT\_COND \*cond, const char \*name, RTIME timeout)  
*Bind to a condition variable.*
- int [rt\\_cond\\_unbind](#) (RT\_COND \*cond)  
*Unbind from a condition variable.*

### 6.81.1 Detailed Description

POSIXish condition variable mechanism.

A condition variable is a synchronization mechanism which allows tasks to suspend execution until some predicate on some arbitrary shared data is satisfied.

The basic operations on conditions are: signal the condition (when the predicate becomes true), and wait for the condition, blocking the task execution until another task signals the condition. A condition variable must always be associated with a mutex, to avoid a well-known race condition where a task prepares to wait on a condition variable and another task signals the condition just before the first task actually waits on it.

### 6.81.2 Function Documentation

#### 6.81.2.1 `rt_cond_bind()`

```
int rt_cond_bind (
    RT_COND * cond,
    const char * name,
    RTIME timeout )
```

Bind to a condition variable.

This routine creates a new descriptor to refer to an existing condition variable identified by its symbolic name. If the object not exist on entry, the caller may block until a condition variable of the given name is created.

Parameters

<i>cond</i>	The address of a condition variable descriptor filled in by the operation. Contents of this memory is undefined upon failure.
<i>name</i>	A valid NULL-terminated name which identifies the condition variable to bind to. This string should match the object name argument passed to <code>rt_cond_create()</code> .
<i>timeout</i>	The number of clock ticks to wait for the registration to occur (see note). Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the object is registered. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if the object is not registered on entry.

Returns

Zero is returned upon success. Otherwise:

- `-EINTR` is returned if `rt_task_unblock()` was called for the current task before the retrieval has completed.
- `-EWOULDBLOCK` is returned if *timeout* is equal to `TM_NONBLOCK` and the searched object is not registered on entry.

- -ETIMEDOUT is returned if the object cannot be retrieved within the specified amount of time.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

The *timeout* value is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.81.2.2 `rt_cond_broadcast()`

```
int rt_cond_broadcast (
    RT_COND * cond )
```

## Broadcast a condition variable

All tasks currently waiting on the condition variable are immediately unblocked.

## Parameters

<i>cond</i>	The condition variable descriptor.
-------------	------------------------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *cond* is not a valid condition variable descriptor.

## Tags

[unrestricted](#), [switch-primary](#)

6.81.2.3 `rt_cond_create()`

```
int rt_cond_create (
    RT_COND * cond,
    const char * name )
```

## Create a condition variable.

Create a synchronization object which allows tasks to suspend execution until some predicate on shared data is satisfied.

## Parameters

<i>cond</i>	The address of a condition variable descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the condition variable. When non-NULL and non-empty, a copy of this string is used for indexing the created condition variable into the object registry.

## Returns

Zero is returned upon success. Otherwise:

- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the condition variable.
- -EEXIST is returned if the *name* is conflicting with an already registered condition variable.
- -EPERM is returned if this service was called from an invalid context, e.g. interrupt or non-Xenomai thread.

## Tags

[xthread-only](#), [mode-unrestricted](#), [switch-secondary](#)

## Note

Condition variables can be shared by multiple processes which belong to the same Xenomai session.

## Attention

If the underlying threading library does not support [pthread\\_condattr\\_setclock\(\)](#), timings with Alchemy condition variables will be based on CLOCK\_REALTIME, and may therefore be affected by updates to the system date (e.g. NTP). This typically concerns legacy setups based on the linuxthreads library. In the normal case, timings are based on CLOCK\_MONOTONIC.

6.81.2.4 `rt_cond_delete()`

```
int rt_cond_delete (
    RT_COND * cond )
```

Delete a condition variable.

This routine deletes a condition variable object previously created by a call to [rt\\_cond\\_create\(\)](#).

## Parameters

<i>cond</i>	The condition variable descriptor.
-------------	------------------------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *alarm* is not a valid condition variable descriptor.
- -EPERM is returned if this service was called from an asynchronous context.
- -EBUSY is returned upon an attempt to destroy the object referenced by *cond* while it is referenced (for example, while being used in a `rt_cond_wait()`, `rt_cond_wait_timed()` or `rt_cond_wait_until()` by another task).

## Tags

[mode-unrestricted](#), [switch-secondary](#)

6.81.2.5 `rt_cond_inquire()`

```
int rt_cond_inquire (
    RT_COND * cond,
    RT_COND_INFO * info )
```

Query condition variable status.

This routine returns the status information about the specified condition variable.

## Parameters

<i>cond</i>	The condition variable descriptor.
<i>info</i>	A pointer to the <a href="#">return</a> buffer" to copy the information to.

## Returns

Zero is returned and status information is written to the structure pointed at by *info* upon success. Otherwise:

- -EINVAL is returned if *cond* is not a valid condition variable descriptor.

## Tags

[unrestricted](#), [switch-primary](#)



6.81.2.6 `rt_cond_signal()`

```
int rt_cond_signal (
    RT_COND * cond )
```

Signal a condition variable.

If the condition variable *cond* is pended, this routine immediately unblocks the first waiting task (by queuing priority order).

Parameters

<i>cond</i>	The condition variable descriptor.
-------------	------------------------------------

Returns

Zero is returned upon success. Otherwise:

- `-EINVAL` is returned if *cond* is not a valid condition variable descriptor.

Tags

[unrestricted](#), [switch-primary](#)

6.81.2.7 `rt_cond_unbind()`

```
int rt_cond_unbind (
    RT_COND * cond )
```

Unbind from a condition variable.

Parameters

<i>cond</i>	The condition variable descriptor.
-------------	------------------------------------

This routine releases a previous binding to a condition variable. After this call has returned, the descriptor is no more valid for referencing this object.

Tags

[thread-unrestricted](#)

6.81.2.8 `rt_cond_wait()`

```
int rt_cond_wait (
    RT_COND * cond,
    RT_MUTEX * mutex,
    RTIME timeout ) [inline], [static]
```

Wait on a condition variable (with relative scalar timeout).

This routine is a variant of `rt_cond_wait_timed()` accepting a relative timeout specification expressed as a scalar value.

## Parameters

<i>cond</i>	The condition variable descriptor.
<i>mutex</i>	The address of the mutex serializing the access to the shared data.
<i>timeout</i>	A delay expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely. Passing <code>TM_NONBLOCK</code> causes the caller to return immediately without block.

## Tags

[xthread-only](#), [switch-primary](#)

6.81.2.9 `rt_cond_wait_timed()`

```
int rt_cond_wait_timed (
    RT_COND * cond,
    RT_MUTEX * mutex,
    const struct timespec * abs_timeout )
```

Wait on a condition variable.

This service atomically releases the mutex and blocks the calling task, until the condition variable *cond* is signaled or a timeout occurs, whichever comes first. The mutex is re-acquired before returning from this service.

## Parameters

<i>cond</i>	The condition variable descriptor.
<i>mutex</i>	The address of the mutex serializing the access to the shared data.
<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for the condition variable to be signaled (see note). Passing <code>NULL</code> causes the caller to block indefinitely. Passing <code>{ .tv_sec = 0, .tv_nsec = 0 }</code> causes the caller to return immediately without block.

## Returns

Zero is returned upon success. Otherwise:

- -ETIMEDOUT is returned if *abs\_timeout* is reached before the condition variable is signaled.
- -EWOULDBLOCK is returned if *abs\_timeout* is { .tv\_sec = 0, .tv\_nsec = 0 } .
- -EINTR is returned if `rt_task_unblock()` was called for the current task.
- -EINVAL is returned if *cond* is not a valid condition variable descriptor.
- -EIDRM is returned if *cond* is deleted while the caller was waiting on the condition variable. In such event, *cond* is no more valid upon return of this service.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-only](#), [switch-primary](#)

## Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.81.2.10 `rt_cond_wait_until()`

```
int rt_cond_wait_until (
    RT_COND * cond,
    RT_MUTEX * mutex,
    RTIME abs_timeout ) [inline], [static]
```

Wait on a condition variable (with absolute scalar timeout).

This routine is a variant of `rt_cond_wait_timed()` accepting an *abs\_timeout* specification expressed as a scalar value.

## Parameters

<i>cond</i>	The condition variable descriptor.
<i>mutex</i>	The address of the mutex serializing the access to the shared data.
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely. Passing <code>TM_NONBLOCK</code> causes the caller to return immediately without block.


## Tags

[xthread-only](#), [switch-primary](#)

## 6.82 Event flag group services

Inter-task notification mechanism based on discrete flags.

Collaboration diagram for Event flag group services:



### Data Structures

- struct [RT\\_EVENT\\_INFO](#)  
*Event status descriptor.*

### Macros

- #define [EV\\_PRIO](#) 0x1 /\* Pend by task priority order. \*/  
*Creation flags.*
- #define [EV\\_ANY](#) 0x1 /\* Disjunctive wait. \*/  
*Operation flags.*

### Functions

- int [rt\\_event\\_delete](#) (RT\_EVENT \*event)  
  
*Delete an event flag group.*
- int [rt\\_event\\_wait\\_timed](#) (RT\_EVENT \*event, unsigned int mask, unsigned int \*mask\_r, int mode, const struct timespec \*abs\_timeout)  
  
*Wait for an arbitrary set of events.*
- static int [rt\\_event\\_wait\\_until](#) (RT\_EVENT \*event, unsigned int mask, unsigned int \*mask\_r, int mode, RTIME timeout)  
  
*Wait for an arbitrary set of events (with absolute scalar timeout).*
- static int [rt\\_event\\_wait](#) (RT\_EVENT \*event, unsigned int mask, unsigned int \*mask\_r, int mode, RTIME timeout)  
  
*Wait for an arbitrary set of events (with relative scalar timeout).*
- int [rt\\_event\\_inquire](#) (RT\_EVENT \*event, [RT\\_EVENT\\_INFO](#) \*info)  
  
*Query event flag group status.*
- int [rt\\_event\\_bind](#) (RT\_EVENT \*event, const char \*name, RTIME timeout)

*Bind to an event flag group.*

- int `rt_event_unbind` (RT\_EVENT \*event)

*Unbind from an event flag group.*

- int `rt_event_create` (RT\_EVENT \*event, const char \*name, unsigned int ivalue, int mode)

*Create an event flag group.*

- int `rt_event_signal` (RT\_EVENT \*event, unsigned int mask)

*Signal an event.*

- int `rt_event_clear` (RT\_EVENT \*event, unsigned int mask, unsigned int \*mask\_r)

*Clear event flags.*

### 6.82.1 Detailed Description

Inter-task notification mechanism based on discrete flags.

An event flag group is a synchronization object represented by a long-word structure; every available bit in this word represents a user-defined event flag.

When a bit is set, the associated event is said to have occurred. Xenomai tasks can use this mechanism to signal the occurrence of particular events to other tasks.

Tasks can either wait for events to occur in a conjunctive manner (all awaited events must have occurred to satisfy the wait request), or in a disjunctive way (at least one of the awaited events must have occurred to satisfy the wait request).

### 6.82.2 Macro Definition Documentation

#### 6.82.2.1 EV\_ANY

```
#define EV_ANY 0x1 /* Disjunctive wait. */
```

Operation flags.

#### 6.82.2.2 EV\_PRIO

```
#define EV_PRIO 0x1 /* Pend by task priority order. */
```

Creation flags.

Referenced by `rt_event_create()`.

### 6.82.3 Function Documentation

#### 6.82.3.1 `rt_event_bind()`

```
int rt_event_bind (
    RT_EVENT * event,
    const char * name,
    RTIME timeout )
```

Bind to an event flag group.

This routine creates a new descriptor to refer to an existing event flag group identified by its symbolic name. If the object does not exist on entry, the caller may block until an event flag group of the given name is created.

#### Parameters

<i>event</i>	The address of an event flag group descriptor filled in by the operation. Contents of this memory is undefined upon failure.
<i>name</i>	A valid NULL-terminated name which identifies the event flag group to bind to. This string should match the object name argument passed to <a href="#">rt_event_create()</a> .
<i>timeout</i>	The number of clock ticks to wait for the registration to occur (see note). Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the object is registered. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if the object is not registered on entry.

#### Returns

Zero is returned upon success. Otherwise:

- `-EINTR` is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task before the retrieval has completed.
- `-EWOULDBLOCK` is returned if *timeout* is equal to `TM_NONBLOCK` and the searched object is not registered on entry.
- `-ETIMEDOUT` is returned if the object cannot be retrieved within the specified amount of time.
- `-EPERM` is returned if this service should block, but was not called from a Xenomai thread.

#### Tags

[xthread-nowait](#), [switch-primary](#)

#### Note

The *timeout* value is interpreted as a multiple of the Alchemy clock resolution (see `--alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.82.3.2 `rt_event_clear()`

```
int rt_event_clear (
    RT_EVENT * event,
    unsigned int mask,
    unsigned int * mask_r )
```

Clear event flags.

This routine clears a set of flags from *event*.

## Parameters

<i>event</i>	The event descriptor.
<i>mask</i>	The set of event flags to be cleared.
<i>mask</i> ↔ <i>_r</i>	If non-NULL, <i>mask_r</i> is the address of a memory location which will receive the previous value of the event flag group before the flags are cleared.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *event* is not a valid event flag group descriptor.

## Tags

[unrestricted](#), [switch-primary](#)

6.82.3.3 `rt_event_create()`

```
int rt_event_create (
    RT_EVENT * event,
    const char * name,
    unsigned int ivalue,
    int mode )
```

Create an event flag group.

Event groups provide for task synchronization by allowing a set of flags (or "events") to be waited for and posted atomically. An event group contains a mask of received events; an arbitrary set of event flags can be pended or posted in a single operation.

## Parameters

<i>event</i>	The address of an event descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the event. When non-NULL and non-empty, a copy of this string is used for indexing the created event into the object registry.
<i>ivalue</i>	The initial value of the group's event mask.
<i>mode</i>	The event group creation mode. The following flags can be OR'ed into this bitmask:

- EV\_FIFO makes tasks pend in FIFO order on the event flag group.
- EV\_PRIO makes tasks pend in priority order on the event flag group.

#### Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *mode* is invalid.
- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the event flag group.
- -EEXIST is returned if the *name* is conflicting with an already registered event flag group.
- -EPERM is returned if this service was called from an invalid context, e.g. interrupt or non-Xenomai thread.

#### Tags

[xthread-only](#), [mode-unrestricted](#), [switch-secondary](#)

#### Note

Event flag groups can be shared by multiple processes which belong to the same Xenomai session.

References EV\_PRIO.

#### 6.82.3.4 rt\_event\_delete()

```
int rt_event_delete (
    RT_EVENT * event )
```

Delete an event flag group.

This routine deletes a event flag group previously created by a call to [rt\\_event\\_create\(\)](#).

#### Parameters

<i>event</i>	The event descriptor.
--------------	-----------------------

#### Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *event* is not a valid event flag group descriptor.
- -EPERM is returned if this service was called from an asynchronous context.



Tags

[mode-unrestricted](#), [switch-secondary](#)

### 6.82.3.5 `rt_event_inquire()`

```
int rt_event_inquire (
    RT_EVENT * event,
    RT_EVENT_INFO * info )
```

Query event flag group status.

This routine returns the status information about *event*.

Parameters

<i>event</i>	The event descriptor.
<i>info</i>	A pointer to the <a href="#">return</a> buffer" to copy the information to.

Returns

Zero is returned and status information is written to the structure pointed at by *info* upon success. Otherwise:

- `-EINVAL` is returned if *event* is not a valid event flag group descriptor.

Tags

[unrestricted](#), [switch-primary](#)

### 6.82.3.6 `rt_event_signal()`

```
int rt_event_signal (
    RT_EVENT * event,
    unsigned int mask )
```

Signal an event.

Post a set of flags to *event*. All tasks having their wait request satisfied as a result of this operation are immediately readied.

Parameters

<i>event</i>	The event descriptor.
<i>mask</i>	The set of events to be posted.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *event* is not an event flag group descriptor.

## Tags

[unrestricted](#), [switch-primary](#)

6.82.3.7 `rt_event_unbind()`

```
int rt_event_unbind (
    RT_EVENT * event )
```

Unbind from an event flag group.

## Parameters

<i>event</i>	The event descriptor.
--------------	-----------------------

This routine releases a previous binding to an event flag group. After this call has returned, the descriptor is no more valid for referencing this object.

## Tags

[thread-unrestricted](#)

6.82.3.8 `rt_event_wait()`

```
int rt_event_wait (
    RT_EVENT * event,
    unsigned int mask,
    unsigned int * mask_r,
    int mode,
    RTIME timeout ) [inline], [static]
```

Wait for an arbitrary set of events (with relative scalar timeout).

This routine is a variant of [rt\\_event\\_wait\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

## Parameters

<i>event</i>	The event descriptor.
<i>mask</i>	The set of bits to wait for.
<i>mask</i> ↔ <i>_r</i>	The value of the event mask at the time the task was readied.
<i>mode</i>	The pend mode.
<i>timeout</i>	A delay expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the request is satisfied. Passing <code>TM_NONBLOCK</code> causes the service to return without blocking in case the request cannot be satisfied immediately.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.82.3.9 `rt_event_wait_timed()`

```
int rt_event_wait_timed (
    RT_EVENT * event,
    unsigned int mask,
    unsigned int * mask_r,
    int mode,
    const struct timespec * abs_timeout )
```

Wait for an arbitrary set of events.

Waits for one or more events to be signaled in *event*, or until a timeout elapses.

## Parameters

<i>event</i>	The event descriptor.
<i>mask</i>	The set of bits to wait for. Passing zero causes this service to return immediately with a success value; the current value of the event mask is also copied to <i>mask_r</i> .
<i>mask</i> ↔ <i>_r</i>	The value of the event mask at the time the task was readied.
<i>mode</i>	The pend mode. The following flags can be OR'ed into this bitmask, each of them affecting the operation:

- `EV_ANY` makes the task pend in disjunctive mode (i.e. OR); this means that the request is fulfilled when at least one bit set into *mask* is set in the current event mask.
- `EV_ALL` makes the task pend in conjunctive mode (i.e. AND); this means that the request is fulfilled when at all bits set into *mask* are set in the current event mask.

## Parameters

<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for the request to be satisfied (see note). Passing NULL causes the caller to block indefinitely until the request is satisfied. Passing { .tv_sec = 0, .tv_nsec = 0 } causes the service to return without blocking in case the request cannot be satisfied immediately.
--------------------	--

## Returns

Zero is returned upon success. Otherwise:

- -ETIMEDOUT is returned if *abs\_timeout* is reached before the request is satisfied.
- -EWOULDBLOCK is returned if *abs\_timeout* is { .tv\_sec = 0, .tv\_nsec = 0 } and the requested flags are not set on entry to the call.
- -EINTR is returned if `rt_task_unblock()` was called for the current task before the request is satisfied.
- -EINVAL is returned if *mode* is invalid, or *event* is not a valid event flag group descriptor.
- -EIDRM is returned if *event* is deleted while the caller was sleeping on it. In such a case, *event* is no more valid upon return of this service.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

*abs\_timeout* value is interpreted as a multiple of the Alchemy clock resolution (see `--alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.82.3.10 `rt_event_wait_until()`

```
int rt_event_wait_until (
    RT_EVENT * event,
    unsigned int mask,
    unsigned int * mask_r,
    int mode,
    RTIME abs_timeout ) [inline], [static]
```

Wait for an arbitrary set of events (with absolute scalar timeout).

This routine is a variant of `rt_event_wait_timed()` accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>event</i>	The event descriptor.
<i>mask</i>	The set of bits to wait for.
<i>mask_r</i>	The value of the event mask at the time the task was readied.
<i>mode</i>	The pend mode.
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the request is satisfied. Passing <code>TM_NONBLOCK</code> causes the service to return without blocking in case the request cannot be satisfied immediately.


## Tags

[xthread-nowait](#), [switch-primary](#)

## 6.83 Heap management services

Region of memory dedicated to real-time allocation.

Collaboration diagram for Heap management services:



### Data Structures

- struct [RT\\_HEAP\\_INFO](#)  
*Heap status descriptor.*

### Macros

- `#define H_PRIO 0x1` /\* Pend by task priority order. \*/  
*Creation flags.*

### Functions

- int [rt\\_heap\\_create](#) (RT\_HEAP \*heap, const char \*name, size\_t heapsize, int mode)  
*Create a heap.*
- int [rt\\_heap\\_delete](#) (RT\_HEAP \*heap)  
*Delete a heap.*
- int [rt\\_heap\\_alloc\\_timed](#) (RT\_HEAP \*heap, size\_t size, const struct timespec \*abs\_timeout, void \*\*blockp)  
*Allocate a block from a heap.*
- static int [rt\\_heap\\_alloc\\_until](#) (RT\_HEAP \*heap, size\_t size, RTIME timeout, void \*\*blockp)  
*Allocate a block from a heap (with absolute scalar timeout).*
- static int [rt\\_heap\\_alloc](#) (RT\_HEAP \*heap, size\_t size, RTIME timeout, void \*\*blockp)  
*Allocate a block from a heap (with relative scalar timeout).*
- int [rt\\_heap\\_free](#) (RT\_HEAP \*heap, void \*block)  
*Release a block to a heap.*
- int [rt\\_heap\\_inquire](#) (RT\_HEAP \*heap, RT\_HEAP\_INFO \*info)  
*Query heap status.*
- int [rt\\_heap\\_bind](#) (RT\_HEAP \*heap, const char \*name, RTIME timeout)  
*Bind to a heap.*
- int [rt\\_heap\\_unbind](#) (RT\_HEAP \*heap)  
*Unbind from a heap.*

### 6.83.1 Detailed Description

Region of memory dedicated to real-time allocation.

Heaps are regions of memory used for dynamic memory allocation in a time-bounded fashion. Blocks of memory are allocated and freed in an arbitrary order and the pattern of allocation and size of blocks is not known until run time.

### 6.83.2 Macro Definition Documentation

#### 6.83.2.1 H\_PRIO

```
#define H_PRIO 0x1 /* Pend by task priority order. */
```

Creation flags.

Referenced by `rt_heap_create()`.

### 6.83.3 Function Documentation

#### 6.83.3.1 `rt_heap_alloc()`

```
int rt_heap_alloc (  
    RT_HEAP * heap,  
    size_t size,  
    RTIME timeout,  
    void ** blockp ) [inline], [static]
```

Allocate a block from a heap (with relative scalar timeout).

This routine is a variant of `rt_heap_alloc_timed()` accepting a relative timeout specification expressed as a scalar value. Passing `TM_INFINITE` in `timeout` causes the caller to block indefinitely until a block is available. Passing `TM_NONBLOCK` causes the service to return immediately without blocking in case a block is not available.

Tags

[xthread-nowait](#), [switch-primary](#)

### 6.83.3.2 `rt_heap_alloc_timed()`

```
int rt_heap_alloc_timed (
    RT_HEAP * heap,
    size_t size,
    const struct timespec * abs_timeout,
    void ** blockp )
```

Allocate a block from a heap.

This service allocates a block from a given heap, or returns the address of the single memory segment if `H_SINGLE` was mentioned in the creation mode to [rt\\_heap\\_create\(\)](#). When not enough memory is available on entry to this service, tasks may be blocked until their allocation request can be fulfilled.



## Parameters

<i>heap</i>	The heap descriptor.
<i>size</i>	The requested size (in bytes) of the block. If the heap is managed as a single-block area (H_SINGLE), this value can be either zero, or the same value given to <a href="#">rt_heap_create()</a> . In that case, the same block covering the entire heap space is returned to all callers of this service.
<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for a block of the requested size to be available from the heap (see note). Passing NULL causes the caller to block indefinitely until a block is available. Passing { .tv_sec = 0, .tv_nsec = 0 } causes the service to return immediately without blocking in case a block is not available.
<i>blockp</i>	A pointer to a memory location which will be written upon success with the address of the allocated block, or the start address of the single memory segment. In the former case, the block can be freed using <a href="#">rt_heap_free()</a> .

## Returns

Zero is returned upon success. Otherwise:

- -ETIMEDOUT is returned if *abs\_timeout* is reached before a block is available.
- -EWOULDBLOCK is returned if *abs\_timeout* is equal to { .tv\_sec = 0, .tv\_nsec = 0 } and no block is immediately available on entry to fulfill the allocation request.
- -EINTR is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task before a block became available.
- -EINVAL is returned if *heap* is not a valid heap descriptor, or *heap* is managed as a single-block area (i.e. H\_SINGLE mode) and *size* is non-zero but does not match the original heap size passed to [rt\\_heap\\_create\(\)](#).
- -EIDRM is returned if *heap* is deleted while the caller was waiting for a block. In such event, *heap* is no more valid upon return of this service.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

If shared multi-processing is enabled (i.e. `–enable-pshared` was passed to the configure script), requests for a block size larger than twice the allocation page size are rounded up to the next page size. The allocation page size is currently 512 bytes long (HOBJ\_PAGE\_SIZE), which means that any request larger than 1k will be rounded up to the next 512 byte boundary.

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.83.3.3 `rt_heap_alloc_until()`

```
int rt_heap_alloc_until (
    RT_HEAP * heap,
    size_t size,
    RTIME abs_timeout,
    void ** blockp ) [inline], [static]
```

Allocate a block from a heap (with absolute scalar timeout).

This routine is a variant of `rt_heap_alloc_timed()` accepting an absolute timeout specification expressed as a scalar value. Passing `TM_INFINITE` in *timeout* causes the caller to block indefinitely until a block is available. Passing `TM_NONBLOCK` causes the service to return immediately without blocking in case a block is not available.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.83.3.4 `rt_heap_bind()`

```
int rt_heap_bind (
    RT_HEAP * heap,
    const char * name,
    RTIME timeout )
```

Bind to a heap.

This routine creates a new descriptor to refer to an existing heap identified by its symbolic name. If the object does not exist on entry, the caller may block until a heap of the given name is created.

## Parameters

<i>heap</i>	The address of a heap descriptor filled in by the operation. Contents of this memory is undefined upon failure.
<i>name</i>	A valid NULL-terminated name which identifies the heap to bind to. This string should match the object name argument passed to <code>rt_heap_create()</code> .
<i>timeout</i>	The number of clock ticks to wait for the registration to occur (see note). Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the object is registered. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if the object is not registered on entry.

## Returns

Zero is returned upon success. Otherwise:

- `-EINTR` is returned if `rt_task_unblock()` was called for the current task before the retrieval has completed.

- -EWOULDBLOCK is returned if *timeout* is equal to TM\_NONBLOCK and the searched object is not registered on entry.
- -ETIMEDOUT is returned if the object cannot be retrieved within the specified amount of time.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

The *timeout* value is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.83.3.5 `rt_heap_create()`

```
int rt_heap_create (
    RT_HEAP * heap,
    const char * name,
    size_t heapsz,
    int mode )
```

## Create a heap.

This routine creates a memory heap suitable for time-bounded allocation requests of RAM chunks. When not enough memory is available, tasks may be blocked until their allocation request can be fulfilled.

By default, heaps support allocation of multiple blocks of memory in an arbitrary order. However, it is possible to ask for single-block management by passing the H\_SINGLE flag into the *mode* parameter, in which case the entire memory space managed by the heap is made available as a unique block. In this mode, all allocation requests made through `rt_heap_alloc()` will return the same block address, pointing at the beginning of the heap memory.

## Parameters

<i>heap</i>	The address of a heap descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the heap. When non-NULL and non-empty, a copy of this string is used for indexing the created heap into the object registry.
<i>heapsz</i>	The size (in bytes) of the memory pool, blocks will be claimed and released to. This area is not extensible, so this value must be compatible with the highest memory pressure that could be expected. The valid range is between 1 byte and 2Gb.
<i>mode</i>	The heap creation mode. The following flags can be OR'ed into this bitmask, each of them affecting the new heap:

- H\_FIFO makes tasks pend in FIFO order on the heap when waiting for available blocks.

- H\_PRIO makes tasks pend in priority order on the heap when waiting for available blocks.
- H\_SINGLE causes the entire heap space to be managed as a single memory block.

#### Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *mode* is invalid, or *heapsz* is zero or larger than 2Gb.
- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the heap.
- -EEXIST is returned if the *name* is conflicting with an already registered heap.
- -EPERM is returned if this service was called from an invalid context, e.g. interrupt or non-Xenomai thread.

#### Tags

[xthread-only](#), [mode-unrestricted](#), [switch-secondary](#)

#### Note

Heaps can be shared by multiple processes which belong to the same Xenomai session.

References H\_PRIO.

#### 6.83.3.6 rt\_heap\_delete()

```
int rt_heap_delete (
    RT_HEAP * heap )
```

Delete a heap.

This routine deletes a heap object previously created by a call to [rt\\_heap\\_create\(\)](#), releasing all tasks currently blocked on it.

#### Parameters

<i>heap</i>	The heap descriptor.
-------------	----------------------

#### Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *heap* is not a valid heap descriptor.

- -EPERM is returned if this service was called from an asynchronous context.

Tags

[mode-unrestricted](#), [switch-secondary](#)

### 6.83.3.7 `rt_heap_free()`

```
int rt_heap_free (
    RT_HEAP * heap,
    void * block )
```

Release a block to a heap.

This service should be used to release a block to the heap it belongs to. An attempt to fulfill the request of every task blocked on [rt\\_heap\\_alloc\(\)](#) is made once *block* is returned to the memory pool.

Parameters

<i>heap</i>	The heap descriptor.
<i>block</i>	The address of the block to free.

Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *heap* is not a valid heap descriptor, or *block* is not a valid block previously allocated by the [rt\\_heap\\_alloc\(\)](#) service from *heap*.

Tags

[unrestricted](#), [switch-primary](#)

### 6.83.3.8 `rt_heap_inquire()`

```
int rt_heap_inquire (
    RT_HEAP * heap,
    RT_HEAP_INFO * info )
```

Query heap status.

This routine returns the status information about *heap*.

## Parameters

<i>heap</i>	The heap descriptor.
<i>info</i>	A pointer to the <a href="#">return</a> buffer" to copy the information to.

## Returns

Zero is returned and status information is written to the structure pointed at by *info* upon success. Otherwise:

- -EINVAL is returned if *heap* is not a valid heap descriptor.

## Tags

[unrestricted](#), [switch-primary](#)

6.83.3.9 `rt_heap_unbind()`

```
int rt_heap_unbind (
    RT_HEAP * heap )
```

Unbind from a heap.

## Parameters

<i>heap</i>	The heap descriptor.
-------------	----------------------

This routine releases a previous binding to a heap. After this call has returned, the descriptor is no more valid for referencing this object.

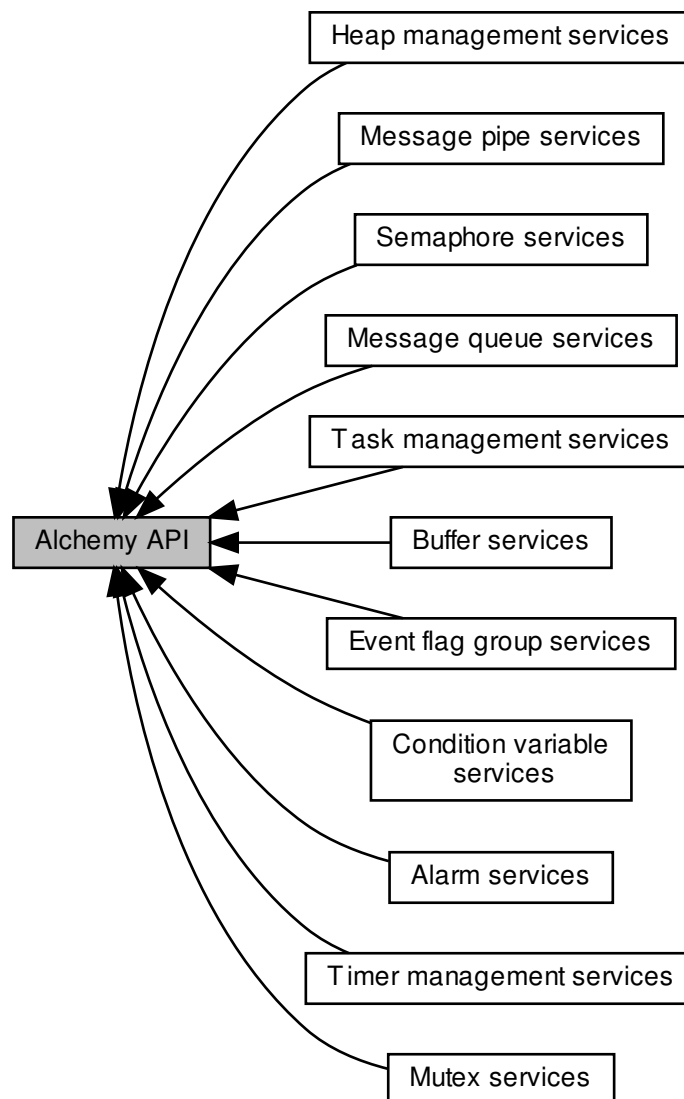
## Tags

[thread-unrestricted](#)

## 6.84 Alchemy API

A programming interface reminiscent from traditional RTOS APIs.

Collaboration diagram for Alchemy API:



### Modules

- [Alarm services](#)  
*General-purpose watchdog timers.*
- [Buffer services](#)  
*Lightweight FIFO IPC mechanism.*

- [Condition variable services](#)  
*POSIXish condition variable mechanism.*
- [Event flag group services](#)  
*Inter-task notification mechanism based on discrete flags.*
- [Heap management services](#)  
*Region of memory dedicated to real-time allocation.*
- [Mutex services](#)  
*POSIXish mutual exclusion servicesl.*
- [Message pipe services](#)  
*Two-way communication channel between Xenomai & Linux domains.*
- [Message queue services](#)  
*real-time IPC mechanism for sending messages of arbitrary size*
- [Semaphore services](#)  
*Counting semaphore IPC mechanism.*
- [Task management services](#)  
*Services dealing with preemptive multi-tasking.*
- [Timer management services](#)  
*Services for reading and spinning on the hardware timer.*

#### 6.84.1 Detailed Description

A programming interface reminiscent from traditional RTOS APIs.


This interface is an evolution of the former *native* API available with the Xenomai 2.x series.



## 6.85 Mutex services

POSIXish mutual exclusion services!

Collaboration diagram for Mutex services:



### Data Structures

- struct [RT\\_MUTEX\\_INFO](#)  
*Mutex status descriptor.*

### Functions

- int [rt\\_mutex\\_create](#) (RT\_MUTEX \*mutex, const char \*name)  
*Create a mutex.*
- int [rt\\_mutex\\_delete](#) (RT\_MUTEX \*mutex)  
*Delete a mutex.*
- int [rt\\_mutex\\_acquire\\_timed](#) (RT\_MUTEX \*mutex, const struct timespec \*abs\_timeout)  
*Acquire/lock a mutex (with absolute timeout date).*
- static int [rt\\_mutex\\_acquire\\_until](#) (RT\_MUTEX \*mutex, RTIME timeout)  
*Acquire/lock a mutex (with absolute scalar timeout).*
- static int [rt\\_mutex\\_acquire](#) (RT\_MUTEX \*mutex, RTIME timeout)  
*Acquire/lock a mutex (with relative scalar timeout).*
- int [rt\\_mutex\\_release](#) (RT\_MUTEX \*mutex)  
*Release/unlock a mutex.*
- int [rt\\_mutex\\_inquire](#) (RT\_MUTEX \*mutex, [RT\\_MUTEX\\_INFO](#) \*info)  
*Query mutex status.*
- int [rt\\_mutex\\_bind](#) (RT\_MUTEX \*mutex, const char \*name, RTIME timeout)  
*Bind to a mutex.*
- int [rt\\_mutex\\_unbind](#) (RT\_MUTEX \*mutex)  
*Unbind from a mutex.*

### 6.85.1 Detailed Description

POSIXish mutual exclusion services!

A mutex is a MUTual EXclusion object, and is useful for protecting shared data structures from concurrent modifications, and implementing critical sections and monitors.

A mutex has two possible states: unlocked (not owned by any task), and locked (owned by one task). A mutex can never be owned by two different tasks simultaneously. A task attempting to lock a mutex that is already locked by another task is blocked until the latter unlocks the mutex first.

Xenomai mutex services enforce a priority inheritance protocol in order to solve priority inversions.

### 6.85.2 Function Documentation

#### 6.85.2.1 `rt_mutex_acquire()`

```
int rt_mutex_acquire (
    RT_MUTEX * mutex,
    RTIME timeout ) [inline], [static]
```

Acquire/lock a mutex (with relative scalar timeout).

This routine is a variant of [rt\\_mutex\\_acquire\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

Parameters

<i>mutex</i>	The mutex descriptor.
<i>timeout</i>	A delay expressed in clock ticks. Passing <code>TM_INFINITE</code> the caller to block indefinitely. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without blocking in case <i>mutex</i> is already locked by another task.

Tags

[xthread-only](#), [switch-primary](#)

#### 6.85.2.2 `rt_mutex_acquire_timed()`

```
int rt_mutex_acquire_timed (
    RT_MUTEX * mutex,
    const struct timespec * abs_timeout )
```

Acquire/lock a mutex (with absolute timeout date).

Attempt to lock a mutex. The calling task is blocked until the mutex is available, in which case it is locked again before this service returns. Xenomai mutexes are implicitly recursive and implement the priority inheritance protocol.

Parameters

<i>mutex</i>	The mutex descriptor.
<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for the mutex to be available (see note). Passing NULL the caller to block indefinitely. Passing { .tv_sec = 0, .tv_nsec = 0 } causes the service to return immediately without blocking in case <i>mutex</i> is already locked by another task.

Returns

Zero is returned upon success. Otherwise:

- -ETIMEDOUT is returned if *abs\_timeout* is reached before the mutex is available.
- -EWOULDBLOCK is returned if *timeout* is { .tv\_sec = 0, .tv\_nsec = 0 } and the mutex is not immediately available.
- -EINTR is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task.
- -EINVAL is returned if *mutex* is not a valid mutex descriptor.
- -EIDRM is returned if *mutex* is deleted while the caller was waiting on it. In such event, *mutex* is no more valid upon return of this service.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

Tags

[xthread-only](#), [switch-primary](#)

Side effects

Over the Cobalt core, an Alchemy task with priority zero keeps running in primary mode until it releases the mutex, at which point it is switched back to secondary mode automatically.

Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

### 6.85.2.3 `rt_mutex_acquire_until()`

```
int rt_mutex_acquire_until (
    RT_MUTEX * mutex,
    RTIME abs_timeout ) [inline], [static]
```

Acquire/lock a mutex (with absolute scalar timeout).

This routine is a variant of [rt\\_mutex\\_acquire\\_timed\(\)](#) accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>mutex</i>	The mutex descriptor.
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing TM_INFINITE the caller to block indefinitely. Passing TM_NONBLOCK causes the service to return immediately without blocking in case <i>mutex</i> is already locked by another task.

## Tags

[xthread-only](#), [switch-primary](#)

## 6.85.2.4 rt\_mutex\_bind()

```
int rt_mutex_bind (
    RT_MUTEX * mutex,
    const char * name,
    RTIME timeout )
```

Bind to a mutex.

This routine creates a new descriptor to refer to an existing mutex identified by its symbolic name. If the object not exist on entry, the caller may block until a mutex of the given name is created.

## Parameters

<i>mutex</i>	The address of a mutex descriptor filled in by the operation. Contents of this memory is undefined upon failure.
<i>name</i>	A valid NULL-terminated name which identifies the mutex to bind to. This string should match the object name argument passed to <a href="#">rt_mutex_create()</a> .
<i>timeout</i>	The number of clock ticks to wait for the registration to occur (see note). Passing TM_INFINITE causes the caller to block indefinitely until the object is registered. Passing TM_NONBLOCK causes the service to return immediately without waiting if the object is not registered on entry.

## Returns

Zero is returned upon success. Otherwise:

- -EINTR is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task before the retrieval has completed.
- -EWOULDBLOCK is returned if *timeout* is equal to TM\_NONBLOCK and the searched object is not registered on entry.
- -ETIMEDOUT is returned if the object cannot be retrieved within the specified amount of time.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

The *timeout* value is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.85.2.5 `rt_mutex_create()`

```
int rt_mutex_create (
    RT_MUTEX * mutex,
    const char * name )
```

Create a mutex.

Create a mutual exclusion object that allows multiple tasks to synchronize access to a shared resource. A mutex is left in an unlocked state after creation.

## Parameters

<i>mutex</i>	The address of a mutex descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the mutex. When non-NULL and non-empty, a copy of this string is used for indexing the created mutex into the object registry.

## Returns

Zero is returned upon success. Otherwise:

- `-ENOMEM` is returned if the system fails to get memory from the main heap in order to create the mutex.
- `-EEXIST` is returned if the *name* is conflicting with an already registered mutex.
- `-EPERM` is returned if this service was called from an invalid context, e.g. interrupt or non-Xenomai thread.

## Tags

[xthread-only](#), [mode-unrestricted](#), [switch-secondary](#)

## Note

Mutexes can be shared by multiple processes which belong to the same Xenomai session.

6.85.2.6 `rt_mutex_delete()`

```
int rt_mutex_delete (
    RT_MUTEX * mutex )
```

Delete a mutex.

This routine deletes a mutex object previously created by a call to [rt\\_mutex\\_create\(\)](#).

Parameters

<i>mutex</i>	The mutex descriptor.
--------------	-----------------------

Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *alarm* is not a valid mutex descriptor.
- -EPERM is returned if this service was called from an asynchronous context.
- -EBUSY is returned upon an attempt to destroy the object referenced by *mutex* while it is referenced (for example, while being used in a [rt\\_mutex\\_acquire\(\)](#), [rt\\_mutex\\_acquire\\_timed\(\)](#) or [rt\\_mutex\\_acquire\\_until\(\)](#) by another task).

Tags

[mode-unrestricted](#), [switch-secondary](#)

6.85.2.7 `rt_mutex_inquire()`

```
int rt_mutex_inquire (
    RT_MUTEX * mutex,
    RT_MUTEX_INFO * info )
```

Query mutex status.

This routine returns the status information about the specified mutex.

Parameters

<i>mutex</i>	The mutex descriptor.
<i>info</i>	A pointer to the <a href="#">return</a> buffer" to copy the information to.

## Returns

Zero is returned and status information is written to the structure pointed at by *info* upon success. Otherwise:

- -EINVAL is returned if *mutex* is not a valid mutex descriptor.
- -EPERM is returned if this service is called from an interrupt context.

## Tags

[xthread-only](#), [switch-primary](#)

6.85.2.8 `rt_mutex_release()`

```
int rt_mutex_release (  
    RT_MUTEX * mutex )
```

Release/unlock a mutex.

This routine releases a mutex object previously locked by a call to [rt\\_mutex\\_acquire\(\)](#) or [rt\\_mutex\\_acquire\\_until\(\)](#). If the mutex is pended, the first waiting task (by priority order) is immediately unblocked and transferred the ownership of the mutex; otherwise, the mutex is left in an unlocked state.

## Parameters

<i>mutex</i>	The mutex descriptor.
--------------	-----------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *alarm* is not a valid mutex descriptor.
- -EPERM is returned if *mutex* is not owned by the current task, or more generally if this service was called from a context which cannot own any mutex (e.g. interrupt context).

## Tags

[xthread-only](#), [switch-primary](#)

6.85.2.9 `rt_mutex_unbind()`

```
int rt_mutex_unbind (  
    RT_MUTEX * mutex )
```

Unbind from a mutex.

## Parameters

<i>mutex</i>	The mutex descriptor.
--------------	-----------------------

This routine releases a previous binding to a mutex. After this call has returned, the descriptor is no more valid for referencing this object.



## 6.86 Message pipe services

Two-way communication channel between Xenomai & Linux domains.

Collaboration diagram for Message pipe services:



### Macros

- `#define P_MINOR_AUTO XNPIPE_MINOR_AUTO`  
*Creation flags.*
- `#define P_URGENT XNPIPE_URGENT`  
*Operation flags.*

### Functions

- `int rt_pipe_delete (RT_PIPE *pipe)`  
*Delete a message pipe.*
- `ssize_t rt_pipe_read_timed (RT_PIPE *pipe, void *buf, size_t size, const struct timespec *abs_↵  
timeout)`  
*Read a message from a pipe.*
- `static ssize_t rt_pipe_read_until (RT_PIPE *pipe, void *buf, size_t size, RTIME timeout)`  
*Read from a pipe (with absolute scalar timeout).*
- `static ssize_t rt_pipe_read (RT_PIPE *pipe, void *buf, size_t size, RTIME timeout)`  
*Read from a pipe (with relative scalar timeout).*
- `ssize_t rt_pipe_write (RT_PIPE *pipe, const void *buf, size_t size, int mode)`  
*Write a message to a pipe.*
- `ssize_t rt_pipe_stream (RT_PIPE *pipe, const void *buf, size_t size)`  
*Stream bytes through a pipe.*
- `int rt_pipe_bind (RT_PIPE *pipe, const char *name, RTIME timeout)`  
*Bind to a message pipe.*
- `int rt_pipe_unbind (RT_PIPE *pipe)`  
*Unbind from a message pipe.*
- `int rt_pipe_create (RT_PIPE *pipe, const char *name, int minor, size_t poolsize)`  
*Create a message pipe.*

### 6.86.1 Detailed Description

Two-way communication channel between Xenomai & Linux domains.

A message pipe is a two-way communication channel between Xenomai threads and normal Linux threads using regular file I/O operations on a pseudo-device. Pipes can be operated in a message-oriented fashion so that message boundaries are preserved, and also in byte-oriented streaming mode from real-time to normal Linux threads for optimal throughput.

Xenomai threads open their side of the pipe using the `rt_pipe_create()` service; regular Linux threads do the same by opening one of the `/dev/rtpN` special devices, where N is the minor number agreed upon between both ends of each pipe.

In addition, named pipes are available through the registry support, which automatically creates a symbolic link from entries under `/proc/xenomai/registry/rtpc/xddp/` to the corresponding special device file.

Note

Alchemy's message pipes are fully based on the [XDDP protocol](#) available from the RTDM/ipc driver.

### 6.86.2 Macro Definition Documentation

#### 6.86.2.1 P\_MINOR\_AUTO

```
#define P_MINOR_AUTO XNPIPE_MINOR_AUTO
```

Creation flags.

#### 6.86.2.2 P\_URGENT

```
#define P_URGENT XNPIPE_URGENT
```

Operation flags.

### 6.86.3 Function Documentation

#### 6.86.3.1 `rt_pipe_bind()`

```
int rt_pipe_bind (
    RT_PIPE * pipe,
    const char * name,
    RTIME timeout )
```

Bind to a message pipe.

This routine creates a new descriptor to refer to an existing message pipe identified by its symbolic name. If the object does not exist on entry, the caller may block until a pipe of the given name is created.

## Parameters

<i>pipe</i>	The address of a pipe descriptor filled in by the operation. Contents of this memory is undefined upon failure.
<i>name</i>	A valid NULL-terminated name which identifies the pipe to bind to. This string should match the object name argument passed to <code>rt_pipe_create()</code> .
<i>timeout</i>	The number of clock ticks to wait for the registration to occur (see note). Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the object is registered. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if the object is not registered on entry.

## Returns

Zero is returned upon success. Otherwise:

- `-EINTR` is returned if `rt_task_unblock()` was called for the current task before the retrieval has completed.
- `-EWOULDBLOCK` is returned if *timeout* is equal to `TM_NONBLOCK` and the searched object is not registered on entry.
- `-ETIMEDOUT` is returned if the object cannot be retrieved within the specified amount of time.
- `-EPERM` is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#)

## Note

The *timeout* value is interpreted as a multiple of the Alchemy clock resolution (see `--alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.86.3.2 `rt_pipe_create()`

```
int rt_pipe_create (
    RT_PIPE * pipe,
    const char * name,
    int minor,
    size_t poolsize )
```

## Create a message pipe.

This service opens a bi-directional communication channel for exchanging messages between Xenomai threads and regular Linux threads. Pipes natively preserve message boundaries, but can also be used in byte-oriented streaming mode from Xenomai to Linux.

`rt_pipe_create()` always returns immediately, even if no thread has opened the associated special device file yet. On the contrary, the non real-time side could block upon attempt to open the special device file until `rt_pipe_create()` is issued on the same pipe from a Xenomai thread, unless `O_NONBLOCK` was given to the `open(2)` system call.

## Parameters

<i>pipe</i>	The address of a pipe descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the pipe. When non-NULL and non-empty, a copy of this string is used for indexing the created pipe into the object registry.

Named pipes are supported through the use of the registry. Passing a valid *name* parameter when creating a message pipe causes a symbolic link to be created from `/proc/xenomai/registry/rtipc/xddp/name` to the associated special device (i.e. `/dev/rtp*`), so that the specific *minor* information does not need to be known from those processes for opening the proper device file. In such a case, both sides of the pipe only need to agree upon a symbolic name to refer to the same data path, which is especially useful whenever the *minor* number is picked up dynamically using an adaptive algorithm, such as passing `P_MINOR_AUTO` as *minor* value.

## Parameters

<i>minor</i>	The minor number of the device associated with the pipe. Passing <code>P_MINOR_AUTO</code> causes the minor number to be auto-allocated. In such a case, a symbolic link will be automatically created from <code>/proc/xenomai/registry/rtipc/xddp/name</code> to the allocated pipe device entry. Valid minor numbers range from 0 to <code>CONFIG_XENO_OPT_PIPE_NRDEV-1</code> .
<i>poolsize</i>	Specifies the size of a dedicated buffer pool for the pipe. Passing 0 means that all message allocations for this pipe are performed on the Cobalt core heap.

## Returns

The *minor* number assigned to the connection is returned upon success. Otherwise:

- `-ENOMEM` is returned if the system fails to get memory from the main heap in order to create the pipe.
- `-ENODEV` is returned if *minor* is different from `P_MINOR_AUTO` and is not a valid minor number.
- `-EEXIST` is returned if the *name* is conflicting with an already registered pipe.
- `-EBUSY` is returned if *minor* is already open.
- `-EPERM` is returned if this service was called from an invalid context, e.g. interrupt or non-Xenomai thread.

## Tags

[xthread-only](#), [mode-unrestricted](#), [switch-secondary](#)

6.86.3.3 `rt_pipe_delete()`

```
int rt_pipe_delete (
    RT_PIPE * pipe )
```

Delete a message pipe.

This routine deletes a pipe object previously created by a call to `rt_pipe_create()`. All resources attached to that pipe are automatically released, all pending data is flushed.

## Parameters

<i>pipe</i>	The pipe descriptor.
-------------	----------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *pipe* is not a valid pipe descriptor.
- -EIDRM is returned if *pipe* is a closed pipe descriptor.
- -EPERM is returned if this service was called from an asynchronous context.

## Tags

[mode-unrestricted](#), [switch-secondary](#)

6.86.3.4 `rt_pipe_read()`

```
ssize_t rt_pipe_read (
    RT_PIPE * pipe,
    void * buf,
    size_t size,
    RTIME timeout ) [inline], [static]
```

Read from a pipe (with relative scalar timeout).

This routine is a variant of [rt\\_queue\\_read\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

## Parameters

<i>pipe</i>	The pipe descriptor.
<i>buf</i>	A pointer to a memory area which will be written upon success with the message received.
<i>size</i>	The count of bytes from the received message to read up into <i>buf</i> . If <i>size</i> is lower than the actual message size, -ENOBUFS is returned since the incompletely received message would be lost. If <i>size</i> is zero, this call returns immediately with no other action.
<i>timeout</i>	A delay expressed in clock ticks. Passing TM_INFINITE causes the caller to block indefinitely until a message is available. Passing TM_NONBLOCK causes the service to return immediately without blocking in case no message is available.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.86.3.5 `rt_pipe_read_timed()`

```
ssize_t rt_pipe_read_timed (
    RT_PIPE * pipe,
    void * buf,
    size_t size,
    const struct timespec * abs_timeout )
```

Read a message from a pipe.

This service reads the next available message from a given pipe.

## Parameters

<i>pipe</i>	The pipe descriptor.
<i>buf</i>	A pointer to a memory area which will be written upon success with the message received.
<i>size</i>	The count of bytes from the received message to read up into <i>buf</i> . If <i>size</i> is lower than the actual message size, <code>-ENOBUFS</code> is returned since the incompletely received message would be lost. If <i>size</i> is zero, this call returns immediately with no other action.
<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for a message to be available from the pipe (see note). Passing <code>NULL</code> causes the caller to block indefinitely until a message is available. Passing <code>{ .tv_sec = 0, .tv_nsec = 0 }</code> causes the service to return immediately without blocking in case no message is available.

## Returns

The number of bytes available from the received message is returned upon success. Otherwise:

- `-ETIMEDOUT` is returned if *abs\_timeout* is reached before a message arrives.
- `-EWOULDBLOCK` is returned if *abs\_timeout* is `{ .tv_sec = 0, .tv_nsec = 0 }` and no message is immediately available on entry to the call.
- `-EINTR` is returned if `rt_task_unblock()` was called for the current task before a message was available.
- `-EINVAL` is returned if *pipe* is not a valid pipe descriptor.
- `-EIDRM` is returned if *pipe* is deleted while the caller was waiting for a message. In such event, *pipe* is no more valid upon return of this service.
- `-EPERM` is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.86.3.6 `rt_pipe_read_until()`

```
ssize_t rt_pipe_read_until (
    RT_PIPE * pipe,
    void * buf,
    size_t size,
    RTIME abs_timeout ) [inline], [static]
```

Read from a pipe (with absolute scalar timeout).

This routine is a variant of `rt_queue_read_timed()` accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>pipe</i>	The pipe descriptor.
<i>buf</i>	A pointer to a memory area which will be written upon success with the message received.
<i>size</i>	The count of bytes from the received message to read up into <i>buf</i> . If <i>size</i> is lower than the actual message size, <code>-ENOBUFS</code> is returned since the incompletely received message would be lost. If <i>size</i> is zero, this call returns immediately with no other action.
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until a message is available. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without blocking in case no message is available.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.86.3.7 `rt_pipe_stream()`

```
ssize_t rt_pipe_stream (
    RT_PIPE * pipe,
    const void * buf,
    size_t size )
```

Stream bytes through a pipe.

This service writes a sequence of bytes to be received from the associated special device. Unlike `rt_queue_send()`, this service does not preserve message boundaries. Instead, an internal buffer is filled on the fly with the data, which will be consumed as soon as the receiver wakes up.

Data buffers sent by the `rt_pipe_stream()` service are always transmitted in FIFO order (i.e. `P_NORMAL` mode).

## Parameters

<i>pipe</i>	The pipe descriptor.
<i>buf</i>	The address of the first data byte to send. The data will be copied to an internal buffer before transmission.
<i>size</i>	The size in bytes of the buffer. Zero is a valid value, in which case the service returns immediately without sending any data.

## Returns

The number of bytes sent upon success; this value may be lower than *size*, depending on the available space in the internal buffer. Otherwise:

- -EINVAL is returned if *mode* is invalid or *pipe* is not a pipe descriptor.
- -ENOMEM is returned if not enough buffer space is available to complete the operation.
- -EIDRM is returned if *pipe* is a closed pipe descriptor.

## Note

Writing data to a pipe before any peer has opened the associated special device is allowed. The output will be buffered until then, only restricted by the available memory in the associated buffer pool (see [rt\\_pipe\\_create\(\)](#)).

## Tags

[xcontext](#), [switch-primary](#)

## 6.86.3.8 rt\_pipe\_unbind()

```
int rt_pipe_unbind (
    RT_PIPE * pipe )
```

Unbind from a message pipe.

## Parameters

<i>pipe</i>	The pipe descriptor.
-------------	----------------------

This routine releases a previous binding to a message pipe. After this call has returned, the descriptor is no more valid for referencing this object.

## Tags

[thread-unrestricted](#)

## 6.86.3.9 rt\_pipe\_write()

```
ssize_t rt_pipe_write (
    RT_PIPE * pipe,
    const void * buf,
```



```
size_t size,  
int mode )
```

Write a message to a pipe.

This service writes a complete message to be received from the associated special device. [rt\\_pipe\\_write\(\)](#) always preserves message boundaries, which means that all data sent through a single call of this service will be gathered in a single `read(2)` operation from the special device.

This service differs from `rt_pipe_send()` in that it accepts a pointer to the raw data to be sent, instead of a canned message buffer.

Parameters

<i>pipe</i>	The pipe descriptor.
<i>buf</i>	The address of the first data byte to send. The data will be copied to an internal buffer before transmission.
<i>size</i>	The size in bytes of the message (payload data only). Zero is a valid value, in which case the service returns immediately without sending any message.
<i>mode</i>	A set of flags affecting the operation:

- `P_URGENT` causes the message to be prepended to the output queue, ensuring a LIFO ordering.
- `P_NORMAL` causes the message to be appended to the output queue, ensuring a FIFO ordering.

Returns

Upon success, this service returns *size*. Upon error, one of the following error codes is returned:

- `-EINVAL` is returned if *mode* is invalid or *pipe* is not a pipe descriptor.
- `-ENOMEM` is returned if not enough buffer space is available to complete the operation.
- `-EIDRM` is returned if *pipe* is a closed pipe descriptor.

Note

Writing data to a pipe before any peer has opened the associated special device is allowed. The output will be buffered until then, only restricted by the available memory in the associated buffer pool (see [rt\\_pipe\\_create\(\)](#)).

Tags

[xcontext](#), [switch-primary](#)

## 6.87 Message queue services

real-time IPC mechanism for sending messages of arbitrary size

Collaboration diagram for Message queue services:



### Data Structures

- struct [RT\\_QUEUE\\_INFO](#)  
*Queue status descriptor.*

### Macros

- `#define Q_PRIO 0x1` /\* Pend by task priority order. \*/  
*Creation flags.*

### Functions

- int [rt\\_queue\\_create](#) (RT\_QUEUE \*queue, const char \*name, size\_t poolsize, size\_t qlimit, int mode)  
*Create a message queue.*
- int [rt\\_queue\\_delete](#) (RT\_QUEUE \*queue)  
*Delete a message queue.*
- void \* [rt\\_queue\\_alloc](#) (RT\_QUEUE \*queue, size\_t size)  
*Allocate a message buffer.*
- int [rt\\_queue\\_free](#) (RT\_QUEUE \*queue, void \*buf)  
*Free a message buffer.*
- int [rt\\_queue\\_send](#) (RT\_QUEUE \*queue, const void \*buf, size\_t size, int mode)  
*Send a message to a queue.*
- ssize\_t [rt\\_queue\\_receive\\_timed](#) (RT\_QUEUE \*queue, void \*\*bufp, const struct timespec \*abs\_↵ timeout)  
*Receive a message from a queue (with absolute timeout date).*
- static ssize\_t [rt\\_queue\\_receive\\_until](#) (RT\_QUEUE \*queue, void \*\*bufp, RTIME timeout)  
*Receive from a queue (with absolute scalar timeout).*

- static ssize\_t `rt_queue_receive` (RT\_QUEUE \*queue, void \*\*bufp, RTIME timeout)

*Receive from a queue (with relative scalar timeout).*

- ssize\_t `rt_queue_read_timed` (RT\_QUEUE \*queue, void \*buf, size\_t size, const struct timespec \*abs\_timeout)

*Read from a queue.*

- static ssize\_t `rt_queue_read_until` (RT\_QUEUE \*queue, void \*buf, size\_t size, RTIME timeout)

*Read from a queue (with absolute scalar timeout).*

- static ssize\_t `rt_queue_read` (RT\_QUEUE \*queue, void \*buf, size\_t size, RTIME timeout)

*Read from a queue (with relative scalar timeout).*

- int `rt_queue_flush` (RT\_QUEUE \*queue)

*Flush pending messages from a queue.*

- int `rt_queue_inquire` (RT\_QUEUE \*queue, RT\_QUEUE\_INFO \*info)

*Query queue status.*

- int `rt_queue_bind` (RT\_QUEUE \*queue, const char \*name, RTIME timeout)

*Bind to a message queue.*

- int `rt_queue_unbind` (RT\_QUEUE \*queue)

*Unbind from a message queue.*

### 6.87.1 Detailed Description

real-time IPC mechanism for sending messages of arbitrary size

Message queueing is a method by which real-time tasks can exchange or pass data through a Xenomai-managed queue of messages. Messages can vary in length and be assigned different types or usages. A message queue can be created by one task and used by multiple tasks that send and/or receive messages to the queue.

### 6.87.2 Macro Definition Documentation

#### 6.87.2.1 Q\_PRIO

```
#define Q_PRIO 0x1 /* Pend by task priority order. */
```

Creation flags.

Referenced by `rt_queue_create()`.

### 6.87.3 Function Documentation

6.87.3.1 `rt_queue_alloc()`

```
void * rt_queue_alloc (
    RT_QUEUE * q,
    size_t size )
```

Allocate a message buffer.

This service allocates a message buffer from the queue's internal pool. This buffer can be filled in with payload information, prior enqueueing it by a call to `rt_queue_send()`. When used in pair, these services provide a zero-copy interface for sending messages.

## Parameters

<i>q</i>	The queue descriptor.
<i>size</i>	The requested size in bytes of the buffer. Zero is an acceptable value, which means that the message conveys no payload; in this case, the receiver will get a zero-sized message.

## Returns

The address of the allocated buffer upon success, or NULL if the call fails.

## Tags

[unrestricted](#), [switch-primary](#)

6.87.3.2 `rt_queue_bind()`

```
int rt_queue_bind (
    RT_QUEUE * q,
    const char * name,
    RTIME timeout )
```

Bind to a message queue.

This routine creates a new descriptor to refer to an existing message queue identified by its symbolic name. If the object does not exist on entry, the caller may block until a queue of the given name is created.

## Parameters

<i>q</i>	The address of a queue descriptor filled in by the operation. Contents of this memory is undefined upon failure.
<i>name</i>	A valid NULL-terminated name which identifies the queue to bind to. This string should match the object name argument passed to <code>rt_queue_create()</code> .
<i>timeout</i>	The number of clock ticks to wait for the registration to occur (see note). Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the object is registered. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if the object is not registered on entry.

## Returns

Zero is returned upon success. Otherwise:

- -EINTR is returned if `rt_task_unblock()` was called for the current task before the retrieval has completed.
- -EWOULDBLOCK is returned if `timeout` is equal to `TM_NONBLOCK` and the searched object is not registered on entry.
- -ETIMEDOUT is returned if the object cannot be retrieved within the specified amount of time.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

The `timeout` value is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.87.3.3 `rt_queue_create()`

```
int rt_queue_create (
    RT_QUEUE * q,
    const char * name,
    size_t poolsize,
    size_t qlimit,
    int mode )
```

Create a message queue.

Create a message queue object which allows multiple tasks to exchange data through the use of variable-sized messages. A message queue is created empty.

## Parameters

<i>q</i>	The address of a queue descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the queue. When non-NULL and non-empty, a copy of this string is used for indexing the created queue into the object registry.
<i>poolsize</i>	The size (in bytes) of the message buffer pool to be pre-allocated for holding messages. Message buffers will be claimed and released to this pool. The buffer pool memory cannot be extended. See note.
<i>qlimit</i>	This parameter allows to limit the maximum number of messages which can be queued at any point in time, sending to a full queue begets an error. The special value <code>Q_UNLIMITED</code> can be passed to disable the limit check.
<i>mode</i>	The queue creation mode. The following flags can be OR'ed into this bitmask, each of them affecting the new queue:

- Q\_FIFO makes tasks pend in FIFO order on the queue for consuming messages.
- Q\_PRIO makes tasks pend in priority order on the queue.

#### Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *mode* is invalid or *poolsize* is zero.
- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the queue.
- -EEXIST is returned if the *name* is conflicting with an already registered queue.
- -EPERM is returned if this service was called from an invalid context, e.g. interrupt or non-Xenomai thread.

#### Tags

[xthread-only](#), [mode-unrestricted](#), [switch-secondary](#)

#### Note

Queues can be shared by multiple processes which belong to the same Xenomai session. Each message pending into the queue consumes four long words plus the actual payload size, aligned to the next long word boundary. e.g. a 6 byte message on a 32 bit platform would require 24 bytes of storage into the pool.

When *qlimit* is given (i.e. different from Q\_UNLIMITED), this overhead is accounted for automatically, so that *qlimit* messages of *poolsize* / *qlimit* bytes can be stored into the pool concurrently. Otherwise, *poolsize* is increased by 5% internally to cope with such overhead.

References Q\_PRIO.

#### 6.87.3.4 rt\_queue\_delete()

```
int rt_queue_delete (
    RT_QUEUE * q )
```

Delete a message queue.

This routine deletes a queue object previously created by a call to [rt\\_queue\\_create\(\)](#). All resources attached to that queue are automatically released, including all pending messages.

#### Parameters

<i>q</i>	The queue descriptor.
----------	-----------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *q* is not a valid queue descriptor.
- -EPERM is returned if this service was called from an asynchronous context.

## Tags

[mode-unrestricted](#), [switch-secondary](#)

6.87.3.5 `rt_queue_flush()`

```
int rt_queue_flush (  
    RT_QUEUE * q )
```

Flush pending messages from a queue.

This routine flushes all messages currently pending in a queue, releasing all message buffers appropriately.

## Parameters

<i>q</i>	The queue descriptor.
----------	-----------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *q* is not a valid queue descriptor.

## Tags

[unrestricted](#), [switch-primary](#)

6.87.3.6 `rt_queue_free()`

```
int rt_queue_free (  
    RT_QUEUE * q,  
    void * buf )
```

Free a message buffer.

This service releases a message buffer to the queue's internal pool.

## Parameters

<i>q</i>	The queue descriptor.
<i>buf</i>	The address of the message buffer to free. Even zero-sized messages carrying no payload data must be freed, since they are assigned a valid memory space to store internal information.

## Returns

Zero is returned upon success, or `-EINVAL` if *buf* is not a valid message buffer previously allocated by the `rt_queue_alloc()` service, or the caller did not get ownership of the message through a successful return from `rt_queue_receive()`.

## Tags

[unrestricted](#), [switch-primary](#)

6.87.3.7 `rt_queue_inquire()`

```
int rt_queue_inquire (
    RT_QUEUE * q,
    RT_QUEUE_INFO * info )
```

Query queue status.

This routine returns the status information about the specified queue.

## Parameters

<i>q</i>	The queue descriptor.
<i>info</i>	A pointer to the <a href="#">return</a> buffer" to copy the information to.

## Returns

Zero is returned and status information is written to the structure pointed at by *info* upon success. Otherwise:

- `-EINVAL` is returned if *q* is not a valid queue descriptor.

## Tags

[unrestricted](#), [switch-primary](#)



6.87.3.8 `rt_queue_read()`

```
ssize_t rt_queue_read (
    RT_QUEUE * q,
    void * buf,
    size_t size,
    RTIME timeout ) [inline], [static]
```

Read from a queue (with relative scalar timeout).

This routine is a variant of [rt\\_queue\\_read\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

## Parameters

<i>q</i>	The queue descriptor.
<i>buf</i>	A pointer to a memory area which will be written upon success with the received message payload.
<i>size</i>	The length in bytes of the memory area pointed to by <i>buf</i> .
<i>timeout</i>	A delay expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until a message is available. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without blocking in case no message is available.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.87.3.9 `rt_queue_read_timed()`

```
ssize_t rt_queue_read_timed (
    RT_QUEUE * q,
    void * buf,
    size_t size,
    const struct timespec * abs_timeout )
```

Read from a queue.

This service reads the next available message from a given queue.

## Parameters

<i>q</i>	The queue descriptor.
<i>buf</i>	A pointer to a memory area which will be written upon success with the received message payload. The internal message buffer conveying the data is automatically freed by this call. If <code>--enable-pshared</code> is enabled in the configuration, <i>buf</i> must have been obtained from the Xenomai memory allocator via <code>xnmalloc()</code> or any service based on it, such as <a href="#">rt_heap_alloc()</a> .

## Parameters

<i>size</i>	The length in bytes of the memory area pointed to by <i>buf</i> . Messages larger than <i>size</i> are truncated appropriately.
<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for a message to be available from the queue (see note). Passing NULL causes the caller to block indefinitely until a message is available. Passing { .tv_sec = 0, .tv_nsec = 0 } causes the service to return immediately without blocking in case no message is available.

## Returns

The number of bytes copied to *buf* is returned upon success. Zero is a possible value corresponding to a zero-sized message passed to `rt_queue_send()` or `rt_queue_write()`. Otherwise:

- -ETIMEDOUT is returned if *abs\_timeout* is reached before a message arrives.
- -EWOULDBLOCK is returned if *abs\_timeout* is { .tv\_sec = 0, .tv\_nsec = 0 } and no message is immediately available on entry to the call.
- -EINTR is returned if `rt_task_unblock()` was called for the current task before a message was available.
- -EINVAL is returned if *q* is not a valid queue descriptor.
- -EIDRM is returned if *q* is deleted while the caller was waiting for a message. In such event, *q* is no more valid upon return of this service.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.87.3.10 `rt_queue_read_until()`

```
ssize_t rt_queue_read_until (
    RT_QUEUE * q,
    void * buf,
    size_t size,
    RTIME abs_timeout ) [inline], [static]
```

Read from a queue (with absolute scalar timeout).

This routine is a variant of `rt_queue_read_timed()` accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>q</i>	The queue descriptor.
<i>buf</i>	A pointer to a memory area which will be written upon success with the received message payload.
<i>size</i>	The length in bytes of the memory area pointed to by <i>buf</i> .
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until a message is available. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without blocking in case no message is available.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.87.3.11 `rt_queue_receive()`

```
ssize_t rt_queue_receive (
    RT_QUEUE * q,
    void ** bufp,
    RTIME timeout ) [inline], [static]
```

Receive from a queue (with relative scalar timeout).

This routine is a variant of [rt\\_queue\\_receive\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

## Parameters

<i>q</i>	The queue descriptor.
<i>bufp</i>	A pointer to a memory location which will be written with the address of the received message.
<i>timeout</i>	A delay expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until a message is available. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without blocking in case no message is available.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.87.3.12 `rt_queue_receive_timed()`

```
ssize_t rt_queue_receive_timed (
    RT_QUEUE * q,
```

```
void ** bufp,
const struct timespec * abs_timeout )
```

Receive a message from a queue (with absolute timeout date).

This service receives the next available message from a given queue.

Parameters

<i>q</i>	The queue descriptor.
<i>bufp</i>	A pointer to a memory location which will be written with the address of the received message, upon success. Once consumed, the message space should be freed using <a href="#">rt_queue_free()</a> .
<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for a message to be available from the queue (see note). Passing NULL causes the caller to block indefinitely until a message is available. Passing { .tv_sec = 0, .tv_nsec = 0 } causes the service to return immediately without blocking in case no message is available.

Returns

The number of bytes available from the received message is returned upon success. Zero is a possible value corresponding to a zero-sized message passed to [rt\\_queue\\_send\(\)](#) or [rt\\_queue\\_write\(\)](#). Otherwise:

- -ETIMEDOUT is returned if *abs\_timeout* is reached before a message arrives.
- -EWOULDBLOCK is returned if *abs\_timeout* is { .tv\_sec = 0, .tv\_nsec = 0 } and no message is immediately available on entry to the call.
- -EINTR is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task before a message was available.
- -EINVAL is returned if *q* is not a valid queue descriptor.
- -EIDRM is returned if *q* is deleted while the caller was waiting for a message. In such event, *q* is no more valid upon return of this service.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

Tags

[xthread-nowait](#), [switch-primary](#)

Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.87.3.13 `rt_queue_receive_until()`

```
ssize_t rt_queue_receive_until (
    RT_QUEUE * q,
    void ** bufp,
    RTIME abs_timeout ) [inline], [static]
```

Receive from a queue (with absolute scalar timeout).

This routine is a variant of `rt_queue_receive_timed()` accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>q</i>	The queue descriptor.
<i>bufp</i>	A pointer to a memory location which will be written with the address of the received message.
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until a message is available. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without blocking in case no message is available.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.87.3.14 `rt_queue_send()`

```
int rt_queue_send (
    RT_QUEUE * q,
    const void * buf,
    size_t size,
    int mode )
```

Send a message to a queue.

This service sends a complete message to a given queue. The message must have been allocated by a previous call to `rt_queue_alloc()`.

## Parameters

<i>q</i>	The queue descriptor.
<i>buf</i>	The address of the message buffer to be sent. The message buffer must have been allocated using the <code>rt_queue_alloc()</code> service. Once passed to <code>rt_queue_send()</code> , the memory pointed to by <i>buf</i> is no more under the control of the sender and thus should not be referenced by it anymore; deallocation of this memory must be handled on the receiving side.
<i>size</i>	The actual size in bytes of the message, which may be lower than the allocated size for the buffer obtained from <code>rt_queue_alloc()</code> . Zero is a valid value, in which case an empty message will be sent.
<i>mode</i>	A set of flags affecting the operation:

- `Q_URGENT` causes the message to be prepended to the message queue, ensuring a LIFO ordering.
- `Q_NORMAL` causes the message to be appended to the message queue, ensuring a FIFO ordering.
- `Q_BROADCAST` causes the message to be sent to all tasks currently waiting for messages. The message is not copied; a reference count is maintained instead so that the message will remain valid until the last receiver releases its own reference using `rt_queue_free()`, after which the message space will be returned to the queue's internal pool.

#### Returns

Upon success, this service returns the number of receivers which got awoken as a result of the operation. If zero is returned, no task was waiting on the receiving side of the queue, and the message has been enqueued. Upon error, one of the following error codes is returned:

- `-EINVAL` is returned if `q` is not a message queue descriptor, `mode` is invalid, or `buf` is `NULL`.
- `-ENOMEM` is returned if queuing the message would exceed the limit defined for the queue at creation.

#### Tags

[unrestricted](#), [switch-primary](#)

#### 6.87.3.15 `rt_queue_unbind()`

```
int rt_queue_unbind (
    RT_QUEUE * q )
```

Unbind from a message queue.

#### Parameters

<code>q</code>	The queue descriptor.
----------------	-----------------------

This routine releases a previous binding to a message queue. After this call has returned, the descriptor is no more valid for referencing this object.


#### Tags

[thread-unrestricted](#)

## 6.88 Semaphore services

Counting semaphore IPC mechanism.

Collaboration diagram for Semaphore services:



### Data Structures

- struct [RT\\_SEM\\_INFO](#)  
*Semaphore status descriptor.*

### Macros

- `#define S_PRIO 0x1` /\* Pend by task priority order. \*/  
*Creation flags.*

### Functions

- int [rt\\_sem\\_create](#) (RT\_SEM \*sem, const char \*name, unsigned long icount, int mode)  
*Create a counting semaphore.*
- int [rt\\_sem\\_delete](#) (RT\_SEM \*sem)  
*Delete a semaphore.*
- int [rt\\_sem\\_p\\_timed](#) (RT\_SEM \*sem, const struct timespec \*abs\_timeout)  
*Pend on a semaphore.*
- static int [rt\\_sem\\_p\\_until](#) (RT\_SEM \*sem, RTIME timeout)  
*Pend on a semaphore (with absolute scalar timeout).*
- static int [rt\\_sem\\_p](#) (RT\_SEM \*sem, RTIME timeout)  
*Pend on a semaphore (with relative scalar timeout).*
- int [rt\\_sem\\_v](#) (RT\_SEM \*sem)  
*Signal a semaphore.*
- int [rt\\_sem\\_broadcast](#) (RT\_SEM \*sem)  
*Broadcast a semaphore.*
- int [rt\\_sem\\_inquire](#) (RT\_SEM \*sem, [RT\\_SEM\\_INFO](#) \*info)

*Query semaphore status.*

- int `rt_sem_bind` (RT\_SEM \*sem, const char \*name, RTIME timeout)

*Bind to a semaphore.*

- int `rt_sem_unbind` (RT\_SEM \*sem)

*Unbind from a semaphore.*

### 6.88.1 Detailed Description

Counting semaphore IPC mechanism.

A counting semaphore is a synchronization object for controlling the concurrency level allowed in accessing a resource from multiple real-time tasks, based on the value of a count variable accessed atomically. The semaphore is used through the P ("Proberen", from the Dutch "test and decrement") and V ("Verhogen", increment) operations. The P operation decrements the semaphore count by one if non-zero, or waits until a V operation is issued by another task. Conversely, the V operation releases a resource by incrementing the count by one, unblocking the heading task waiting on the P operation if any. Waiting on a semaphore may cause a priority inversion.

If no more than a single resource is made available at any point in time, the semaphore enforces mutual exclusion and thus can be used to serialize access to a critical section. However, mutexes should be used instead in order to prevent priority inversions, based on the priority inheritance protocol.

### 6.88.2 Macro Definition Documentation

#### 6.88.2.1 S\_PRIO

```
#define S_PRIO 0x1 /* Pend by task priority order. */
```

Creation flags.

Referenced by `rt_sem_create()`.

### 6.88.3 Function Documentation

#### 6.88.3.1 `rt_sem_bind()`

```
int rt_sem_bind (
    RT_SEM * sem,
    const char * name,
    RTIME timeout )
```

Bind to a semaphore.

This routine creates a new descriptor to refer to an existing semaphore identified by its symbolic name. If the object does not exist on entry, the caller may block until a semaphore of the given name is created.



## Parameters

<i>sem</i>	The address of a semaphore descriptor filled in by the operation. Contents of this memory is undefined upon failure.
<i>name</i>	A valid NULL-terminated name which identifies the semaphore to bind to. This string should match the object name argument passed to <a href="#">rt_sem_create()</a> .
<i>timeout</i>	The number of clock ticks to wait for the registration to occur (see note). Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the object is registered. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if the object is not registered on entry.

## Returns

Zero is returned upon success. Otherwise:

- `-EINTR` is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task before the retrieval has completed.
- `-EWOULDBLOCK` is returned if *timeout* is equal to `TM_NONBLOCK` and the searched object is not registered on entry.
- `-ETIMEDOUT` is returned if the object cannot be retrieved within the specified amount of time.
- `-EPERM` is returned if this service should block, but was not called from a Xenomai thread.

## Tags

[xthread-nowait](#), [switch-primary](#)

## Note

The *timeout* value is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.88.3.2 `rt_sem_broadcast()`

```
int rt_sem_broadcast (
    RT_SEM * sem )
```

Broadcast a semaphore.

All tasks currently waiting on the semaphore are immediately unblocked. The semaphore count is set to zero.

## Parameters

<i>sem</i>	The semaphore descriptor.
------------	---------------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *sem* is not a valid semaphore descriptor.

## Tags

unrestricted

## 6.88.3.3 rt\_sem\_create()

```
int rt_sem_create (
    RT_SEM * sem,
    const char * name,
    unsigned long icount,
    int mode )
```

Create a counting semaphore.

## Parameters

<i>sem</i>	The address of a semaphore descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the semaphore. When non-NULL and non-empty, a copy of this string is used for indexing the created semaphore into the object registry.
<i>icount</i>	The initial value of the counting semaphore.
<i>mode</i>	The semaphore creation mode. The following flags can be OR'ed into this bitmask:

- S\_FIFO makes tasks pend in FIFO order on the semaphore.
- S\_PRIO makes tasks pend in priority order on the semaphore.
- S\_PULSE causes the semaphore to behave in "pulse" mode. In this mode, the V (signal) operation attempts to release a single waiter each time it is called, without incrementing the semaphore count, even if no waiter is pending. For this reason, the semaphore count in pulse mode remains zero.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *icount* is non-zero and S\_PULSE is set in *mode*, or *mode* is otherwise invalid.
- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the semaphore.

- -EEXIST is returned if the *name* is conflicting with an already registered semaphore.
- -EPERM is returned if this service was called from an invalid context, e.g. interrupt or non-Xenomai thread.

## Tags

[xthread-only](#), [mode-unrestricted](#), [switch-secondary](#)

## Note

Semaphores can be shared by multiple processes which belong to the same Xenomai session.

References `S_PRIO`.

6.88.3.4 `rt_sem_delete()`

```
int rt_sem_delete (
    RT_SEM * sem )
```

Delete a semaphore.

This routine deletes a semaphore previously created by a call to [rt\\_sem\\_create\(\)](#).

## Parameters

<i>sem</i>	The semaphore descriptor.
------------	---------------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *sem* is not a valid semaphore descriptor.
- -EPERM is returned if this service was called from an asynchronous context.

## Tags

[mode-unrestricted](#), [switch-secondary](#)

6.88.3.5 `rt_sem_inquire()`

```
int rt_sem_inquire (
    RT_SEM * sem,
    RT_SEM_INFO * info )
```

Query semaphore status.

This routine returns the status information about the specified semaphore.

## Parameters

<i>sem</i>	The semaphore descriptor.
<i>info</i>	A pointer to the <a href="#">return buffer</a> to copy the information to.

## Returns

Zero is returned and status information is written to the structure pointed at by *info* upon success. Otherwise:

- -EINVAL is returned if *sem* is not a valid semaphore descriptor.

## Tags

[unrestricted](#)

6.88.3.6 `rt_sem_p()`

```
int rt_sem_p (
    RT_SEM * sem,
    RTIME timeout ) [inline], [static]
```

Pend on a semaphore (with relative scalar timeout).

This routine is a variant of [rt\\_sem\\_p\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

## Parameters

<i>sem</i>	The semaphore descriptor.
<i>timeout</i>	A delay expressed in clock ticks. Passing TM_INFINITE causes the caller to block indefinitely until the request is satisfied. Passing TM_NONBLOCK causes the service to return without blocking in case the request cannot be satisfied immediately.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.88.3.7 `rt_sem_p_timed()`

```
int rt_sem_p_timed (
    RT_SEM * sem,
    const struct timespec * abs_timeout )
```

Pend on a semaphore.

Test and decrement the semaphore count. If the semaphore value is greater than zero, it is decremented by one and the service immediately returns to the caller. Otherwise, the caller is blocked until the semaphore is either signaled or destroyed, unless a non-blocking operation was required.

Parameters

<i>sem</i>	The semaphore descriptor.
<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for the request to be satisfied (see note). Passing NULL causes the caller to block indefinitely until the request is satisfied. Passing { .tv_sec = 0, .tv_nsec = 0 } causes the service to return without blocking in case the request cannot be satisfied immediately.

Returns

Zero is returned upon success. Otherwise:

- -ETIMEDOUT is returned if *abs\_timeout* is reached before the request is satisfied.
- -EWOULDBLOCK is returned if *abs\_timeout* is { .tv\_sec = 0, .tv\_nsec = 0 } and the semaphore count is zero on entry.
- -EINTR is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task before the request is satisfied.
- -EINVAL is returned if *sem* is not a valid semaphore descriptor.
- -EIDRM is returned if *sem* is deleted while the caller was sleeping on it. In such a case, *sem* is no more valid upon return of this service.
- -EPERM is returned if this service should block, but was not called from a Xenomai thread.

Tags

[xthread-nowait](#), [switch-primary](#)

Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

### 6.88.3.8 `rt_sem_p_until()`

```
int rt_sem_p_until (
    RT_SEM * sem,
    RTIME abs_timeout ) [inline], [static]
```

Pend on a semaphore (with absolute scalar timeout).

This routine is a variant of [rt\\_sem\\_p\\_timed\(\)](#) accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>sem</i>	The semaphore descriptor.
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing TM_INFINITE causes the caller to block indefinitely until the request is satisfied. Passing TM_NONBLOCK causes the service to return without blocking in case the request cannot be satisfied immediately.

## Tags

[xthread-nowait](#), [switch-primary](#)

6.88.3.9 `rt_sem_unbind()`

```
int rt_sem_unbind (
    RT_SEM * sem )
```

Unbind from a semaphore.

## Parameters

<i>sem</i>	The semaphore descriptor.
------------	---------------------------

This routine releases a previous binding to a semaphore. After this call has returned, the descriptor is no more valid for referencing this object.

## Tags

[thread-unrestricted](#)

6.88.3.10 `rt_sem_v()`

```
int rt_sem_v (
    RT_SEM * sem )
```

Signal a semaphore.

If the semaphore is pended, the task heading the wait queue is immediately unblocked. Otherwise, the semaphore count is incremented by one, unless the semaphore is used in "pulse" mode (see [rt\\_sem\\_create\(\)](#)).

## Parameters

<i>sem</i>	The semaphore descriptor.
------------	---------------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *sem* is not a valid semaphore descriptor.


## Tags

[unrestricted](#)

## 6.89 Task management services

Services dealing with preemptive multi-tasking.

Collaboration diagram for Task management services:



### Data Structures

- struct [RT\\_TASK\\_INFO](#)  
*Task status descriptor.*

### Macros

- #define [T\\_LOPRIO](#) 0  
*Task priorities.*
- #define [T\\_LOCK](#) \_\_THREAD\_M\_LOCK  
*Task mode bits.*
- #define [T\\_WARNSW](#) \_\_THREAD\_M\_WARNSW  
*Cobalt only, nop over Mercury.*

### Functions

- int [rt\\_task\\_delete](#) (RT\_TASK \*task)  
*Delete a real-time task.*
- int [rt\\_task\\_set\\_affinity](#) (RT\_TASK \*task, const cpu\_set\_t \*cpus)  
*Set CPU affinity of real-time task.*
- int [rt\\_task\\_start](#) (RT\_TASK \*task, void(\*entry)(void \*arg), void \*arg)  
*Start a real-time task.*
- int [rt\\_task\\_shadow](#) (RT\_TASK \*task, const char \*name, int prio, int mode)  
*Turn caller into a real-time task.*
- int [rt\\_task\\_join](#) (RT\_TASK \*task)
- int [rt\\_task\\_wait\\_period](#) (unsigned long \*overruns\_r)  
*Wait on the termination of a real-time task.*
- int [rt\\_task\\_wait\\_period](#) (unsigned long \*overruns\_r)  
*Wait for the next periodic release point.*



- int `rt_task_sleep` (RTIME delay)
 

*Delay the current real-time task (with relative delay).*
- int `rt_task_sleep_until` (RTIME date)
 

*Delay the current real-time task (with absolute wakeup date).*
- int `rt_task_same` (RT\_TASK \*task1, RT\_TASK \*task2)
 

*Compare real-time task descriptors.*
- int `rt_task_suspend` (RT\_TASK \*task)
 

*Suspend a real-time task.*
- int `rt_task_resume` (RT\_TASK \*task)
 

*Resume a real-time task.*
- RT\_TASK \* `rt_task_self` (void)
 

*Retrieve the current task descriptor.*
- int `rt_task_set_priority` (RT\_TASK \*task, int prio)
 

*Change the base priority of a real-time task.*
- int `rt_task_set_mode` (int clrmask, int setmask, int \*mode\_r)
 

*Change the current task mode.*
- int `rt_task_yield` (void)
 

*Manual round-robin.*
- int `rt_task_unblock` (RT\_TASK \*task)
 

*Unblock a real-time task.*
- int `rt_task_slice` (RT\_TASK \*task, RTIME quantum)
 

*Set a task's round-robin quantum.*
- int `rt_task_inquire` (RT\_TASK \*task, RT\_TASK\_INFO \*info)
 

*Retrieve information about a real-time task.*
- ssize\_t `rt_task_send_timed` (RT\_TASK \*task, RT\_TASK\_MCB \*mcb\_s, RT\_TASK\_MCB \*mcb\_r, const struct timespec \*abs\_timeout)
 

*Send a message to a real-time task.*
- static ssize\_t `rt_task_send_until` (RT\_TASK \*task, RT\_TASK\_MCB \*mcb\_s, RT\_TASK\_MCB \*mcb\_r, RTIME timeout)
 

*Send a message to a real-time task (with absolute scalar timeout).*
- static ssize\_t `rt_task_send` (RT\_TASK \*task, RT\_TASK\_MCB \*mcb\_s, RT\_TASK\_MCB \*mcb\_r, RTIME timeout)
 

*Send a message to a real-time task (with relative scalar timeout).*
- int `rt_task_receive_timed` (RT\_TASK\_MCB \*mcb\_r, const struct timespec \*abs\_timeout)
 

*Receive a message from a real-time task.*
- static int `rt_task_receive_until` (RT\_TASK\_MCB \*mcb\_r, RTIME timeout)
 

*Receive a message from a real-time task (with absolute scalar timeout).*

- static int `rt_task_receive` (RT\_TASK\_MCB \*mcb\_r, RTIME timeout)

*Receive a message from a real-time task (with relative scalar timeout).*

- int `rt_task_reply` (int flowid, RT\_TASK\_MCB \*mcb\_s)

*Reply to a remote task message.*

- int `rt_task_bind` (RT\_TASK \*task, const char \*name, RTIME timeout)

*Bind to a task.*

- int `rt_task_unbind` (RT\_TASK \*task)

*Unbind from a task.*

- int `rt_task_create` (RT\_TASK \*task, const char \*name, int stksize, int prio, int mode)

*Create a task with Alchemy personality.*

- int `rt_task_set_periodic` (RT\_TASK \*task, RTIME idate, RTIME period)

*Make a real-time task periodic.*

- int `rt_task_spawn` (RT\_TASK \*task, const char \*name, int stksize, int prio, int mode, void(\*entry)(void \*arg), void \*arg)

*Create and start a real-time task.*

### 6.89.1 Detailed Description

Services dealing with preemptive multi-tasking.

Each Alchemy task is an independent portion of the overall application code embodied in a C procedure, which executes on its own stack context.

### 6.89.2 Macro Definition Documentation

#### 6.89.2.1 T\_LOCK

```
#define T_LOCK __THREAD_M_LOCK
```

Task mode bits.

Referenced by `rt_task_create()`.

#### 6.89.2.2 T\_LOPRIO

```
#define T_LOPRIO 0
```

Task priorities.

## 6.89.2.3 T\_WARNSW

```
#define T_WARNSW __THREAD_M_WARNSW
```

Cobalt only, nop over Mercury.

Referenced by `rt_task_create()`.

## 6.89.3 Function Documentation

6.89.3.1 `rt_task_bind()`

```
int rt_task_bind (
    RT_TASK * task,
    const char * name,
    RTIME timeout )
```

Bind to a task.

This routine creates a new descriptor to refer to an existing Alchemy task identified by its symbolic name. If the object does not exist on entry, the caller may block until a task of the given name is created.

Parameters

<i>task</i>	The address of a task descriptor filled in by the operation. Contents of this memory is undefined upon failure.
<i>name</i>	A valid NULL-terminated name which identifies the task to bind to. This string should match the object name argument passed to <code>rt_task_create()</code> , or <code>rt_task_shadow()</code> .
<i>timeout</i>	The number of clock ticks to wait for the registration to occur (see note). Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until the object is registered. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if the object is not registered on entry.

Returns

Zero is returned upon success. Otherwise:

- `-EINTR` is returned if `rt_task_unblock()` was called for the current task before the retrieval has completed.
- `-EWOULDBLOCK` is returned if `timeout` is equal to `TM_NONBLOCK` and the searched object is not registered on entry.
- `-ETIMEDOUT` is returned if the object cannot be retrieved within the specified amount of time.
- `-EPERM` is returned if this service should block, but was not called from a Xenomai thread.

Tags

[xthread-nowait](#), [switch-primary](#)

## Note

The *timeout* value is interpreted as a multiple of the Alchemy clock resolution (see `--alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.89.3.2 `rt_task_create()`

```
int rt_task_create (
    RT_TASK * task,
    const char * name,
    int stksize,
    int prio,
    int mode )
```

Create a task with Alchemy personality.

This service creates a task with access to the full set of Alchemy services. If *prio* is non-zero, the new task belongs to Xenomai's real-time FIFO scheduling class, aka `SCHED_FIFO`. If *prio* is zero, the task belongs to the regular `SCHED_OTHER` class.

Creating tasks with zero priority is useful for running non real-time processes which may invoke blocking real-time services, such as pending on a semaphore, reading from a message queue or a buffer, and so on.

Once created, the task is left dormant until it is actually started by [rt\\_task\\_start\(\)](#).

## Parameters

<i>task</i>	The address of a task descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the task. When non-NULL and non-empty, a copy of this string is used for indexing the created task into the object registry.
<i>stksize</i>	The size of the stack (in bytes) for the new task. If zero is passed, a system-dependent default size will be substituted.
<i>prio</i>	The base priority of the new task. This value must be in the [0 .. 99] range, where 0 is the lowest effective priority.
<i>mode</i>	The task creation mode. The following flags can be OR'ed into this bitmask:

- `T_JOINABLE` allows another task to wait on the termination of the new task. [rt\\_task\\_join\(\)](#) shall be called for this task to clean up any resources after its termination.
- `T_LOCK` causes the new task to lock the scheduler prior to entering the user routine specified by [rt\\_task\\_start\(\)](#). A call to [rt\\_task\\_set\\_mode\(\)](#) from the new task is required to drop this lock.
- When running over the Cobalt core, `T_WARNSW` causes the `SIGDEBUG` signal to be sent to the current task whenever it switches to the secondary mode. This feature is useful to detect unwanted migrations to the Linux domain. This flag has no effect over the Mercury core.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if either *prio*, *mode* or *stksize* are invalid.
- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the task.
- -EEXIST is returned if the *name* is conflicting with an already registered task.
- -EPERM is returned if this service was called from an invalid context, e.g. interrupt or non-Xenomai thread.

## Tags

[xthread-only](#), [mode-unrestricted](#), [switch-secondary](#)

## Side effects

- When running over the Cobalt core:
  - calling [rt\\_task\\_create\(\)](#) causes SCHED\_FIFO tasks to switch to secondary mode.
  - members of Xenomai's SCHED\_FIFO class running in the primary domain have utmost priority over all Linux activities in the system, including Linux interrupt handlers.
- When running over the Mercury core, the new task belongs to the regular POSIX SCHED\_FIFO class.

## Note

Tasks can be referred to from multiple processes which all belong to the same Xenomai session.

## Examples:

[cross-link.c](#).

References [T\\_LOCK](#), and [T\\_WARNSW](#).

6.89.3.3 [rt\\_task\\_delete\(\)](#)

```
int rt_task_delete (
    RT_TASK * task )
```

Delete a real-time task.

This call terminates a task previously created by [rt\\_task\\_create\(\)](#).

Tasks created with the [T\\_JOINABLE](#) flag shall be joined by a subsequent call to [rt\\_task\\_join\(\)](#) once successfully deleted, to reclaim all resources.

## Parameters

<i>task</i>	The task descriptor.
-------------	----------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *task* is not a valid task descriptor.
- -EPERM is returned if *task* is NULL and this service was called from an invalid context. In addition, this error is always raised when this service is called from asynchronous context, such as a timer/alarm handler.

## Tags

[mode-unrestricted](#), [switch-secondary](#)

## Note

The caller must be an Alchemy task if *task* is NULL.

## Examples:

[cross-link.c](#).

6.89.3.4 `rt_task_inquire()`

```
int rt_task_inquire (
    RT_TASK * task,
    RT_TASK_INFO * info )
```

Retrieve information about a real-time task.

Return various information about an Alchemy task. This service may also be used to probe for task existence.

## Parameters

<i>task</i>	The task descriptor. If <i>task</i> is NULL, the information about the current task is returned.
<i>info</i>	The address of a structure the task information will be written to. Passing NULL is valid, in which case the system is only probed for existence of the specified task.

## Returns

Zero is returned if the task exists. In addition, if *info* is non-NULL, it is filled in with task information.

- -EINVAL is returned if *task* is not a valid task descriptor, or if *prio* is invalid.
- -EPERM is returned if *task* is NULL and this service was called from an invalid context.

## Tags

[mode-unrestricted](#), [switch-primary](#)

## Note

The caller must be an Alchemy task if *task* is NULL.

6.89.3.5 `rt_task_join()`

```
int rt_task_join (
    RT_TASK * task )
```

Wait on the termination of a real-time task.

This service blocks the caller in non-real-time context until *task* has terminated. All resources are released after successful completion of this service.

The specified task must have been created by the same process that wants to join it, and the T\_JOINABLE mode flag must have been set on creation to [rt\\_task\\_create\(\)](#).

Tasks created with the T\_JOINABLE flag shall be joined by a subsequent call to [rt\\_task\\_join\(\)](#) once successfully deleted, to reclaim all resources.

## Parameters

<i>task</i>	The task descriptor.
-------------	----------------------

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *task* is not a valid task descriptor.
- -EINVAL is returned if the task was not created with T\_JOINABLE set or some other task is already waiting on the termination.
- -EDEADLK is returned if *task* refers to the caller.
- -ESRCH is returned if *task* no longer exists or refers to task created by a different process.

## Tags

[mode-unrestricted](#), [switch-primary](#)

## Note

After successful completion of this service, it is neither required nor valid to additionally invoke [rt\\_task\\_delete\(\)](#) on the same task.

6.89.3.6 `rt_task_receive()`

```
ssize_t rt_task_receive (
    RT_TASK_MCB * mcb_r,
    RTIME timeout ) [inline], [static]
```

Receive a message from a real-time task (with relative scalar timeout).

This routine is a variant of [rt\\_task\\_receive\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

## Parameters

<i>mcb_r</i>	The address of a message control block referring to the receive message area.
<i>timeout</i>	A delay expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until a remote task eventually sends a message. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if no remote task is currently waiting for sending a message.

## Tags

[xthread-only](#), [switch-primary](#)

6.89.3.7 `rt_task_receive_timed()`

```
int rt_task_receive_timed (
    RT_TASK_MCB * mcb_r,
    const struct timespec * abs_timeout )
```

Receive a message from a real-time task.

This service is part of the synchronous message passing support available to Alchemy tasks. The caller receives a variable-sized message from another task. The sender is blocked until the caller invokes [rt\\_task\\_reply\(\)](#) to finish the transaction.

A basic message control block is used to store the location and size of the data area to receive from the client, in addition to a user-defined operation code.



## Parameters

<i>mcb_r</i>	The address of a message control block referring to the receive message area. The fields from this control block should be set as follows:
--------------	--

- *mcb\_r->data* should contain the address of a buffer large enough to collect the data sent by the remote task;
- *mcb\_r->size* should contain the size in bytes of the buffer space pointed at by *mcb\_r->data*. If *mcb\_r->size* is lower than the actual size of the received message, no data copy takes place and `-ENOBUFS` is returned to the caller. See note.

Upon return, *mcb\_r->opcode* will contain the operation code sent from the remote task using [rt\\_task\\_send\(\)](#).

## Parameters

<i>abs_timeout</i>	The number of clock ticks to wait for receiving a message (see note). Passing <code>NULL</code> causes the caller to block indefinitely until a remote task eventually sends a message. Passing <code>{ .tv_sec = 0, .tv_nsec = 0 }</code> causes the service to return immediately without waiting if no remote task is currently waiting for sending a message.
--------------------	---

## Returns

A strictly positive value is returned upon success, representing a flow identifier for the opening transaction; this token should be passed to [rt\\_task\\_reply\(\)](#), in order to send back a reply to and unblock the remote task appropriately. Otherwise:

- `-EPERM` is returned if this service was called from an invalid context.
- `-EINTR` is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task before a message was received.
- `-ENOBUFS` is returned if *mcb\_r* does not point at a message area large enough to collect the remote task's message.
- `-EWOULDBLOCK` is returned if *abs\_timeout* is `{ .tv_sec = 0, .tv_nsec = 0 }` and no remote task is currently waiting for sending a message to the caller.
- `-ETIMEDOUT` is returned if no message was received within the *timeout*.

## Tags

[xthread-only](#), [switch-primary](#)

## Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.89.3.8 `rt_task_receive_until()`

```
ssize_t rt_task_receive_until (
    RT_TASK_MCB * mcb_r,
    RTIME abs_timeout ) [inline], [static]
```

Receive a message from a real-time task (with absolute scalar timeout).

This routine is a variant of [rt\\_task\\_receive\\_timed\(\)](#) accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>mcb_r</i>	The address of a message control block referring to the receive message area.
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until a remote task eventually sends a message. Passing <code>TM_NONBLOCK</code> causes the service to return immediately without waiting if no remote task is currently waiting for sending a message.

## Tags

[xthread-only](#), [switch-primary](#)

6.89.3.9 `rt_task_reply()`

```
int rt_task_reply (
    int flowid,
    RT_TASK_MCB * mcb_s )
```

Reply to a remote task message.

This service is part of the synchronous message passing support available to Alchemy tasks. The caller sends a variable-sized message back to a remote task, in response to this task's initial message received by a call to [rt\\_task\\_receive\(\)](#). As a consequence of calling [rt\\_task\\_reply\(\)](#), the remote task will be unblocked from the [rt\\_task\\_send\(\)](#) service.

A basic message control block is used to store the location and size of the data area to send back, in addition to a user-defined status code.

## Parameters

<i>flowid</i>	The flow identifier returned by a previous call to <a href="#">rt_task_receive()</a> which uniquely identifies the current transaction.
<i>mcb_s</i>	The address of an optional message control block referring to the message to be sent back. If <i>mcb_s</i> is NULL, the remote will be unblocked without getting any reply data. When <i>mcb_s</i> is valid, the fields from this control block should be set as follows:

- `mcb_s->data` should contain the address of the payload data to send to the remote task.
- `mcb_s->size` should contain the size in bytes of the payload data pointed at by `mcb_s->data`. Zero is a legitimate value, and indicates that no payload data will be transferred. In the latter case, `mcb_s->data` will be ignored.
- `mcb_s->opcode` is an opaque status code carried during the message transfer the caller can fill with any appropriate value. It will be made available "as is" to the remote task into the status code field by the `rt_task_send()` service. If `mcb_s` is NULL, Zero will be returned to the remote task into the status code field.

#### Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if `flowid` is invalid.
- -ENXIO is returned if `flowid` does not match the expected identifier returned from the latest call of the current task to `rt_task_receive()`, or if the remote task stopped waiting for the reply in the meantime (e.g. the remote could have been deleted or forcibly unblocked).
- -ENOBUFS is returned if the reply data referred to by `mcb_s` is larger than the reply area mentioned by the remote task when calling `rt_task_send()`. In such a case, the remote task also receives `-ENOBUFS` on return from `rt_task_send()`.
- -EPERM is returned if this service was called from an invalid context.

#### Tags

[xthread-only](#), [switch-primary](#)

#### 6.89.3.10 `rt_task_resume()`

```
int rt_task_resume (
    RT_TASK * task )
```

Resume a real-time task.

Forcibly resume the execution of a task which was previously suspended by a call to `rt_task_suspend()`, if the suspend nesting count decrements to zero.

#### Parameters

<code>task</code>	The task descriptor.
-------------------	----------------------

#### Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if `task` is not a valid task descriptor.

Tags

[unrestricted](#), [switch-primary](#)

Note

Blocked and suspended task states are cumulative. Therefore, resuming a task currently waiting on a synchronization object (e.g. semaphore, queue) does not make it eligible for scheduling until the awaited resource is eventually acquired, or a timeout elapses.

#### 6.89.3.11 `rt_task_same()`

```
int rt_task_same (
    RT_TASK * task1,
    RT_TASK * task2 )
```

Compare real-time task descriptors.

This predicate returns true if *task1* and *task2* refer to the same task.

Parameters

<i>task1</i>	First task descriptor to compare.
<i>task2</i>	Second task descriptor to compare.

Returns

A non-zero value is returned if both descriptors refer to the same task, zero otherwise.

Tags

[unrestricted](#)

#### 6.89.3.12 `rt_task_self()`

```
RT_TASK * rt_task_self (
    void )
```

Retrieve the current task descriptor.

Return the address of the current Alchemy task descriptor.

## Returns

The address of the task descriptor referring to the current Alchemy task is returned upon success, or NULL if not called from a valid Alchemy task context.

## Tags

[xthread-only](#)

6.89.3.13 `rt_task_send()`

```
ssize_t rt_task_send (
    RT_TASK * task,
    RT_TASK_MCB * mcb_s,
    RT_TASK_MCB * mcb_r,
    RTIME timeout ) [inline], [static]
```

Send a message to a real-time task (with relative scalar timeout).

This routine is a variant of [rt\\_task\\_send\\_timed\(\)](#) accepting a relative timeout specification expressed as a scalar value.

## Parameters

<i>task</i>	The task descriptor.
<i>mcb_s</i>	The address of the message control block referring to the message to be sent.
<i>mcb_r</i>	The address of an optional message control block referring to the reply message area.
<i>timeout</i>	A delay expressed in clock ticks. Passing TM_INFINITE causes the caller to block indefinitely until a reply is received. Passing TM_NONBLOCK causes the service to return without blocking in case the recipient task is not waiting for messages at the time of the call.

## Tags

[xthread-only](#), [switch-primary](#)

6.89.3.14 `rt_task_send_timed()`

```
ssize_t rt_task_send_timed (
    RT_TASK * task,
    RT_TASK_MCB * mcb_s,
    RT_TASK_MCB * mcb_r,
    const struct timespec * abs_timeout )
```

Send a message to a real-time task.

This service is part of the synchronous message passing support available to Alchemy tasks. The caller sends a variable-sized message to another task, waiting for the remote to receive the initial message by a call to `rt_task_receive()`, then reply to it using `rt_task_reply()`.

A basic message control block is used to store the location and size of the data area to send or retrieve upon reply, in addition to a user-defined operation code.

Parameters

<i>task</i>	The task descriptor.
<i>mcb_s</i>	The address of the message control block referring to the message to be sent. The fields from this control block should be set as follows:

- *mcb\_s->data* should contain the address of the payload data to send to the remote task.
- *mcb\_s->size* should contain the size in bytes of the payload data pointed at by *mcb\_s->data*. Zero is a legitimate value, and indicates that no payload data will be transferred. In the latter case, *mcb\_s->data* will be ignored.
- *mcb\_s->opcode* is an opaque operation code carried during the message transfer, the caller can fill with any appropriate value. It will be made available "as is" to the remote task into the operation code field by the `rt_task_receive()` service.

Parameters

<i>mcb_r</i>	The address of an optional message control block referring to the reply message area. If <i>mcb_r</i> is NULL and a reply is sent back by the remote task, the reply message will be discarded, and -ENOBUFS will be returned to the caller. When <i>mcb_r</i> is valid, the fields from this control block should be set as follows:
--------------	---

- *mcb\_r->data* should contain the address of a buffer large enough to collect the reply data from the remote task.
- *mcb\_r->size* should contain the size in bytes of the buffer space pointed at by *mcb\_r->data*. If *mcb\_r->size* is lower than the actual size of the reply message, no data copy takes place and -ENOBUFS is returned to the caller.

Upon return, *mcb\_r->opcode* will contain the status code sent back from the remote task using `rt_task_reply()`, or zero if unspecified.

Parameters

<i>abs_timeout</i>	An absolute date expressed in clock ticks, specifying a time limit to wait for the recipient task to reply to the initial message (see note). Passing NULL causes the caller to block indefinitely until a reply is received. Passing { .tv_sec = 0, .tv_nsec = 0 } causes the service to return without blocking in case the recipient task is not waiting for messages at the time of the call.
--------------------	---

## Returns

A positive value is returned upon success, representing the length (in bytes) of the reply message returned by the remote task. Zero is a success status, meaning either that *mcb\_r* was NULL on entry, or that no actual message was passed to the remote call to [rt\\_task\\_reply\(\)](#). Otherwise:

- -EINVAL is returned if *task* is not a valid task descriptor.
- -EPERM is returned if this service was called from an invalid context.
- -ENOBUFS is returned if *mcb\_r* does not point at a message area large enough to collect the remote task's reply. This includes the case where *mcb\_r* is NULL on entry, despite the remote task attempts to send a reply message.
- -EWOULDBLOCK is returned if *abs\_timeout* is { .tv\_sec = 0, .tv\_nsec = 0 } and the recipient *task* is not currently waiting for a message on the [rt\\_task\\_receive\(\)](#) service.
- -EIDRM is returned if *task* has been deleted while waiting for a reply.
- -EINTR is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task before any reply was received from the recipient *task*.

## Tags

[xthread-only](#), [switch-primary](#)

## Note

*abs\_timeout* is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.89.3.15 [rt\\_task\\_send\\_until\(\)](#)

```
ssize_t rt_task_send_until (
    RT_TASK * task,
    RT_TASK_MCB * mcb_s,
    RT_TASK_MCB * mcb_r,
    RTIME abs_timeout ) [inline], [static]
```

Send a message to a real-time task (with absolute scalar timeout).

This routine is a variant of [rt\\_task\\_send\\_timed\(\)](#) accepting an absolute timeout specification expressed as a scalar value.

## Parameters

<i>task</i>	The task descriptor.
<i>mcb_s</i>	The address of the message control block referring to the message to be sent.
<i>mcb_r</i>	The address of an optional message control block referring to the reply message area.
<i>abs_timeout</i>	An absolute date expressed in clock ticks. Passing <code>TM_INFINITE</code> causes the caller to block indefinitely until a reply is received. Passing <code>TM_NONBLOCK</code> causes the service to return without blocking in case the recipient task is not waiting for messages at the time of the call.

Tags

[xthread-only](#), [switch-primary](#)

### 6.89.3.16 `rt_task_set_affinity()`

```
int rt_task_set_affinity (
    RT_TASK * task,
    const cpu_set_t * cpus )
```

Set CPU affinity of real-time task.

This call makes *task* affine to the set of CPUs defined by *cpus*.

Parameters

<i>task</i>	The task descriptor. If <i>task</i> is NULL, the CPU affinity of the current task is changed.
<i>cpus</i>	The set of CPUs <i>task</i> should be affine to.

Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *task* is NULL but the caller is not a Xenomai task, or if *task* is non-NULL but not a valid task descriptor.
- -EINVAL is returned if *cpus* contains no processors that are currently physically on the system and permitted to the process according to any restrictions that may be imposed by the "cpuset" mechanism described in `cpuset(7)`.

Tags

[mode-unrestricted](#), [switch-secondary](#)

Note

The caller must be an Alchemy task if *task* is NULL.

### 6.89.3.17 `rt_task_set_mode()`

```
int rt_task_set_mode (
    int clrmask,
    int setmask,
    int * mode_r )
```



Change the current task mode.

Each Alchemy task has a set of internal flags determining several operating conditions. `rt_task_set_mode()` takes a bitmask of mode bits to clear for disabling the corresponding modes for the current task, and another one to set for enabling them. The mode bits which were previously in effect before the change can be returned upon request.

The following bits can be part of the bitmask:

- `T_LOCK` causes the current task to lock the scheduler on the current CPU, preventing all further involuntary task switches on this CPU. Clearing this bit unlocks the scheduler.
- Only when running over the Cobalt core:
  - `T_WARNSW` causes the `SIGDEBUG` signal to be sent to the current task whenever it switches to the secondary mode. This feature is useful to detect unwanted migrations to the Linux domain.
  - `T_CONFORMING` can be passed in `setmask` to switch the current Alchemy task to its preferred runtime mode. The only meaningful use of this switch is to force a real-time task back to primary mode (see note). Any other use leads to a nop.

These two last flags have no effect over the Mercury core, and are simply ignored.

Parameters

<code>clrmask</code>	A bitmask of mode bits to clear for the current task, before <code>setmask</code> is applied. Zero is an acceptable value which leads to a no-op.
<code>setmask</code>	A bitmask of mode bits to set for the current task. Zero is an acceptable value which leads to a no-op.
<code>mode↔ _r</code>	If non-NULL, <code>mode_r</code> must be a pointer to a memory location which will be written upon success with the previous set of active mode bits. If NULL, the previous set of active mode bits will not be returned.

Returns

Zero is returned upon success. Otherwise:

- `-EINVAL` is returned if `task` is not a valid task descriptor, or if any bit from `clrmask` or `setmask` is invalid.
- `-EPERM` is returned if this service was called from an invalid context.

Tags

[xthread-only](#), [switch-primary](#)

Note

The caller must be an Alchemy task.

Forcing the task mode using the `T_CONFORMING` bit from user code is almost always wrong, since the Xenomai/cobalt core handles mode switches internally when/if required. Most often, manual mode switching from applications introduces useless overhead. This mode bit is part of the API only to cover rare use cases in middleware code based on the Alchemy interface.

6.89.3.18 `rt_task_set_periodic()`

```
int rt_task_set_periodic (
    RT_TASK * task,
    RTIME idate,
    RTIME period )
```

Make a real-time task periodic.

Make a task periodic by programming its first release point and its period in the processor time line. *task* should then call [rt\\_task\\_wait\\_period\(\)](#) to sleep until the next periodic release point in the processor timeline is reached.

## Parameters

<i>task</i>	The task descriptor. If <i>task</i> is NULL, the current task is made periodic. <i>task</i> must belong to the current process.
<i>idate</i>	The initial (absolute) date of the first release point, expressed in clock ticks (see note). If <i>idate</i> is equal to TM_NOW, the current system date is used.
<i>period</i>	The period of the task, expressed in clock ticks (see note). Passing TM_INFINITE stops the task's periodic timer if enabled, then returns successfully.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *task* is NULL but the caller is not a Xenomai task, or if *task* is non-NULL but not a valid task descriptor.
- -ETIMEDOUT is returned if *idate* is different from TM\_INFINITE and represents a date in the past.

## Tags

[mode-unrestricted](#), [switch-primary](#)

## Note

The caller must be an Alchemy task if *task* is NULL.

Over Cobalt, -EINVAL is returned if *period* is different from TM\_INFINITE but shorter than the user scheduling latency value for the target system, as displayed by `/proc/xenomai/latency`.

The *idate* and *period* values are interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

## Attention

Unlike its Xenomai 2.x counterpart, [rt\\_task\\_set\\_periodic\(\)](#) will **NOT** block *task* until *idate* is reached. The first beat in the periodic timeline should be awaited for by a call to [rt\\_task\\_wait\\_period\(\)](#).

## Examples:

[cross-link.c](#).

6.89.3.19 `rt_task_set_priority()`

```
int rt_task_set_priority (
    RT_TASK * task,
    int prio )
```

Change the base priority of a real-time task.

The base priority of a task defines the relative importance of the work being done by each task, which gains control of the CPU accordingly.

Changing the base priority of a task does not affect the priority boost the target task might have obtained as a consequence of a priority inheritance undergoing.

Parameters

<i>task</i>	The task descriptor. If <i>task</i> is NULL, the priority of the current task is changed.
<i>prio</i>	The new priority. This value must range from [T_LOPRIO .. T_HIPRIO] (inclusive) where T_LOPRIO is the lowest effective priority.

Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *task* is not a valid task descriptor, or if *prio* is invalid.
- -EPERM is returned if *task* is NULL and this service was called from an invalid context.

Tags

[mode-unrestricted](#), [switch-primary](#), [switch-secondary](#)

Note

The caller must be an Alchemy task if *task* is NULL.

Assigning the same priority to a running or ready task moves it to the end of its priority group, thus causing a manual round-robin.

6.89.3.20 `rt_task_shadow()`

```
int rt_task_shadow (
    RT_TASK * task,
    const char * name,
    int prio,
    int mode )
```

Turn caller into a real-time task.

Set the calling thread personality to the Alchemy API, enabling the full set of Alchemy services. Upon success, the caller is no more a regular POSIX thread, but a Xenomai-extended thread.

If *prio* is non-zero, the new task moves to Xenomai's real-time FIFO scheduling class, aka SCHED\_FIFO. If *prio* is zero, the task moves to the regular SCHED\_OTHER class.

Running Xenomai tasks with zero priority is useful for running non real-time processes which may invoke blocking real-time services, such as pending on a semaphore, reading from a message queue or a buffer, and so on.

## Parameters

<i>task</i>	If non-NULL, the address of a task descriptor which can be later used to identify uniquely the task, upon success of this call. If NULL, no descriptor is returned.
<i>name</i>	An ASCII string standing for the symbolic name of the task. When non-NULL and non-empty, a copy of this string is used for indexing the task into the object registry.
<i>prio</i>	The base priority of the task. This value must be in the [0 .. 99] range, where 0 is the lowest effective priority.
<i>mode</i>	The task shadowing mode. The following flags can be OR'ed into this bitmask:

- T\_LOCK causes the current task to lock the scheduler before returning to the caller, preventing all further involuntary task switches on the current CPU. A call to [rt\\_task\\_set\\_mode\(\)](#) from the current task is required to drop this lock.
- When running over the Cobalt core, T\_WARNSW causes the SIGDEBUG signal to be sent to the current task whenever it switches to the secondary mode. This feature is useful to detect unwanted migrations to the Linux domain. This flag has no effect over the Mercury core.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *prio* is invalid.
- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the task extension.
- -EEXIST is returned if the *name* is conflicting with an already registered task.
- -EBUSY is returned if the caller is not a regular POSIX thread.
- -EPERM is returned if this service was called from an invalid context, e.g. interrupt handler.

## Tags

[pthread-only](#), [switch-secondary](#), [switch-primary](#)

## Side effects

Over Cobalt, if the caller is a plain POSIX thread, it is turned into a Xenomai *shadow* thread, with full access to all Cobalt services. The caller always returns from this service in primary mode.

## Note

Tasks can be referred to from multiple processes which all belong to the same Xenomai session.

## Examples:

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

6.89.3.21 `rt_task_sleep()`

```
int rt_task_sleep (
    RTIME delay )
```

Delay the current real-time task (with relative delay).

This routine is a variant of [rt\\_task\\_sleep\\_until\(\)](#) accepting a relative timeout specification.

## Parameters

<i>delay</i>	A relative delay expressed in clock ticks (see note). A zero delay causes this service to return immediately to the caller with a success status.
--------------	---

## Returns

See [rt\\_task\\_sleep\\_until\(\)](#).

## Tags

[xthread-only](#), [switch-primary](#)

## Note

The *delay* value is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

## Examples:

[cross-link.c](#), and [rtcansend.c](#).

6.89.3.22 `rt_task_sleep_until()`

```
int rt_task_sleep_until (
    RTIME date )
```

Delay the current real-time task (with absolute wakeup date).

Delay the execution of the calling task until a given date is reached. The caller is put to sleep, and does not consume any CPU time in such a state.

## Parameters

<i>date</i>	An absolute date expressed in clock ticks, specifying a wakeup date (see note). As a special case, <code>TM_INFINITE</code> is an acceptable value that causes the caller to block indefinitely, until <a href="#">rt_task_unblock()</a> is called against it. Otherwise, any wake up date in the past causes the task to return immediately.
-------------	---

## Returns

Zero is returned upon success. Otherwise:

- `-EINTR` is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task.
- `-ETIMEDOUT` is returned if *date* has already elapsed.

- -EPERM is returned if this service was called from an invalid context.

## Tags

[xthread-only](#), [switch-primary](#)

## Note

The caller must be an Alchemy task if *task* is NULL.

The *date* value is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.89.3.23 `rt_task_slice()`

```
int rt_task_slice (
    RT_TASK * task,
    RTIME quantum )
```

Set a task's round-robin quantum.

Set the time credit allotted to a task undergoing the round-robin scheduling. If *quantum* is non-zero, `rt_task_slice()` also refills the current quantum for the target task, otherwise, time-slicing is stopped for that task.

In other words, `rt_task_slice()` should be used to toggle round-robin scheduling for an Alchemy task.

## Parameters

<i>task</i>	The task descriptor. If <i>task</i> is NULL, the time credit of the current task is changed. <i>task</i> must belong to the current process.
<i>quantum</i>	The round-robin quantum for the task expressed in clock ticks (see note).

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *task* is not a valid task descriptor, or if *prio* is invalid.
- -EPERM is returned if *task* is NULL and this service was called from an invalid context.

## Tags

[mode-unrestricted](#), [switch-primary](#)

## Note

The caller must be an Alchemy task if *task* is NULL.

The *quantum* value is interpreted as a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

6.89.3.24 `rt_task_spawn()`

```
int rt_task_spawn (
    RT_TASK * task,
    const char * name,
    int stksize,
    int prio,
    int mode,
    void(*)(void *arg) entry,
    void * arg )
```

Create and start a real-time task.

This service spawns a task by combining calls to [rt\\_task\\_create\(\)](#) and [rt\\_task\\_start\(\)](#) for the new task.

## Parameters

<i>task</i>	The address of a task descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the task. When non-NULL and non-empty, a copy of this string is used for indexing the created task into the object registry.
<i>stksize</i>	The size of the stack (in bytes) for the new task. If zero is passed, a system-dependent default size will be substituted.
<i>prio</i>	The base priority of the new task. This value must be in the [0 .. 99] range, where 0 is the lowest effective priority.
<i>mode</i>	The task creation mode. See <a href="#">rt_task_create()</a> .
<i>entry</i>	The address of the task entry point.
<i>arg</i>	A user-defined opaque argument <i>entry</i> will receive.

## Returns

See [rt\\_task\\_create\(\)](#).

## Tags

[mode-unrestricted](#), [switch-secondary](#)

## Side effects

see [rt\\_task\\_create\(\)](#).

6.89.3.25 `rt_task_start()`

```
int rt_task_start (
    RT_TASK * task,
    void(*)(void *arg) entry,
    void * arg )
```

Start a real-time task.

This call starts execution of a task previously created by [rt\\_task\\_create\(\)](#). This service causes the started task to leave the initial dormant state.



## Parameters

<i>task</i>	The task descriptor.
<i>entry</i>	The address of the task entry point.
<i>arg</i>	A user-defined opaque argument <i>entry</i> will receive.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *task* is not a valid task descriptor.

## Tags

[mode-unrestricted](#), [switch-primary](#)

## Note

Starting an already started task leads to a nop, returning a success status.

## Examples:

[cross-link.c](#).

6.89.3.26 `rt_task_suspend()`

```
int rt_task_suspend (
    RT_TASK * task )
```

Suspend a real-time task.

Forcibly suspend the execution of a task. This task will not be eligible for scheduling until it is explicitly resumed by a call to [rt\\_task\\_resume\(\)](#). In other words, the suspended state caused by a call to [rt\\_task\\_suspend\(\)](#) is cumulative with respect to the delayed and blocked states caused by other services, and is managed separately from them.

A nesting count is maintained so that [rt\\_task\\_suspend\(\)](#) and [rt\\_task\\_resume\(\)](#) must be used in pairs.

Receiving a Linux signal causes the suspended task to resume immediately.

## Parameters

<i>task</i>	The task descriptor. If <i>task</i> is NULL, the current task is suspended.
-------------	---

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *task* is NULL but the caller is not a Xenomai task, or if *task* is non-NULL but not a valid task descriptor.
- -EINTR is returned if a Linux signal has been received by the caller if suspended.
- -EPERM is returned if *task* is NULL and this service was called from an invalid context.

## Tags

[mode-unrestricted](#), [switch-primary](#)

## Note

The caller must be an Alchemy task if *task* is NULL.

Blocked and suspended task states are cumulative. Therefore, suspending a task currently waiting on a synchronization object (e.g. semaphore, queue) holds its execution until it is resumed, despite the awaited resource may have been acquired, or a timeout has elapsed in the meantime.

6.89.3.27 `rt_task_unbind()`

```
int rt_task_unbind (  
    RT_TASK * task )
```

Unbind from a task.

## Parameters

<i>task</i>	The task descriptor.
-------------	----------------------

This routine releases a previous binding to an Alchemy task. After this call has returned, the descriptor is no more valid for referencing this object.

## Tags

[thread-unrestricted](#)

6.89.3.28 `rt_task_unblock()`

```
int rt_task_unblock (  
    RT_TASK * task )
```

Unblock a real-time task.

Break the task out of any wait it is currently in. This call clears all delay and/or resource wait condition for the target task.

However, `rt_task_unblock()` does not resume a task which has been forcibly suspended by a previous call to `rt_task_suspend()`. If all suspensive conditions are gone, the task becomes eligible anew for scheduling.

Parameters

<code>task</code>	The task descriptor.
-------------------	----------------------

Returns

Zero is returned upon success. Otherwise:

- `-EINVAL` is returned if `task` is not a valid task descriptor.

Tags

[unrestricted](#), [switch-primary](#)

### 6.89.3.29 `rt_task_wait_period()`

```
int rt_task_wait_period (
    unsigned long * overruns_r )
```

Wait for the next periodic release point.

Delay the current task until the next periodic release point is reached. The periodic timer should have been previously started for `task` by a call to `rt_task_set_periodic()`.

Parameters

<code>overruns_r</code>	If non-NULL, <code>overruns_r</code> shall be a pointer to a memory location which will be written with the count of pending overruns. This value is written to only when <code>rt_task_wait_period()</code> returns <code>-ETIMEDOUT</code> or success. The memory location remains unmodified otherwise. If NULL, this count will not be returned.
-------------------------	--

Returns

Zero is returned upon success. If `overruns_r` is non-NULL, zero is written to the pointed memory location. Otherwise:

- `-EWOULDBLOCK` is returned if `rt_task_set_periodic()` was not called for the current task.

- -EINTR is returned if `rt_task_unblock()` was called for the waiting task before the next periodic release point was reached. In this case, the overrun counter is also cleared.
- -ETIMEDOUT is returned if a timer overrun occurred, which indicates that a previous release point was missed by the calling task. If `overruns_r` is non-NULL, the count of pending overruns is written to the pointed memory location.
- -EPERM is returned if this service was called from an invalid context.

#### Tags

[xthread-only](#), [switch-primary](#)

#### Note

If the current release point has already been reached at the time of the call, the current task immediately returns from this service with no delay.

#### Examples:

[cross-link.c](#).

#### 6.89.3.30 `rt_task_yield()`

```
int rt_task_yield (
    void )
```

#### Manual round-robin.

Move the current task to the end of its priority group, so that the next equal-priority task in ready state is switched in.

#### Returns

Zero is returned upon success. Otherwise:

- -EPERM is returned if this service was called from an invalid context.


#### Tags

[xthread-only](#), [switch-primary](#)

## 6.90 Timer management services

Services for reading and spinning on the hardware timer.

Collaboration diagram for Timer management services:



### Data Structures

- struct `rt_timer_info`  
*Timer status descriptor.*

### Typedefs

- typedef struct `rt_timer_info` `RT_TIMER_INFO`  
*Timer status descriptor.*

### Functions

- static `RTIME` `rt_timer_read` (void)  
*Return the current system time.*
- `SRTIME` `rt_timer_ns2ticks` (`SRTIME` ns)  
*Convert nanoseconds to Alchemy clock ticks.*
- `SRTIME` `rt_timer_ticks2ns` (`SRTIME` ticks)  
*Convert Alchemy clock ticks to nanoseconds.*
- void `rt_timer_inquire` (`RT_TIMER_INFO` \*info)  
*Inquire about the Alchemy clock.*
- void `rt_timer_spin` (`RTIME` ns)  
*Busy wait burning CPU cycles.*

#### 6.90.1 Detailed Description

Services for reading and spinning on the hardware timer.

## 6.90.2 Typedef Documentation

### 6.90.2.1 RT\_TIMER\_INFO

```
typedef struct rt\_timer\_info RT_TIMER_INFO
```

Timer status descriptor.

This structure reports information about the Alchemy clock, returned by a call to [rt\\_timer\\_inquire\(\)](#).

## 6.90.3 Function Documentation

### 6.90.3.1 [rt\\_timer\\_inquire\(\)](#)

```
void rt\_timer\_inquire (  
    RT\_TIMER\_INFO * info )
```

Inquire about the Alchemy clock.

Return status information about the Alchemy clock.

Parameters

<i>info</i>	The address of a <a href="#">structure</a> to fill with the clock information.
-------------	--

Tags

[unrestricted](#)

### 6.90.3.2 [rt\\_timer\\_ns2ticks\(\)](#)

```
SRTIME rt\_timer\_ns2ticks (  
    SRTIME ns )
```

Convert nanoseconds to Alchemy clock ticks.

Convert a count of nanoseconds to Alchemy clock ticks. This routine operates on signed nanosecond values. This is the converse call to [rt\\_timer\\_ticks2ns\(\)](#).

## Parameters

<i>ns</i>	The count of nanoseconds to convert.
-----------	--------------------------------------

## Returns

The corresponding value expressed in clock ticks of the Alchemy clock. The resolution of the Alchemy clock can be set using the `–alchemy-clock-resolution` option when starting the application process (defaults to 1 nanosecond).

## Tags

[unrestricted](#)

## Examples:

[cross-link.c](#), and [rtcansend.c](#).

6.90.3.3 `rt_timer_read()`

```
RTIME rt_timer_read (  
    void ) [inline], [static]
```

Return the current system time.

Return the current time maintained by the Xenomai core clock.

## Returns

The current time expressed in clock ticks (see note).

## Tags

[unrestricted](#)

## Note

The *time* value is a multiple of the Alchemy clock resolution (see `–alchemy-clock-resolution` option, defaults to 1 nanosecond).

## Examples:

[cross-link.c](#).

6.90.3.4 `rt_timer_spin()`

```
void rt_timer_spin (  
    RTIME ns )
```

Busy wait burning CPU cycles.

Enter a busy waiting loop for a count of nanoseconds.

Since this service is always called with interrupts enabled, the caller might be preempted by other real-time activities, therefore the actual delay might be longer than specified.

## Parameters

<i>ns</i>	The time to wait expressed in nanoseconds.
-----------	--

## Tags

[unrestricted](#)

6.90.3.5 `rt_timer_ticks2ns()`

```
SRTIME rt_timer_ticks2ns (  
    SRTIME ns )
```

Convert Alchemy clock ticks to nanoseconds.

Convert a count of Alchemy clock ticks to nanoseconds. This routine operates on signed nanosecond values. This is the converse call to [rt\\_timer\\_ns2ticks\(\)](#).

## Parameters

<i>ns</i>	The count of nanoseconds to convert.
-----------	--------------------------------------

## Returns

The corresponding value expressed in nanoseconds. The resolution of the Alchemy clock can be set using the `-alchemy-clock-resolution` option when starting the application process (defaults to 1 nanosecond).

## Tags

[unrestricted](#)



## 6.91 VxWorks® emulator

A VxWorks® emulation library on top of Xenomai.

A VxWorks® emulation library on top of Xenomai.

The emulator mimicks the behavior described in the public documentation of the WIND 5.x API for the following class of services:

- taskLib, taskInfoLib, taskHookLib,
- semLib, msgQLib, wdLib, memPartLib
- intLib, tickLib, sysLib (partial)
- errnoLib, lstLib, kernelLib (partial)

## 6.92 pSOS® emulator

A pSOS® emulation library on top of Xenomai.

A pSOS® emulation library on top of Xenomai.

The emulator mimicks the behavior described in the public documentation of the pSOS 2.x API for the following class of services:

- Tasks, Events, Queues, Semaphores
- Partitions, Regions, Timers

## 6.93 Transition Kit

A set of wrappers and services easing the transition from Xenomai 2.x to 3.x.

- int `COMPAT__rt_task_create` (RT\_TASK \*task, const char \*name, int stksize, int prio, int mode)
  - Create a real-time task (compatibility service).*
- int `COMPAT__rt_task_set_periodic` (RT\_TASK \*task, RTIME idate, RTIME period)
  - Make a real-time task periodic (compatibility service).*
- int `COMPAT__rt_alarm_create` (RT\_ALARM \*alarm, const char \*name)
  - Create an alarm object (compatibility service).*
- int `rt_alarm_wait` (RT\_ALARM \*alarm)
  - Wait for the next alarm shot (compatibility service).*
- int `COMPAT__rt_event_create` (RT\_EVENT \*event, const char \*name, unsigned long ivalue, int mode)
  - Create an event flag group.*
- int `COMPAT__rt_event_signal` (RT\_EVENT \*event, unsigned long mask)
  - Signal an event.*
- int `COMPAT__rt_event_clear` (RT\_EVENT \*event, unsigned long mask, unsigned long \*mask\_r)
  - Clear event flags.*
- int `COMPAT__rt_pipe_create` (RT\_PIPE \*pipe, const char \*name, int minor, size\_t poolsize)
  - Create a message pipe.*
- int `pthread_make_periodic_np` (pthread\_t thread, struct timespec \*starttp, struct timespec \*periodtp)
  - Make a thread periodic (compatibility service).*
- int `pthread_wait_np` (unsigned long \*overruns\_r)
  - Wait for the next periodic release point (compatibility service)*

### 6.93.1 Detailed Description

A set of wrappers and services easing the transition from Xenomai 2.x to 3.x.

This interface provides a source compatibility layer for building applications based on the Xenomai 2.x *posix* and *native* APIs over Xenomai 3.x.

### 6.93.2 Function Documentation

## 6.93.2.1 COMPAT\_\_rt\_alarm\_create()

```
int COMPAT__rt_alarm_create (
    RT_ALARM * alarm,
    const char * name )
```

Create an alarm object (compatibility service).

This routine creates an object triggering an alarm routine at a specified time in the future. Alarms can be periodic or oneshot, depending on the reload interval value passed to [rt\\_alarm\\_start\(\)](#). A task can wait for timeouts using the [rt\\_alarm\\_wait\(\)](#) service.

## Parameters

<i>alarm</i>	The address of an alarm descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the alarm. When non-NULL and non-empty, a copy of this string is used for indexing the created alarm into the object registry.

## Returns

Zero is returned upon success. Otherwise:

- -ENOMEM is returned if the system fails to get memory from the local pool in order to create the alarm.
- -EEXIST is returned if the *name* is conflicting with an already registered alarm.
- -EPERM is returned if this service was called from an asynchronous context.

## Tags

[thread-unrestricted](#), [switch-secondary](#)

## Note

Alarms are process-private objects and thus cannot be shared by multiple processes, even if they belong to the same Xenomai session.

**Deprecated** This is a compatibility service from the Transition Kit.

## 6.93.2.2 COMPAT\_\_rt\_event\_clear()

```
int COMPAT__rt_event_clear (
    RT_EVENT * event,
    unsigned long mask,
    unsigned long * mask_r )
```

Clear event flags.

This call is the legacy form of the [rt\\_event\\_clear\(\)](#) service, using a long event mask. The new form uses a regular integer to hold the event mask instead.

## Parameters

<i>event</i>	The event descriptor.
<i>mask</i>	The set of event flags to be cleared.
<i>mask</i> ↔ <i>_r</i>	If non-NULL, <i>mask_r</i> is the address of a memory location which will receive the previous value of the event flag group before the flags are cleared.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *event* is not a valid event flag group descriptor.

## Tags

unrestricted, switch-primary

**Deprecated** This is a compatibility service from the Transition Kit.

## 6.93.2.3 COMPAT\_\_rt\_event\_create()

```
int COMPAT__rt_event_create (
    RT_EVENT * event,
    const char * name,
    unsigned long ivalue,
    int mode )
```

Create an event flag group.

This call is the legacy form of the [rt\\_event\\_create\(\)](#) service, using a long event mask. The new form uses a regular integer to hold the event mask instead.

## Parameters

<i>event</i>	The address of an event descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the event. When non-NULL and non-empty, a copy of this string is used for indexing the created event into the object registry.
<i>ivalue</i>	The initial value of the group's event mask.
<i>mode</i>	The event group creation mode. The following flags can be OR'ed into this bitmask:

- EV\_FIFO makes tasks pend in FIFO order on the event flag group.
- EV\_PRIO makes tasks pend in priority order on the event flag group.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *mode* is invalid.
- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the event flag group.
- -EEXIST is returned if the *name* is conflicting with an already registered event flag group.
- -EPERM is returned if this service was called from an asynchronous context.

## Tags

[thread-unrestricted](#), [switch-secondary](#)

## Note

Event flag groups can be shared by multiple processes which belong to the same Xenomai session.

**Deprecated** This is a compatibility service from the Transition Kit.

## 6.93.2.4 COMPAT\_\_rt\_event\_signal()

```
int COMPAT__rt_event_signal (
    RT_EVENT * event,
    unsigned long mask )
```

Signal an event.

This call is the legacy form of the [rt\\_event\\_signal\(\)](#) service, using a long event mask. The new form uses a regular integer to hold the event mask instead.

## Parameters

<i>event</i>	The event descriptor.
<i>mask</i>	The set of events to be posted.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *event* is not an event flag group descriptor.

## Tags

[unrestricted](#), [switch-primary](#)

**Deprecated** This is a compatibility service from the Transition Kit.

### 6.93.2.5 COMPAT\_\_rt\_pipe\_create()

```
int COMPAT__rt_pipe_create (
    RT_PIPE * pipe,
    const char * name,
    int minor,
    size_t poolsize )
```

Create a message pipe.

This call is the legacy form of the [rt\\_pipe\\_create\(\)](#) service, which returns a zero status upon success. The new form returns the *minor* number assigned to the connection instead, which is useful when `P_↔MINOR_AUTO` is specified in the call (see the discussion about the *minor* parameter).

This service opens a bi-directional communication channel for exchanging messages between Xenomai threads and regular Linux threads. Pipes natively preserve message boundaries, but can also be used in byte-oriented streaming mode from Xenomai to Linux.

[rt\\_pipe\\_create\(\)](#) always returns immediately, even if no thread has opened the associated special device file yet. On the contrary, the non real-time side could block upon attempt to open the special device file until [rt\\_pipe\\_create\(\)](#) is issued on the same pipe from a Xenomai thread, unless `O_NONBLOCK` was given to the `open(2)` system call.

#### Parameters

<i>pipe</i>	The address of a pipe descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the pipe. When non-NULL and non-empty, a copy of this string is used for indexing the created pipe into the object registry.

Named pipes are supported through the use of the registry. Passing a valid *name* parameter when creating a message pipe causes a symbolic link to be created from `/proc/xenomai/registry/rtipc/xddp/name` to the associated special device (i.e. `/dev/rtp*`), so that the specific *minor* information does not need to be known from those processes for opening the proper device file. In such a case, both sides of the pipe only need to agree upon a symbolic name to refer to the same data path, which is especially useful whenever the *minor* number is picked up dynamically using an adaptive algorithm, such as passing `P_MINOR_AUTO` as *minor* value.

#### Parameters

<i>minor</i>	The minor number of the device associated with the pipe. Passing <code>P_MINOR_AUTO</code> causes the minor number to be auto-allocated. In such a case, a symbolic link will be automatically created from <code>/proc/xenomai/registry/rtipc/xddp/name</code> to the allocated pipe device entry. Valid minor numbers range from 0 to <code>CONFIG_XENO_OPT_PIPE_NRDEV-1</code> .
<i>poolsize</i>	Specifies the size of a dedicated buffer pool for the pipe. Passing 0 means that all message allocations for this pipe are performed on the Cobalt core heap.

## Returns

This compatibility call returns zero upon success. Otherwise:

- -ENOMEM is returned if the system fails to get memory from the main heap in order to create the pipe.
- -ENODEV is returned if *minor* is different from P\_MINOR\_AUTO and is not a valid minor number.
- -EEXIST is returned if the *name* is conflicting with an already registered pipe.
- -EBUSY is returned if *minor* is already open.
- -EPERM is returned if this service was called from an asynchronous context.

## Tags

[thread-unrestricted](#), [switch-secondary](#)

## 6.93.2.6 COMPAT\_\_rt\_task\_create()

```
int COMPAT__rt_task_create (
    RT_TASK * task,
    const char * name,
    int stksize,
    int prio,
    int mode )
```

Create a real-time task (compatibility service).

This service creates a task with access to the full set of Xenomai real-time services.

This service creates a task with access to the full set of Xenomai real-time services. If *prio* is non-zero, the new task belongs to Xenomai's real-time FIFO scheduling class, aka SCHED\_FIFO. If *prio* is zero, the task belongs to the regular SCHED\_OTHER class.

Creating tasks with zero priority is useful for running non real-time processes which may invoke blocking real-time services, such as pending on a semaphore, reading from a message queue or a buffer, and so on.

Once created, the task is left dormant until it is actually started by [rt\\_task\\_start\(\)](#).

## Parameters

<i>task</i>	The address of a task descriptor which can be later used to identify uniquely the created object, upon success of this call.
<i>name</i>	An ASCII string standing for the symbolic name of the task. When non-NULL and non-empty, a copy of this string is used for indexing the created task into the object registry.
<i>stksize</i>	The size of the stack (in bytes) for the new task. If zero is passed, a system-dependent default size will be substituted.
<i>prio</i>	The base priority of the new task. This value must be in the [0 .. 99] range, where 0 is the lowest effective priority.
<i>mode</i>	The task creation mode. The following flags can be OR'ed into this bitmask:



- `T_FPU` allows the task to use the FPU whenever available on the platform. This flag may be omitted, as it is automatically set when a FPU is present on the platform, cleared otherwise.
- `T_SUSP` causes the task to start in suspended mode. In such a case, the thread will have to be explicitly resumed using the `rt_task_resume()` service for its execution to actually begin.
- `T_CPU(cpuid)` makes the new task affine to CPU # **cpuid**. CPU identifiers range from 0 to 7 (inclusive).
- `T_JOINABLE` allows another task to wait on the termination of the new task. `rt_task_join()` shall be called for this task to clean up any resources after its termination.

Passing `T_FPU|T_CPU(1)` in the *mode* parameter thus creates a task with FPU support enabled and which will be affine to CPU #1.

- When running over the Cobalt core, `T_WARNSW` causes the `SIGDEBUG` signal to be sent to the current task whenever it switches to the secondary mode. This feature is useful to detect unwanted migrations to the Linux domain. This flag has no effect over the Mercury core.

#### Returns

Zero is returned upon success. Otherwise:

- `-EINVAL` is returned if either *prio*, *mode* or *stksize* are invalid.
- `-ENOMEM` is returned if the system fails to get memory from the main heap in order to create the task.
- `-EEXIST` is returned if the *name* is conflicting with an already registered task.

#### Tags

[thread-unrestricted](#), [switch-secondary](#)

#### Side effects

- calling `rt_task_create()` causes `SCHED_FIFO` tasks to switch to

secondary mode.

- members of Xenomai's `SCHED_FIFO` class running in the primary domain have utmost priority over all Linux activities in the system, including Linux interrupt handlers.

#### Note

Tasks can be referred to from multiple processes which all belong to the same Xenomai session.

**Deprecated** This is a compatibility service from the Transition Kit.

## 6.93.2.7 COMPAT\_\_rt\_task\_set\_periodic()

```
int COMPAT__rt_task_set_periodic (  
    RT_TASK * task,  
    RTIME idate,  
    RTIME period )
```

Make a real-time task periodic (compatibility service).

Make a task periodic by programing its first release point and its period in the processor time line. *task* should then call [rt\\_task\\_wait\\_period\(\)](#) to sleep until the next periodic release point in the processor timeline is reached.

## Parameters

<i>task</i>	The task descriptor. If <i>task</i> is NULL, the current task is made periodic. <i>task</i> must belong to the current process.
<i>idate</i>	The initial (absolute) date of the first release point, expressed in clock ticks (see note). If <i>idate</i> is equal to TM_NOW, the current system date is used. Otherwise, if <i>task</i> is NULL or equal to <code>rt_task_self()</code> , the caller is delayed until <i>idate</i> has elapsed.
<i>period</i>	The period of the task, expressed in clock ticks (see note). Passing TM_INFINITE stops the task's periodic timer if enabled, then returns successfully.

## Returns

Zero is returned upon success. Otherwise:

- -EINVAL is returned if *task* is NULL but the caller is not a Xenomai task, or if *task* is non-NULL but not a valid task descriptor.
- -ETIMEDOUT is returned if *idate* is different from TM\_INFINITE and represents a date in the past.

## Tags

[thread-unrestricted](#), [switch-primary](#)

## Note

The caller must be an Alchemy task if *task* is NULL.

Unlike the original Xenomai 2.x call, this emulation delays the caller until *idate* has elapsed only if *task* is NULL or equal to `rt_task_self()`.

## Side effects

Over Cobalt, -EINVAL is returned if *period* is different from TM\_INFINITE but shorter than the user scheduling latency value for the target system, as displayed by `/proc/xenomai/latency`.

## Note

The *idate* and *period* values are interpreted as a multiple of the Alchemy clock resolution (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

**Deprecated** This is a compatibility service from the Transition Kit.

## 6.93.2.8 pthread\_make\_periodic\_np()

```
int pthread_make_periodic_np (
    pthread_t thread,
    struct timespec * starttp,
    struct timespec * periodtp )
```

Make a thread periodic (compatibility service).

This service makes the POSIX *thread* periodic.

## Parameters

<i>thread</i>	thread to arm a periodic timer for.
<i>starttp</i>	start time, expressed as an absolute value of the CLOCK_REALTIME clock.
<i>periodtp</i>	period, expressed as a time interval.

## Returns

0 on success;

an error number if:

- ESRCH, *thread* is invalid.
- ETIMEDOUT, the start time has already passed.
- EPERM, the caller is not a Xenomai thread.
- EINVAL, *thread* does not refer to the current thread.

## Note

Unlike the original Xenomai 2.x call, this emulation does not delay the caller waiting for the first periodic release point. In addition, *thread* must be equal to `pthread_self()`.

**Deprecated** This service is a non-portable extension of the Xenomai 2.x POSIX interface, not available with Xenomai 3.x. Instead, Cobalt-based applications should set up a periodic timer using the `timer_create()`, `timer_settime()` call pair, then wait for release points via `sigwaitinfo()`. Overruns can be detected by looking at the `siginfo.si_overrun` field. Alternatively, applications may obtain a file descriptor referring to a Cobalt timer via the `timerfd()` call, and `read()` from it to wait for timeouts.

References `timer_create()`, and `timer_settime()`.

6.93.2.9 `pthread_wait_np()`

```
int pthread_wait_np (
    unsigned long * overruns_r )
```

Wait for the next periodic release point (compatibility service)

Delay the current thread until the next periodic release point is reached. The periodic timer should have been previously started for *thread* by a call to `pthread_make_periodic_np()`.

## Parameters

<i>overruns_r</i>	If non-NULL, <i>overruns_r</i> shall be a pointer to a memory location which will be written with the count of pending overruns. This value is written to only when <code>pthread_wait_np()</code> returns ETIMEDOUT or success. The memory location remains unmodified otherwise. If NULL, this count will not be returned.
-------------------	--

## Returns

Zero is returned upon success. If *overruns\_r* is non-NULL, zero is written to the pointed memory location. Otherwise:

- EWOULDBLOCK is returned if [pthread\\_make\\_periodic\\_np\(\)](#) was not called for the current thread.
- EINTR is returned if *thread* was interrupted by a signal before the next periodic release point was reached.
- ETIMEDOUT is returned if a timer overrun occurred, which indicates that a previous release point was missed by the calling thread. If *overruns\_r* is non-NULL, the count of pending overruns is written to the pointed memory location.
- EPERM is returned if this service was called from an invalid context.

## Note

If the current release point has already been reached at the time of the call, the current thread immediately returns from this service with no delay.

**Deprecated** This service is a non-portable extension of the Xenomai 2.x POSIX interface, not available with Xenomai 3.x. Instead, Cobalt-based applications should set up a periodic timer using the [timer\\_create\(\)](#), [timer\\_settime\(\)](#) call pair, then wait for release points via [sigwaitinfo\(\)](#). Overruns can be detected by looking at the `siginfo.si_overrun` field. Alternatively, applications may obtain a file descriptor referring to a Cobalt timer via the [timerfd\(\)](#) call, and [read\(\)](#) from it to wait for timeouts.

6.93.2.10 `rt_alarm_wait()`

```
int rt_alarm_wait (
    RT_ALARM * alarm )
```

Wait for the next alarm shot (compatibility service).

This service allows the current task to suspend execution until the specified alarm triggers. The priority of the current task is raised above all other tasks - except those also undergoing an alarm wait.

## Returns

Zero is returned upon success, after the alarm timed out. Otherwise:

- -EINVAL is returned if *alarm* is not a valid alarm descriptor.
- -EPERM is returned if this service was called from an invalid context.
- -EINTR is returned if [rt\\_task\\_unblock\(\)](#) was called for the current task before the request is satisfied.
- -EIDRM is returned if *alarm* is deleted while the caller was sleeping on it. In such a case, *alarm* is no more valid upon return of this service.

## Tags

[xthread-only](#), [switch-primary](#)

**Deprecated** This is a compatibility service from the Transition Kit.



# Chapter 7

## Data Structure Documentation

### 7.1 a4l\_channel Struct Reference

Structure describing some channel's characteristics.

#### Data Fields

- unsigned long [flags](#)
- unsigned long [nb\\_bits](#)

#### 7.1.1 Detailed Description

Structure describing some channel's characteristics.

#### 7.1.2 Field Documentation

##### 7.1.2.1 flags

```
unsigned long a4l_channel::flags
```

Channel flags to define the reference.

##### 7.1.2.2 nb\_bits

```
unsigned long a4l_channel::nb_bits
```

Channel resolution.

Referenced by `a4l_get_chan()`.


The documentation for this struct was generated from the following file:

- `include/cobalt/kernel/rtdm/analogy/channel_range.h`

## 7.2 a4l\_channels\_desc Struct Reference

Structure describing a channels set.

Collaboration diagram for a4l\_channels\_desc:



### Data Fields

- unsigned long [mode](#)
- unsigned long [length](#)
- struct [a4l\\_channel](#) [chans](#) []

### 7.2.1 Detailed Description

Structure describing a channels set.

### 7.2.2 Field Documentation

#### 7.2.2.1 chans

```
struct a4l\_channel a4l_channels_desc::chans[]
```

Channels tab

Referenced by [a4l\\_get\\_chan\(\)](#).



## 7.2.2.2 length

unsigned long a4l\_channels\_desc::length

Channels count

## 7.2.2.3 mode

unsigned long a4l\_channels\_desc::mode

Declaration mode (global or per channel)

Referenced by `a4l_get_chan()`.

The documentation for this struct was generated from the following file:

- `include/cobalt/kernel/rtdm/analogy/channel_range.h`

## 7.3 a4l\_cmd\_desc Struct Reference

Structure describing the asynchronous instruction.

### Data Fields

- unsigned char `idx_subd`  
*Subdevice to which the command will be applied.*
- unsigned long `flags`  
*Command flags.*
- unsigned int `start_src`  
*Start trigger type.*
- unsigned int `start_arg`  
*Start trigger argument.*
- unsigned int `scan_begin_src`  
*Scan begin trigger type.*
- unsigned int `scan_begin_arg`  
*Scan begin trigger argument.*
- unsigned int `convert_src`  
*Convert trigger type.*
- unsigned int `convert_arg`  
*Convert trigger argument.*
- unsigned int `scan_end_src`  
*Scan end trigger type.*
- unsigned int `scan_end_arg`  
*Scan end trigger argument.*
- unsigned int `stop_src`  
*Stop trigger type.*
- unsigned int `stop_arg`  
*Stop trigger argument.*
- unsigned char `nb_chan`  
*Count of channels related with the command.*
- unsigned int \* `chan_descs`  
*Tab containing channels descriptors.*
- unsigned int `data_len`  
*< cmd simulation valid stages (driver dependent)*
- `sampl_t` \* `data`  
*Driver specific buffer pointer.*

### 7.3.1 Detailed Description

Structure describing the asynchronous instruction.

See also

[a4l\\_snd\\_command\(\)](#)

### 7.3.2 Field Documentation

#### 7.3.2.1 data\_len

```
unsigned int a4l_cmd_desc::data_len
```

< cmd simulation valid stages (driver dependent)

Driver specific buffer size

#### 7.3.2.2 idx\_subd

```
unsigned char a4l_cmd_desc::idx_subd
```

Subdevice to which the command will be applied.

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/analogy.h](#)

## 7.4 a4l\_descriptor Struct Reference

Structure containing device-information useful to users.

### Data Fields

- char [board\\_name](#) [A4L\_NAMELEN]  
*Board name.*
- char [driver\\_name](#) [A4L\_NAMELEN]  
*Driver name.*
- int [nb\\_subd](#)  
*Subdevices count.*
- int [idx\\_read\\_subd](#)  
*Input subdevice index.*
- int [idx\\_write\\_subd](#)  
*Output subdevice index.*
- int [fd](#)  
*File descriptor.*
- unsigned int [magic](#)  
*Opaque field.*
- int [sbsize](#)  
*Data buffer size.*
- void \* [sbdata](#)  
*Data buffer pointer.*

### 7.4.1 Detailed Description

Structure containing device-information useful to users.

See also

`a4l_get_desc()`

### 7.4.2 Field Documentation

#### 7.4.2.1 board\_name

```
char a4l_descriptor::board_name[A4L_NAMELEN]
```

Board name.

#### 7.4.2.2 driver\_name

```
char a4l_descriptor::driver_name[A4L_NAMELEN]
```

Driver name.

#### 7.4.2.3 fd

```
int a4l_descriptor::fd
```

File descriptor.

Referenced by `a4l_close()`, `a4l_fill_desc()`, `a4l_get_bufsize()`, `a4l_mark_bufrw()`, `a4l_mmap()`, `a4l_open()`, `a4l_poll()`, `a4l_set_bufsize()`, `a4l_snd_cancel()`, `a4l_snd_command()`, `a4l_snd_insn()`, and `a4l_snd_insnlist()`.

#### 7.4.2.4 idx\_read\_subd

```
int a4l_descriptor::idx_read_subd
```

Input subdevice index.

#### 7.4.2.5 idx\_write\_subd

```
int a4l_descriptor::idx_write_subd
```

Output subdevice index.

#### 7.4.2.6 magic

```
unsigned int a4l_descriptor::magic
```

Opaque field.

Referenced by `a4l_fill_desc()`, `a4l_find_range()`, `a4l_get_chinfo()`, `a4l_get_rnginfo()`, `a4l_get_subdinfo()`, and `a4l_sys_desc()`.

#### 7.4.2.7 nb\_subd

```
int a4l_descriptor::nb_subd
```

Subdevices count.

#### 7.4.2.8 sbdata

```
void* a4l_descriptor::sbdata
```

Data buffer pointer.

#### 7.4.2.9 sbsize

```
int a4l_descriptor::sbsize
```

Data buffer size.

The documentation for this struct was generated from the following file:

- [include/rtdm/analogy.h](#)

## 7.5 a4l\_driver Struct Reference

Structure containing driver declaration data.

## Data Fields

- struct list\_head [list](#)  
*List stuff.*
- struct module \* [owner](#)  
*Pointer to module containing the code.*
- unsigned int [flags](#)  
*Type / status driver's flags.*
- char \* [board\\_name](#)  
*Board name.*
- char \* [driver\\_name](#)  
*driver name*
- int [privdata\\_size](#)  
*Size of the driver's private data.*
- int(\* [attach](#) )(struct a4l\_device \*, struct a4l\_link\_desc \*)  
*Attach procedure.*
- int(\* [detach](#) )(struct a4l\_device \*)  
*Detach procedure.*

### 7.5.1 Detailed Description

Structure containing driver declaration data.

See also

[rt\\_task\\_inquire\(\)](#)

The documentation for this struct was generated from the following file:

- include/cobalt/kernel/rtdm/analogy/[driver.h](#)

## 7.6 a4l\_instruction Struct Reference

Structure describing the synchronous instruction.

## Data Fields

- unsigned int [type](#)  
*Instruction type.*
- unsigned int [idx\\_subd](#)  
*Subdevice to which the instruction will be applied.*
- unsigned int [chan\\_desc](#)  
*Channel descriptor.*
- unsigned int [data\\_size](#)  
*Size of the intruction data.*
- void \* [data](#)  
*Instruction data.*

### 7.6.1 Detailed Description

Structure describing the synchronous instruction.

See also

[a4l\\_snd\\_insn\(\)](#)

### 7.6.2 Field Documentation

#### 7.6.2.1 idx\_subd

```
unsigned int a4l_instruction::idx_subd
```

Subdevice to which the instruction will be applied.


The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/analogy.h](#)

## 7.7 a4l\_instruction\_list Struct Reference

Structure describing the list of synchronous instructions.

Collaboration diagram for a4l\_instruction\_list:



### Data Fields

- unsigned int [count](#)  
*Instructions count.*
- [a4l\\_insn\\_t](#) \* [insns](#)  
*Tab containing the instructions pointers.*

### 7.7.1 Detailed Description

Structure describing the list of synchronous instructions.

See also

[a4l\\_snd\\_insnlist\(\)](#)

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/analogy.h](#)

## 7.8 a4l\_range Struct Reference

Structure describing a (unique) range.

### Data Fields

- long [min](#)
- long [max](#)
- unsigned long [flags](#)

### 7.8.1 Detailed Description

Structure describing a (unique) range.

### 7.8.2 Field Documentation

#### 7.8.2.1 flags

```
unsigned long a4l_range::flags
```

Range flags (unit, etc.)

#### 7.8.2.2 max

```
long a4l_range::max
```

Maximal false

## 7.8.2.3 min

```
long a4l_range::min
```

Minimal value

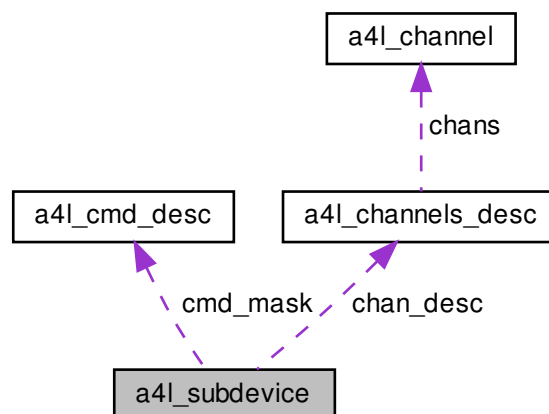
The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/analog/channel\\_range.h](#)

## 7.9 a4l\_subdevice Struct Reference

Structure describing the subdevice.

Collaboration diagram for a4l\_subdevice:



## Data Fields

- struct list\_head [list](#)  
*List stuff.*
- struct a4l\_device \* [dev](#)  
*Containing device.*
- unsigned int [idx](#)  
*Subdevice index.*
- struct a4l\_buffer \* [buf](#)  
*Linked buffer.*
- unsigned long [status](#)  
*Subdevice's status.*
- unsigned long [flags](#)



*Type flags.*

- struct [a4l\\_channels\\_desc](#) \* [chan\\_desc](#)  
*Tab of channels descriptors pointers.*
- struct [a4l\\_rngdesc](#) \* [rng\\_desc](#)  
*Tab of ranges descriptors pointers.*
- struct [a4l\\_cmd\\_desc](#) \* [cmd\\_mask](#)  
*Command capabilities mask.*
- int(\* [insn\\_read](#) )(struct [a4l\\_subdevice](#) \*, struct [a4l\\_kernel\\_instruction](#) \*)  
*Callback for the instruction "read".*
- int(\* [insn\\_write](#) )(struct [a4l\\_subdevice](#) \*, struct [a4l\\_kernel\\_instruction](#) \*)  
*Callback for the instruction "write".*
- int(\* [insn\\_bits](#) )(struct [a4l\\_subdevice](#) \*, struct [a4l\\_kernel\\_instruction](#) \*)  
*Callback for the instruction "bits".*
- int(\* [insn\\_config](#) )(struct [a4l\\_subdevice](#) \*, struct [a4l\\_kernel\\_instruction](#) \*)  
*Callback for the configuration instruction.*
- int(\* [do\\_cmd](#) )(struct [a4l\\_subdevice](#) \*, struct [a4l\\_cmd\\_desc](#) \*)  
*Callback for command handling.*
- int(\* [do\\_cmdtest](#) )(struct [a4l\\_subdevice](#) \*, struct [a4l\\_cmd\\_desc](#) \*)  
*Callback for command checking.*
- void(\* [cancel](#) )(struct [a4l\\_subdevice](#) \*)  
*Callback for asynchronous transfer cancellation.*
- void(\* [mung](#) )(struct [a4l\\_subdevice](#) \*, void \*, unsigned long)  
*Callback for munge operation.*
- int(\* [trigger](#) )(struct [a4l\\_subdevice](#) \*, [lsampl\\_t](#))  
*Callback for trigger operation.*
- char [priv](#) [0]  
*Private data.*

### 7.9.1 Detailed Description

Structure describing the subdevice.

See also

[a4l\\_add\\_subd\(\)](#)

The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/analogy/subdevice.h](#)

## 7.10 atomic\_t Struct Reference

### 7.10.1 Detailed Description

Copyright © 2011 Gilles Chanteperrin [gilles.chanteperrin@xenomai.org](mailto:gilles.chanteperrin@xenomai.org).

Copyright © 2013 Philippe Gerum [rpm@xenomai.org](mailto:rpm@xenomai.org).

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.


The documentation for this struct was generated from the following file:

- `include/boilerplate/atomic.h`

## 7.11 `can_bittime` Struct Reference

Custom CAN bit-time definition.

Collaboration diagram for `can_bittime`:



### 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.*

### 7.11.1 Detailed Description

Custom CAN bit-time definition.

Examples:

[rtcanconfig.c](#).

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/can.h](#)

## 7.12 `can_bittime_btr` Struct Reference

Hardware-specific BTR bit-times.

### Data Fields

- `uint8_t btr0`  
*Bus timing register 0.*
- `uint8_t btr1`  
*Bus timing register 1.*

### 7.12.1 Detailed Description

Hardware-specific BTR bit-times.

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/can.h](#)

## 7.13 `can_bittime_std` Struct Reference

Standard bit-time parameters according to Bosch.

### 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*

### 7.13.1 Detailed Description

Standard bit-time parameters according to Bosch.

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/can.h](#)

## 7.14 can\_filter Struct Reference

Filter for reception of CAN messages.

### 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.*

### 7.14.1 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:

[can-rtt.c](#), and [rtcanrecv.c](#).

### 7.14.2 Field Documentation

## 7.14.2.1 can\_id

```
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](#).

## 7.14.2.2 can\_mask

```
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/uapi/can.h](#)

## 7.15 can\_frame Struct Reference

Raw CAN frame.

### 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.*

### 7.15.1 Detailed Description

Raw CAN frame.

Central structure for receiving and sending CAN frames.

Examples:

[can-rtt.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

### 7.15.2 Field Documentation

#### 7.15.2.1 can\_id

`can_id_t` `can_frame::can_id`

CAN ID of the frame.

See [CAN ID flags](#) for special bits.

Examples:

[can-rtt.c](#).

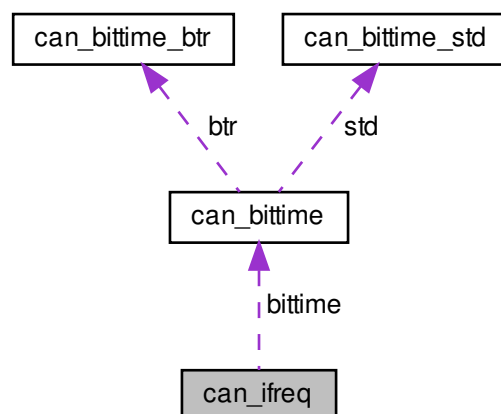
The documentation for this struct was generated from the following file:

- `include/rtdm/uapi/can.h`

## 7.16 can\_ifreq Struct Reference

CAN interface request descriptor.

Collaboration diagram for `can_ifreq`:



### 7.16.1 Detailed Description

CAN interface request descriptor.

Parameter block for submitting CAN control requests.

Examples:

[can-rtt.c](#), [rtcanconfig.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/can.h](#)

## 7.17 macb\_dma\_desc Struct Reference

Hardware DMA descriptor.

### Data Fields

- u32 [addr](#)  
*DMA address of data buffer.*
- u32 [ctrl](#)  
*Control and status bits.*

### 7.17.1 Detailed Description

Hardware DMA descriptor.

The documentation for this struct was generated from the following file:

- [kernel/drivers/net/drivers/rt\\_macb.h](#)

## 7.18 macb\_tx\_skb Struct Reference

Data about an skb which is being transmitted.

### Data Fields

- struct rtskb \* [skb](#)  
*skb currently being transmitted*
- dma\_addr\_t [mapping](#)  
*DMA address of the skb's data buffer.*

### 7.18.1 Detailed Description

Data about an skb which is being transmitted.

The documentation for this struct was generated from the following file:

- kernel/drivers/net/drivers/rt\_macb.h

## 7.19 RT\_ALARM\_INFO Struct Reference

Alarm status descriptor.

### Data Fields

- unsigned long [expiries](#)  
*Number of past expiries.*
- char [name](#) [XNOBJECT\_NAME\_LEN]  
*Name of alarm object.*
- int [active](#)  
*Active flag.*

### 7.19.1 Detailed Description

Alarm status descriptor.

This structure reports various static and runtime information about a real-time alarm, returned by a call to [rt\\_alarm\\_inquire\(\)](#).

The documentation for this struct was generated from the following file:

- include/alchemy/alarm.h

## 7.20 RT\_BUFFER\_INFO Struct Reference

Buffer status descriptor.

### Data Fields

- int [iwaiters](#)  
*Number of tasks waiting on the read side of the buffer for input data.*
- int [owaiters](#)  
*Number of tasks waiting on the write side of the buffer for sending out data.*
- size\_t [totalmem](#)  
*Overall size of buffer (in bytes).*
- size\_t [availmem](#)  
*Amount of memory currently available for holding more data.*
- char [name](#) [XNOBJECT\_NAME\_LEN]  
*Name of the buffer.*



### 7.20.1 Detailed Description

Buffer status descriptor.

This structure reports various static and runtime information about a real-time buffer, returned by a call to [rt\\_buffer\\_inquire\(\)](#).

The documentation for this struct was generated from the following file:

- `include/alchemy/buffer.h`

## 7.21 RT\_COND\_INFO Struct Reference

Condition variable status descriptor.

### Data Fields

- char [name](#) [XNOBJECT\_NAME\_LEN]  
*Name of condition variable.*

### 7.21.1 Detailed Description

Condition variable status descriptor.

This structure reports various static and runtime information about a condition variable, returned by a call to [rt\\_cond\\_inquire\(\)](#).

The documentation for this struct was generated from the following file:

- `include/alchemy/cond.h`

## 7.22 RT\_EVENT\_INFO Struct Reference

Event status descriptor.

### Data Fields

- unsigned int [value](#)  
*Current value of the event flag group.*
- int [nwaiters](#)  
*Number of tasks currently waiting for events.*
- char [name](#) [XNOBJECT\_NAME\_LEN]  
*Name of event flag group.*

### 7.22.1 Detailed Description

Event status descriptor.

This structure reports various static and runtime information about an event flag group, returned by a call to [rt\\_event\\_inquire\(\)](#).

The documentation for this struct was generated from the following file:

- `include/alchemy/event.h`

## 7.23 RT\_HEAP\_INFO Struct Reference

Heap status descriptor.

### Data Fields

- `int nwaiters`  
*Number of tasks waiting for available memory in [rt\\_heap\\_alloc\(\)](#).*
- `int mode`  
*Creation mode flags as given to [rt\\_heap\\_create\(\)](#).*
- `size_t heapsize`  
*Size of heap (in bytes) as given to [rt\\_heap\\_create\(\)](#).*
- `size_t usablemem`  
*Maximum amount of memory available from the heap.*
- `size_t usedmem`  
*Amount of heap memory currently consumed.*
- `char name [XNOBJECT_NAME_LEN]`  
*Name of heap.*

### 7.23.1 Detailed Description

Heap status descriptor.

This structure reports various static and runtime information about a real-time heap, returned by a call to [rt\\_heap\\_inquire\(\)](#).

### 7.23.2 Field Documentation

#### 7.23.2.1 heapsize

```
size_t RT_HEAP_INFO::heapsize
```

Size of heap (in bytes) as given to [rt\\_heap\\_create\(\)](#).

The maximum amount of memory available from this heap may be larger, due to internal padding.

### 7.23.2.2 usablemem

size\_t RT\_HEAP\_INFO::usablemem

Maximum amount of memory available from the heap.

This value accounts for the overhead of internal data structures required to maintain the heap.

### 7.23.2.3 usedmem

size\_t RT\_HEAP\_INFO::usedmem

Amount of heap memory currently consumed.

info.usablemem - info.usedmem computes the current amount of free memory in the relevant heap.

The documentation for this struct was generated from the following file:

- include/alchemy/heap.h

## 7.24 RT\_MUTEX\_INFO Struct Reference

Mutex status descriptor.

### Data Fields

- RT\_TASK [owner](#)  
*Current mutex owner, or NO\_ALCHEMY\_TASK if unlocked.*
- char [name](#) [XNOBJECT\_NAME\_LEN]  
*Name of mutex.*

### 7.24.1 Detailed Description

Mutex status descriptor.

This structure reports various static and runtime information about a mutex, returned by a call to [rt\\_mutex\\_inquire\(\)](#).

### 7.24.2 Field Documentation

### 7.24.2.1 owner

`RT_TASK RT_MUTEX_INFO::owner`

Current mutex owner, or `NO_ALCHEMY_TASK` if unlocked.

This information is in essence transient, and may not be valid anymore once used by the caller.

The documentation for this struct was generated from the following file:

- `include/alchemy/mutex.h`

## 7.25 RT\_QUEUE\_INFO Struct Reference

Queue status descriptor.

### Data Fields

- `int nwaiters`  
*Number of tasks currently waiting on the queue for messages.*
- `int nmessages`  
*Number of messages pending in queue.*
- `int mode`  
*Queue mode bits, as given to `rt_queue_create()`.*
- `size_t qlimit`  
*Maximum number of messages in queue, zero if unlimited.*
- `size_t poolsize`  
*Size of memory pool for holding message buffers (in bytes).*
- `size_t usedmem`  
*Amount of memory consumed from the buffer pool.*
- `char name [XNOBJECT_NAME_LEN]`  
*Name of message queue.*

### 7.25.1 Detailed Description

Queue status descriptor.

This structure reports various static and runtime information about a real-time queue, returned by a call to `rt_queue_inquire()`.

The documentation for this struct was generated from the following file:

- `include/alchemy/queue.h`

## 7.26 RT\_SEM\_INFO Struct Reference

Semaphore status descriptor.

## Data Fields

- unsigned long `count`  
*Current semaphore value.*
- int `nwaiters`  
*Number of tasks waiting on the semaphore.*
- char `name` [XNOBJECT\_NAME\_LEN]  
*Name of semaphore.*

### 7.26.1 Detailed Description

Semaphore status descriptor.

This structure reports various static and runtime information about a semaphore, returned by a call to [rt\\_sem\\_inquire\(\)](#).

The documentation for this struct was generated from the following file:

- `include/alchemy/sem.h`

## 7.27 RT\_TASK\_INFO Struct Reference

Task status descriptor.

## Data Fields

- int `prio`  
*Task priority.*
- struct `threadobj_stat` `stat`  
*Task status.*
- char `name` [XNOBJECT\_NAME\_LEN]  
*Name of task.*
- `pid_t` `pid`  
*Host pid.*

### 7.27.1 Detailed Description

Task status descriptor.

This structure reports various static and runtime information about a real-time task, returned by a call to [rt\\_task\\_inquire\(\)](#).

The documentation for this struct was generated from the following file:

- `include/alchemy/task.h`

## 7.28 rt\_timer\_info Struct Reference

Timer status descriptor.

### Data Fields

- [RTIME period](#)  
*Clock resolution in nanoseconds.*
- [RTIME date](#)  
*Current monotonic date expressed in clock ticks.*

### 7.28.1 Detailed Description

Timer status descriptor.

This structure reports information about the Alchemy clock, returned by a call to [rt\\_timer\\_inquire\(\)](#).

### 7.28.2 Field Documentation

#### 7.28.2.1 date

```
RTIME rt_timer_info::date
```

Current monotonic date expressed in clock ticks.

The duration of a tick depends on the Alchemy clock resolution for the process (see `-alchemy-clock-resolution` option, defaults to 1 nanosecond).

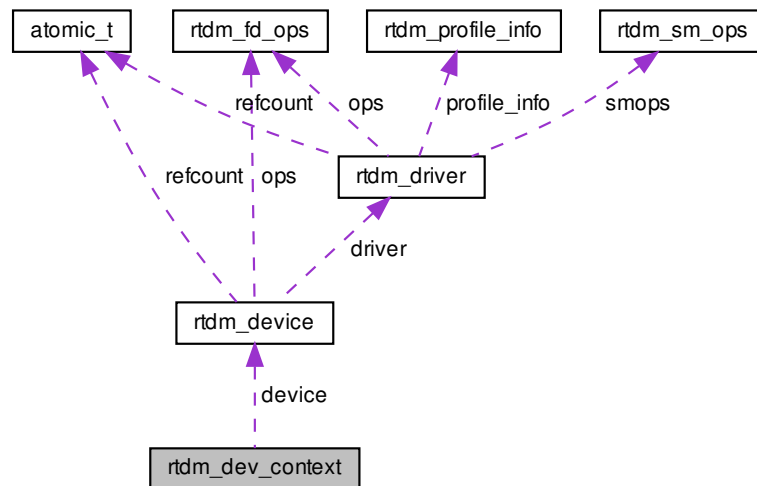
The documentation for this struct was generated from the following file:

- `include/alchemy/timer.h`

## 7.29 rtdm\_dev\_context Struct Reference

Device context.

Collaboration diagram for rtdm\_dev\_context:



### Data Fields

- struct `rtdm_device` \* `device`  
Set of active device operation handlers.
- char `dev_private` [0]  
Begin of driver defined context data structure.

#### 7.29.1 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 `rtdm_driver.context_size` during device registration.

#### 7.29.2 Field Documentation

## 7.29.2.1 device

```
struct rtdm_device* rtdm_dev_context::device
```

Set of active device operation handlers.

Reference to owning device

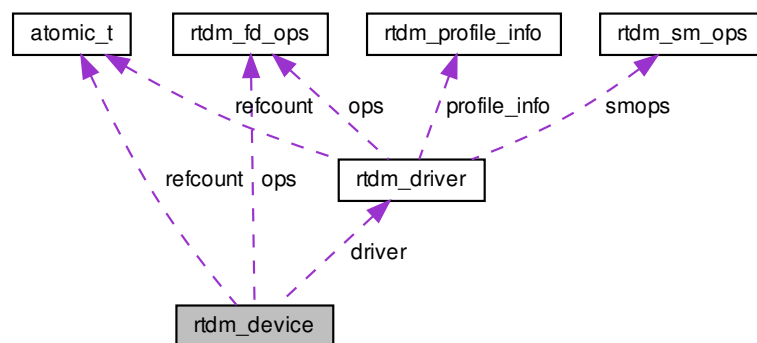
The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/driver.h](#)

## 7.30 rtdm\_device Struct Reference

RTDM device.

Collaboration diagram for rtdm\_device:



## Data Fields

- struct `rtdm_driver` \* `driver`  
*Device driver.*
- void \* `device_data`  
*Driver definable device data.*
- const char \* `label`  
*Device label template for composing the device name.*
- int `minor`  
*Minor number of the device.*
- struct {  
};  
*Reserved area.*



### 7.30.1 Detailed Description

RTDM device.

This descriptor describes a RTDM device instance. The structure holds runtime data, therefore it must reside in writable memory.

### 7.30.2 Field Documentation

#### 7.30.2.1 "@15

```
struct { ... }
```

Reserved area.

#### 7.30.2.2 driver

```
struct rtdm\_driver* rtdm_device::driver
```

Device driver.

#### 7.30.2.3 label

```
const char* rtdm_device::label
```

Device label template for composing the device name.

A limited printf-like format string is assumed, with a provision for replacing the first d/i placeholder found in the string by the device minor number. It is up to the driver to actually mention this placeholder or not, depending on the naming convention for its devices. For named devices, the corresponding device node will automatically appear in the /dev/rtdm hierarchy with hotplug-enabled device filesystems (DEVTMPFS).

#### 7.30.2.4 minor

```
int rtdm_device::minor
```

Minor number of the device.

If RTDM\_FIXED\_MINOR is present in the driver flags, the value stored in this field is used verbatim by [rtdm\\_dev\\_register\(\)](#). Otherwise, the RTDM core automatically assigns minor numbers to all devices managed by the driver referred to by *driver*, in order of registration, storing the resulting values into this field.

Device nodes created for named devices in the Linux /dev hierarchy are assigned this minor number.

The minor number of the current device handling an I/O request can be retrieved by a call to [rtdm\\_fd\\_get\\_minor\(\)](#).

The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/driver.h](#)

## 7.31 rtdm\_device\_info Struct Reference

Device information.

### 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.*

### 7.31.1 Detailed Description

Device information.


The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/rtdm.h](#)

## 7.32 rtdm\_driver Struct Reference

RTDM driver.

Collaboration diagram for rtdm\_driver:



## Data Fields

- struct [rtdm\\_profile\\_info](#) `profile_info`  
*Class profile information.*
- int [device\\_flags](#)  
*Device flags, see [Device Flags](#) for details.*
- size\_t [context\\_size](#)  
*Size of the private memory area the core should automatically allocate for each open file descriptor, which is usable for storing the context data associated to each connection.*
- int [protocol\\_family](#)  
*Protocol device identification: protocol family (PF\_XXX)*
- int [socket\\_type](#)  
*Protocol device identification: socket type (SOCK\_XXX)*
- struct [rtdm\\_fd\\_ops](#) `ops`  
*I/O operation handlers.*
- struct [rtdm\\_sm\\_ops](#) `smops`  
*State management handlers.*
- int [device\\_count](#)  
*Count of devices this driver manages.*
- int [base\\_minor](#)  
*Base minor for named devices.*
- ```
struct {
};
```

  
*Reserved area.*

## 7.32.1 Detailed Description

RTDM driver.

This descriptor describes a RTDM device driver. The structure holds runtime data, therefore it must reside in writable memory.

## 7.32.2 Field Documentation

## 7.32.2.1 base\_minor

```
int rtdm_driver::base_minor
```

Base minor for named devices.

### 7.32.2.2 context\_size

```
size_t rtdm_driver::context_size
```

Size of the private memory area the core should automatically allocate for each open file descriptor, which is usable for storing the context data associated to each connection.

The allocated memory is zero-initialized. The start of this area can be retrieved by a call to [rtdm\\_fd\\_to\\_private\(\)](#).

### 7.32.2.3 device\_count

```
int rtdm_driver::device_count
```

Count of devices this driver manages.

This value is used to allocate a chrdev region for named devices.

### 7.32.2.4 device\_flags

```
int rtdm_driver::device_flags
```

Device flags, see [Device Flags](#) for details.

### 7.32.2.5 profile\_info

```
struct rtdm_profile_info rtdm_driver::profile_info
```

Class profile information.

The [RTDM\\_PROFILE\\_INFO\(\)](#) macro **must** be used for filling up this field.

The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/driver.h](#)

## 7.33 rtdm\_fd\_ops Struct Reference

RTDM file operation descriptor.

## Data Fields

- `int(* open )(struct rtdm_fd *fd, int oflags)`  
See `rtdm_open_handler()`.
- `int(* socket )(struct rtdm_fd *fd, int protocol)`  
See `rtdm_socket_handler()`.
- `void(* close )(struct rtdm_fd *fd)`  
See `rtdm_close_handler()`.
- `int(* ioctl_rt )(struct rtdm_fd *fd, unsigned int request, void __user *arg)`  
See `rtdm_ioctl_handler()`.
- `int(* ioctl_nrt )(struct rtdm_fd *fd, unsigned int request, void __user *arg)`  
See `rtdm_ioctl_handler()`.
- `ssize_t(* read_rt )(struct rtdm_fd *fd, void __user *buf, size_t size)`  
See `rtdm_read_handler()`.
- `ssize_t(* read_nrt )(struct rtdm_fd *fd, void __user *buf, size_t size)`  
See `rtdm_read_handler()`.
- `ssize_t(* write_rt )(struct rtdm_fd *fd, const void __user *buf, size_t size)`  
See `rtdm_write_handler()`.
- `ssize_t(* write_nrt )(struct rtdm_fd *fd, const void __user *buf, size_t size)`  
See `rtdm_write_handler()`.
- `ssize_t(* recvmmsg_rt )(struct rtdm_fd *fd, struct user_msghdr *msg, int flags)`  
See `rtdm_recvmmsg_handler()`.
- `ssize_t(* recvmmsg_nrt )(struct rtdm_fd *fd, struct user_msghdr *msg, int flags)`  
See `rtdm_recvmmsg_handler()`.
- `ssize_t(* sendmsg_rt )(struct rtdm_fd *fd, const struct user_msghdr *msg, int flags)`  
See `rtdm_sendmsg_handler()`.
- `ssize_t(* sendmsg_nrt )(struct rtdm_fd *fd, const struct user_msghdr *msg, int flags)`  
See `rtdm_sendmsg_handler()`.
- `int(* select )(struct rtdm_fd *fd, struct xselector *selector, unsigned int type, unsigned int index)`  
See `rtdm_select_handler()`.
- `int(* mmap )(struct rtdm_fd *fd, struct vm_area_struct *vma)`  
See `rtdm_mmap_handler()`.
- `unsigned long(* get_unmapped_area )(struct rtdm_fd *fd, unsigned long len, unsigned long pgoff, unsigned long flags)`  
See `rtdm_get_unmapped_area_handler()`.

## 7.33.1 Detailed Description

RTDM file operation descriptor.

This structure describes the operations available with a RTDM device, defining handlers for submitting I/O requests. Those handlers are implemented by RTDM device drivers.

## 7.33.2 Field Documentation

### 7.33.2.1 close

`void(* rtdm_fd_ops::close) (struct rtdm_fd *fd)`

See [rtdm\\_close\\_handler\(\)](#).

### 7.33.2.2 get\_unmapped\_area

`unsigned long(* rtdm_fd_ops::get_unmapped_area) (struct rtdm_fd *fd, unsigned long len, unsigned long pgoff, unsigned long flags)`

See [rtdm\\_get\\_unmapped\\_area\\_handler\(\)](#).

### 7.33.2.3 ioctl\_nrt

`int(* rtdm_fd_ops::ioctl_nrt) (struct rtdm_fd *fd, unsigned int request, void __user *arg)`

See [rtdm\\_ioctl\\_handler\(\)](#).

### 7.33.2.4 ioctl\_rt

`int(* rtdm_fd_ops::ioctl_rt) (struct rtdm_fd *fd, unsigned int request, void __user *arg)`

See [rtdm\\_ioctl\\_handler\(\)](#).

### 7.33.2.5 mmap

`int(* rtdm_fd_ops::mmap) (struct rtdm_fd *fd, struct vm_area_struct *vma)`

See [rtdm\\_mmap\\_handler\(\)](#).

### 7.33.2.6 open

`int(* rtdm_fd_ops::open) (struct rtdm_fd *fd, int oflags)`

See [rtdm\\_open\\_handler\(\)](#).

## 7.33.2.7 read\_nrt

ssize\_t(\* rtdm\_fd\_ops::read\_nrt) (struct rtdm\_fd \*fd, void \_\_user \*buf, size\_t size)

See [rtdm\\_read\\_handler\(\)](#).

## 7.33.2.8 read\_rt

ssize\_t(\* rtdm\_fd\_ops::read\_rt) (struct rtdm\_fd \*fd, void \_\_user \*buf, size\_t size)

See [rtdm\\_read\\_handler\(\)](#).

## 7.33.2.9 recvmsg\_nrt

ssize\_t(\* rtdm\_fd\_ops::recvmsg\_nrt) (struct rtdm\_fd \*fd, struct user\_msghdr \*msg, int flags)

See [rtdm\\_recvmsg\\_handler\(\)](#).

## 7.33.2.10 recvmsg\_rt

ssize\_t(\* rtdm\_fd\_ops::recvmsg\_rt) (struct rtdm\_fd \*fd, struct user\_msghdr \*msg, int flags)

See [rtdm\\_recvmsg\\_handler\(\)](#).

## 7.33.2.11 select

int(\* rtdm\_fd\_ops::select) (struct rtdm\_fd \*fd, struct xselector \*selector, unsigned int type, unsigned int index)

See [rtdm\\_select\\_handler\(\)](#).

## 7.33.2.12 sendmsg\_nrt

ssize\_t(\* rtdm\_fd\_ops::sendmsg\_nrt) (struct rtdm\_fd \*fd, const struct user\_msghdr \*msg, int flags)

See [rtdm\\_sendmsg\\_handler\(\)](#).

### 7.33.2.13 sendmsg\_rt

ssize\_t(\* rtdm\_fd\_ops::sendmsg\_rt) (struct rtdm\_fd \*fd, const struct user\_msghdr \*msg, int flags)

See [rtdm\\_sendmsg\\_handler\(\)](#).

### 7.33.2.14 socket

int(\* rtdm\_fd\_ops::socket) (struct rtdm\_fd \*fd, int protocol)

See [rtdm\\_socket\\_handler\(\)](#).

### 7.33.2.15 write\_nrt

ssize\_t(\* rtdm\_fd\_ops::write\_nrt) (struct rtdm\_fd \*fd, const void \_\_user \*buf, size\_t size)

See [rtdm\\_write\\_handler\(\)](#).

### 7.33.2.16 write\_rt

ssize\_t(\* rtdm\_fd\_ops::write\_rt) (struct rtdm\_fd \*fd, const void \_\_user \*buf, size\_t size)

See [rtdm\\_write\\_handler\(\)](#).

The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/fd.h](#)

## 7.34 rtdm\_gpio\_readout Struct Reference

### 7.34.1 Detailed Description

Note

Copyright (C) 2016 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.

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/gpio.h](#)



## 7.35 rtdm\_profile\_info Struct Reference

RTDM profile information.

### Data Fields

- const char \* [name](#)  
*Device class name.*
- int [class\\_id](#)  
*Device class ID, see [RTDM\\_CLASS\\_xxx](#).*
- int [subclass\\_id](#)  
*Device sub-class, see [RTDM\\_SUBCLASS\\_xxx](#) definition in the [Device Profiles](#).*
- int [version](#)  
*Supported device profile version.*
- unsigned int [magic](#)  
*Reserved.*

### 7.35.1 Detailed Description

RTDM profile information.

This descriptor details the profile information associated to a RTDM class of device managed by a driver.

The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/driver.h](#)

## 7.36 rtdm\_sm\_ops Struct Reference

RTDM state management handlers.

### Data Fields

- int(\* [start](#) )(struct [rtdm\\_driver](#) \*drv)  
*Handler called upon transition to COBALT\_STATE\_WARMUP.*
- int(\* [stop](#) )(struct [rtdm\\_driver](#) \*drv)  
*Handler called upon transition to COBALT\_STATE\_TEARDOWN.*

### 7.36.1 Detailed Description

RTDM state management handlers.

The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/driver.h](#)

## 7.37 rtdm\_spi\_config Struct Reference

### 7.37.1 Detailed Description

Note

Copyright (C) 2016 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.

The documentation for this struct was generated from the following file:

- `include/rtdm/uapi/spi.h`

## 7.38 rtipc\_port\_label Struct Reference

Port label information structure.

### Data Fields

- `char label [XNOBJECT_NAME_LEN]`  
*Port label string, null-terminated.*

### 7.38.1 Detailed Description

Port label information structure.

Examples:

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

### 7.38.2 Field Documentation

## 7.38.2.1 label

```
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/uapi/ipc.h](#)

## 7.39 rtser\_config Struct Reference

Serial device configuration.

### 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](#)*

### 7.39.1 Detailed Description

Serial device configuration.

Examples:

[cross-link.c](#).

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/serial.h](#)

## 7.40 rtser\_event Struct Reference

Additional information about serial device events.

### 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*

### 7.40.1 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/uapi/serial.h](#)

## 7.41 rtser\_status Struct Reference

Serial device status.

## Data Fields

- int [line\\_status](#)  
*line status register, see [RTSER\\_LSR\\_xxx](#)*
- int [modem\\_status](#)  
*modem status register, see [RTSER\\_MSR\\_xxx](#)*

### 7.41.1 Detailed Description

Serial device status.

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/serial.h](#)

## 7.42 sockaddr\_can Struct Reference

Socket address structure for the CAN address family.

## Data Fields

- `sa_family_t` [can\\_family](#)  
*CAN address family, must be `AF_CAN`.*
- int [can\\_ifindex](#)  
*Interface index of CAN controller.*

### 7.42.1 Detailed Description

Socket address structure for the CAN address family.

Examples:

[can-rtt.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

### 7.42.2 Field Documentation

#### 7.42.2.1 can\_ifindex

```
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/uapi/can.h](#)

## 7.43 sockaddr\_ipc Struct Reference

Socket address structure for the RTIPC address family.

### Data Fields

- [sa\\_family\\_t sipc\\_family](#)  
*RTIPC address family, must be AF\_RTIPC.*
- [rtipc\\_port\\_t sipc\\_port](#)  
*Port number.*

### 7.43.1 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](#).

### 7.43.2 Field Documentation

#### 7.43.2.1 sipc\_port

[rtipc\\_port\\_t](#) `sockaddr_ipc::sipc_port`

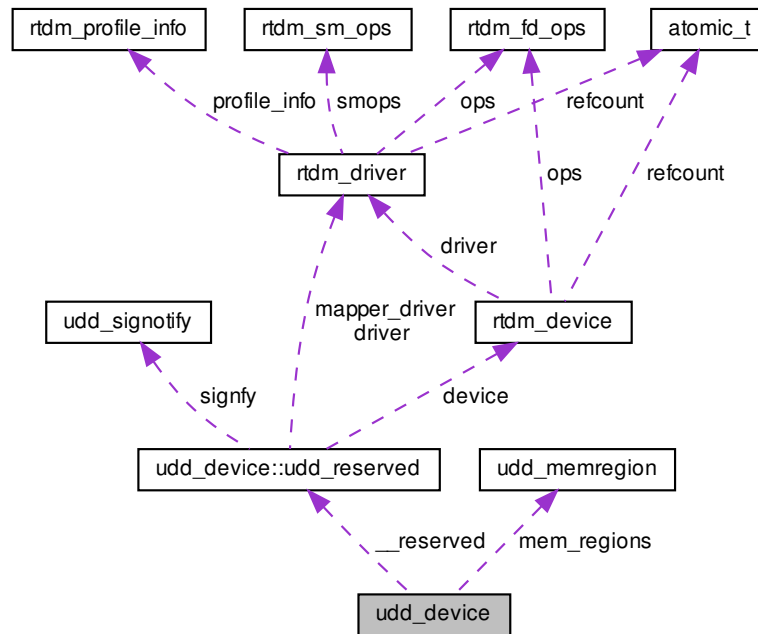
Port number.

The documentation for this struct was generated from the following file:

- [include/rtdm/uapi/ipc.h](#)

## 7.44 udd\_device Struct Reference

Collaboration diagram for udd\_device:



## Data Structures

- struct `udd_reserved`  
*Reserved to the UDD core.*

## Data Fields

- const char \* `device_name`  
*Name of the device managed by the mini-driver, appears automatically in the `/dev/rtdm` namespace upon creation.*
- int `device_flags`  
*Additional device flags (e.g.*
- int `device_subclass`  
*Subclass code of the device managed by the mini-driver (see `RTDM_SUBCLASS_xxx` definition in the [Device Profiles](#)).*
- int `irq`  
*IRQ number.*
- struct `udd_memregion mem_regions` [UDD\_NR\_MAPS]  
*Array of memory regions defined by the device.*
- int(\* `open` )(struct `rtdm_fd` \*`fd`, int `oflags`)  
*Ancillary `open()` handler, optional.*

- `void(* close )(struct rtdm_fd *fd)`  
*Ancillary `close()` handler, optional.*
- `int(* ioctl )(struct rtdm_fd *fd, unsigned int request, void *arg)`  
*Ancillary `ioctl()` handler, optional.*
- `int(* mmap )(struct rtdm_fd *fd, struct vm_area_struct *vma)`  
*Ancillary `mmap()` handler for the mapper device, optional.*
- `int(* interrupt )(struct udd_device *udd)`

### 7.44.1 Detailed Description

UDD device descriptor.

This descriptor defines the characteristics of a UDD-based mini-driver when registering via a call to `udd_register_device()`.

### 7.44.2 Field Documentation

#### 7.44.2.1 close

```
void(* udd_device::close) (struct rtdm_fd *fd)
```

Ancillary `close()` handler, optional.

See `rtdm_close_handler()`.

Note

This handler is called from secondary mode only.

#### 7.44.2.2 device\_flags

```
int udd_device::device_flags
```

Additional device flags (e.g.

`RTDM_EXCLUSIVE`) `RTDM_NAMED_DEVICE` may be omitted).

Referenced by `udd_register_device()`.



## 7.44.2.3 device\_subclass

```
int udd_device::device_subclass
```

Subclass code of the device managed by the mini-driver (see `RTDM_SUBCLASS_XXX` definition in the [Device Profiles](#)).

The main class code is pre-set to `RTDM_CLASS_UDD`.

## 7.44.2.4 interrupt

```
int(* udd_device::interrupt) (struct udd_device *udd)
```

Ancillary handler for receiving interrupts. This handler must be provided if the mini-driver hands over IRQ handling to the UDD core, by setting the `irq` field to a valid value, different from `UDD_IRQ_CUSTOM` and `UDD_IRQ_NONE`.

The `->interrupt()` handler shall return one of the following status codes:

- `RTDM_IRQ_HANDLED`, if the mini-driver successfully handled the IRQ. This flag can be combined with `RTDM_IRQ_DISABLE` to prevent the Cobalt kernel from re-enabling the interrupt line upon return, otherwise it is re-enabled automatically.
- `RTDM_IRQ_NONE`, if the interrupt does not match any IRQ the mini-driver can handle.

Once the `->interrupt()` handler has returned, the UDD core notifies user-space Cobalt threads waiting for IRQ events (if any).

Note

This handler is called from primary mode only.

Referenced by `udd_register_device()`.

## 7.44.2.5 ioctl

```
int(* udd_device::ioctl) (struct rtdm_fd *fd, unsigned int request, void *arg)
```

Ancillary `ioctl()` handler, optional.

See `rtdm_ioctl_handler()`.

If this routine returns `-ENOSYS`, the default action implemented by the UDD core for the corresponding request will be applied, as if no `ioctl` handler had been defined.

Note

This handler is called from primary mode only.

## 7.44.2.6 irq

```
int udd_device::irq
```

IRQ number.

If valid, the UDD core manages the corresponding interrupt line, installing a base handler. Otherwise, a special value can be passed for declaring [unmanaged IRQs](#).

Referenced by `udd_register_device()`.

## 7.44.2.7 mem\_regions

```
struct udd_memregion udd_device::mem_regions[UDD_NR_MAPS]
```

Array of memory regions defined by the device.

The array can be sparse, with some entries bearing the `UDD_MEM_NONE` type interleaved with valid ones. See the discussion about [UDD memory regions](#).

## 7.44.2.8 mmap

```
int(* udd_device::mmap) (struct rtdm_fd *fd, struct vm_area_struct *vma)
```

Ancillary `mmap()` handler for the mapper device, optional.

See `rtdm_mmap_handler()`. The mapper device operates on a valid region defined in the `mem_regions[]` array. A pointer to the region can be obtained by a call to `udd_get_region()`.

If this handler is NULL, the UDD core establishes the mapping automatically, depending on the memory type defined for the region.

Note

This handler is called from secondary mode only.

## 7.44.2.9 open

```
int(* udd_device::open) (struct rtdm_fd *fd, int oflags)
```

Ancillary `open()` handler, optional.

See `rtdm_open_handler()`.

Note

This handler is called from secondary mode only.

The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/udd.h](#)

## 7.45 udd\_memregion Struct Reference

### Data Fields

- const char \* [name](#)  
*Name of the region (informational but required)*
- unsigned long [addr](#)  
*Start address of the region.*
- size\_t [len](#)  
*Length (in bytes) of the region.*
- int [type](#)  
*Type of the region.*

### 7.45.1 Detailed Description

UDD memory region descriptor.

This descriptor defines the characteristics of a memory region declared to the UDD core by the mini-driver. All valid regions should be declared in the `udd_device.mem_regions[]` array, invalid/unassigned ones should bear the `UDD_MEM_NONE` type.

The UDD core exposes each region via the `mmap(2)` interface to the application. To this end, a companion mapper device is created automatically when registering the mini-driver.

The mapper device creates special files in the RTDM namespace for reaching the individual regions, which the application can open then map to its address space via the `mmap(2)` system call.

For instance, declaring a region of physical memory at index #2 of the memory region array could be done as follows:

```
static struct udd_device udd;

static int foocard_pci_probe(struct pci_dev *dev, const struct pci_device_id *id)
{
    udd.device_name = "foocard";
    ...
    udd.mem_regions[2].name = "ADC";
    udd.mem_regions[2].addr = pci_resource_start(dev, 1);
    udd.mem_regions[2].len = pci_resource_len(dev, 1);
    udd.mem_regions[2].type = UDD_MEM_PHYS;
    ...
    return udd_register_device(&udd);
}
```

This will make such region accessible via the mapper device using the following sequence of code (see note), via the default `->mmap()` handler from the UDD core:

```
int fd, fdm;
void *p;

fd = open("/dev/rtdm/foocard", O_RDWR);
fdm = open("/dev/rtdm/foocard,mapper2", O_RDWR);
p = mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_SHARED, fdm, 0);
```

if no valid region has been declared in the `udd_device.mem_regions[]` array, no mapper device is created.

### Note

The example code assumes that `cobalt_api` POSIX symbol wrapping is in effect, so that RTDM performs the memory mapping operation (not the regular kernel).

## 7.45.2 Field Documentation

### 7.45.2.1 addr

unsigned long udd\_memregion::addr

Start address of the region.

This may be a physical or virtual address, depending on the [memory type](#).

### 7.45.2.2 len

size\_t udd\_memregion::len

Length (in bytes) of the region.

This value must be PAGE\_SIZE aligned.

### 7.45.2.3 type

int udd\_memregion::type

Type of the region.

See the discussion about [UDD memory types](#) for possible values.

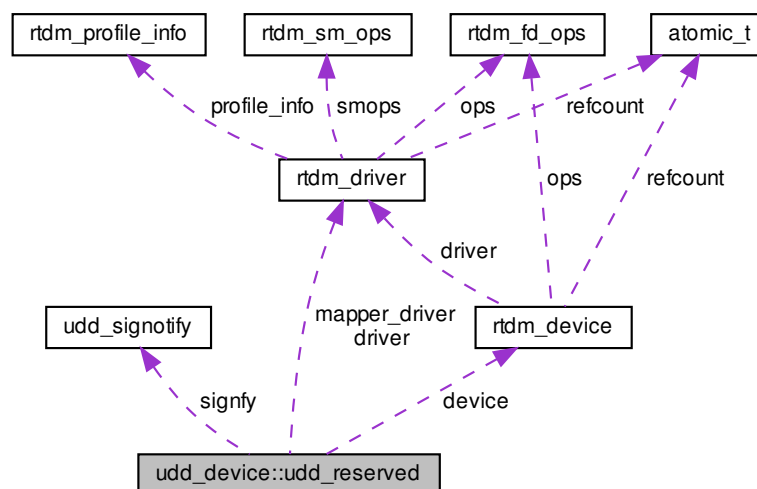
The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/udd.h](#)

## 7.46 udd\_device::udd\_reserved Struct Reference

Reserved to the UDD core.

Collaboration diagram for udd\_device::udd\_reserved:



### 7.46.1 Detailed Description

Reserved to the UDD core.

The documentation for this struct was generated from the following file:

- [include/cobalt/kernel/rtdm/udd.h](#)

## 7.47 udd\_signotify Struct Reference

UDD event notification descriptor.

### Data Fields

- [pid\\_t pid](#)  
*PID of the Cobalt thread to notify upon interrupt receipt.*
- [int sig](#)  
*Signal number to send to PID for notifying, which must be in the range [SIGRTMIN .*

### 7.47.1 Detailed Description

UDD event notification descriptor.

This structure shall be used to pass the information required to enable/disable the notification by signal upon interrupt receipt.

If PID is zero or negative, the notification is disabled. Otherwise, the Cobalt thread whose PID is given will receive the Cobalt signal also mentioned, along with the count of interrupts at the time of the receipt stored in `siginfo.si_int`. A Cobalt thread must explicitly wait for notifications using the `sigwaitinfo()` or `sigtimedwait()` services (no asynchronous mode available).

### 7.47.2 Field Documentation

#### 7.47.2.1 pid

```
pid_t udd_signotify::pid
```

PID of the Cobalt thread to notify upon interrupt receipt.

If `pid` is zero or negative, the notification is disabled.

### 7.47.2.2 sig

```
int udd_notify::sig
```

Signal number to send to PID for notifying, which must be in the range [SIGRTMIN .

. SIGRTMAX] inclusive. This value is not considered if *pid* is zero or negative.

The documentation for this struct was generated from the following file:

- include/rtdm/uapi/[udd.h](#)

## 7.48 xnsched Struct Reference

Scheduling information structure.

### Data Fields

- unsigned long [status](#)
- unsigned long [lflags](#)
- struct xnthread \* [curr](#)
- int [cpu](#)
- cpumask\_t [resched](#)
- struct xnsched\_rt [rt](#)
- volatile unsigned [inesting](#)
- struct xntimer [htimer](#)
- struct xntimer [rrbtimer](#)

### 7.48.1 Detailed Description

Scheduling information structure.

### 7.48.2 Field Documentation

#### 7.48.2.1 cpu

```
int xnsched::cpu
```

Mask of CPUs needing rescheduling.

## 7.48.2.2 curr

```
struct xntimer* xnsched::curr
```

Owner CPU id.

## 7.48.2.3 htimer

```
struct xntimer xnsched::htimer
```

Round-robin timer.

## 7.48.2.4 inesting

```
volatile unsigned xnsched::inesting
```

Host timer.

## 7.48.2.5 lflags

```
unsigned long xnsched::lflags
```

Current thread.

## 7.48.2.6 resched

```
cpumask_t xnsched::resched
```

Context of built-in real-time class.

## 7.48.2.7 rrbtimer

```
struct xntimer xnsched::rrbtimer
```

Root thread control block.

## 7.48.2.8 rt

```
struct xnsched_rt xnsched::rt
```

Interrupt nesting level.

### 7.48.2.9 status

unsigned long xnsched::status

< Scheduler specific status bitmask. Scheduler specific local flags bitmask.

The documentation for this struct was generated from the following file:

- include/cobalt/kernel/sched.h

## 7.49 xnvfile\_lock\_ops Struct Reference

Vfile locking operations.

### Data Fields

- int(\* [get](#))(struct xnvfile \*vfile)
- void(\* [put](#))(struct xnvfile \*vfile)

### 7.49.1 Detailed Description

Vfile locking operations.

This structure describes the operations to be provided for implementing locking support on vfiles. They apply to both snapshot-driven and regular vfiles.

### 7.49.2 Field Documentation

#### 7.49.2.1 get

int(\* xnvfile\_lock\_ops::get)(struct xnvfile \*vfile)

This handler should grab the desired lock.

Parameters

|              |                                                    |
|--------------|----------------------------------------------------|
| <i>vfile</i> | A pointer to the virtual file which needs locking. |
|--------------|----------------------------------------------------|



## Returns

zero should be returned if the call succeeds. Otherwise, a negative error code can be returned; upon error, the current vfile operation is aborted, and the user-space caller is passed back the error value.

7.49.2.2 `put`

```
void(* xnvfile_lock_ops::put) (struct xnvfile *vfile)
```

This handler should release the lock previously grabbed by the [get\(\) handler](#).

## Parameters

|                    |                                                                          |
|--------------------|--------------------------------------------------------------------------|
| <code>vfile</code> | A pointer to the virtual file which currently holds the lock to release. |
|--------------------|--------------------------------------------------------------------------|

The documentation for this struct was generated from the following file:

- `include/cobalt/kernel/vfile.h`

## 7.50 `xnvfile_regular_iterator` Struct Reference

Regular vfile iterator.

## Data Fields

- `loff_t pos`  
*Current record position while iterating.*
- `struct seq_file * seq`  
*Backlink to the host sequential file supporting the vfile.*
- `struct xnvfile_regular * vfile`  
*Backlink to the vfile being read.*
- `char private [0]`  
*Start of private area.*

### 7.50.1 Detailed Description

Regular vfile iterator.

This structure defines an iterator over a regular vfile.

### 7.50.2 Field Documentation

### 7.50.2.1 pos

```
loff_t xnvfile_regular_iterator::pos
```

Current record position while iterating.

### 7.50.2.2 private

```
char xnvfile_regular_iterator::private[0]
```

Start of private area.

Use `xnvfile_iterator_priv()` to address it.

### 7.50.2.3 seq

```
struct seq_file* xnvfile_regular_iterator::seq
```

Backlink to the host sequential file supporting the vfile.

### 7.50.2.4 vfile

```
struct xnvfile_regular* xnvfile_regular_iterator::vfile
```

Backlink to the vfile being read.

The documentation for this struct was generated from the following file:

- `include/cobalt/kernel/vfile.h`

## 7.51 xnvfile\_regular\_ops Struct Reference

Regular vfile operation descriptor.

### Data Fields

- `int(* rewind )(struct xnvfile\_regular\_iterator *it)`
- `void *(* begin )(struct xnvfile\_regular\_iterator *it)`
- `void *(* next )(struct xnvfile\_regular\_iterator *it)`
- `void(* end )(struct xnvfile\_regular\_iterator *it)`
- `int(* show )(struct xnvfile\_regular\_iterator *it, void *data)`
- `ssize_t(* store )(struct xnvfile\_input *input)`

### 7.51.1 Detailed Description

Regular vfile operation descriptor.

This structure describes the operations available with a regular vfile. It defines handlers for sending back formatted kernel data upon a user-space read request, and for obtaining user data upon a user-space write request.

### 7.51.2 Field Documentation

#### 7.51.2.1 `begin`

```
void>(* xnvfile_regular_ops::begin) (struct xnvfile\_regular\_iterator *it)
```

This handler should prepare for iterating over the records upon a read request, starting from the specified position.

Parameters

|           |                                                                                                                                                                                                                        |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>it</i> | A pointer to the current vfile iterator. On entry, <code>it-&gt;pos</code> is set to the (0-based) position of the first record to output. This handler may be called multiple times with different position requests. |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Returns

A pointer to the first record to format and output, to be passed to the [show\(\) handler](#) as its *data* parameter, if the call succeeds. Otherwise:

- NULL in case no record is available, in which case the read operation will terminate immediately with no output.
- `VFILE_SEQ_START`, a special value indicating that [the show\(\) handler](#) should receive a NULL data pointer first, in order to output a header.
- `ERR_PTR(errno)`, where `errno` is a negative error code; upon error, the current operation will be aborted immediately.

Note

This handler is optional; if none is given in the operation descriptor (i.e. NULL value), the [show\(\) handler\(\)](#) will be called only once for a read operation, with a NULL *data* parameter. This particular setting is convenient for simple regular vfiles having a single, fixed record to output.

#### 7.51.2.2 `end`

```
void(* xnvfile_regular_ops::end) (struct xnvfile\_regular\_iterator *it)
```

This handler is called after all records have been output.

## Parameters

|           |                                          |
|-----------|------------------------------------------|
| <i>it</i> | A pointer to the current vfile iterator. |
|-----------|------------------------------------------|

## Note

This handler is optional and the pointer may be NULL.

## 7.51.2.3 next

```
void>(* xnvfile_regular_ops::next) (struct xnvfile\_regular\_iterator *it)
```

This handler should return the address of the next record to format and output by the [show\(\)](#) handler".

## Parameters

|           |                                                                                                                                           |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <i>it</i> | A pointer to the current vfile iterator. On entry, <code>it-&gt;pos</code> is set to the (0-based) position of the next record to output. |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------|

## Returns

A pointer to the next record to format and output, to be passed to the [show\(\) handler](#) as its *data* parameter, if the call succeeds. Otherwise:

- NULL in case no record is available, in which case the read operation will terminate immediately with no output.
- ERR\_PTR(errno), where errno is a negative error code; upon error, the current operation will be aborted immediately.

## Note

This handler is optional; if none is given in the operation descriptor (i.e. NULL value), the read operation will stop after the first invocation of the [show\(\) handler](#).

## 7.51.2.4 rewind

```
int(* xnvfile_regular_ops::rewind) (struct xnvfile\_regular\_iterator *it)
```

This handler is called only once, when the virtual file is opened, before the [begin\(\) handler](#) is invoked.

## Parameters

|           |                                                                               |
|-----------|-------------------------------------------------------------------------------|
| <i>it</i> | A pointer to the vfile iterator which will be used to read the file contents. |
|-----------|-------------------------------------------------------------------------------|

## Returns

Zero should be returned upon success. Otherwise, a negative error code aborts the operation, and is passed back to the reader.

## Note

This handler is optional. It should not be used to allocate resources but rather to perform consistency checks, since no closure call is issued in case the open sequence eventually fails.

## 7.51.2.5 show

```
int(* xnvfile_regular_ops::show) (struct xnvfile_regular_iterator *it, void *data)
```

This handler should format and output a record.

`xnvfile_printf()`, `xnvfile_write()`, `xnvfile_puts()` and `xnvfile_putc()` are available to format and/or emit the output. All routines take the iterator argument *it* as their first parameter.

## Parameters

|             |                                                                                                                                                                                                                                                                                                                                |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>it</i>   | A pointer to the current vfile iterator.                                                                                                                                                                                                                                                                                       |
| <i>data</i> | A pointer to the record to format then output. The first call to the handler may receive a NULL <i>data</i> pointer, depending on the presence and/or return of a <a href="#">handler</a> ; the show handler should test this special value to output any header that fits, prior to receiving more calls with actual records. |

## Returns

zero if the call succeeds, also indicating that the handler should be called for the next record if any. Otherwise:

- A negative error code. This will abort the output phase, and return this status to the reader.
- `VFILE_SEQ_SKIP`, a special value indicating that the current record should be skipped and will not be output.

## 7.51.2.6 store

```
ssize_t(* xnvfile_regular_ops::store) (struct xnvfile_input *input)
```

This handler receives data written to the vfile, likely for updating some kernel setting, or triggering any other action which fits. This is the only handler which deals with the write-side of a vfile. It is called when writing to the `/proc` entry of the vfile from a user-space process.

The input data is described by a descriptor passed to the handler, which may be subsequently passed to parsing helper routines. For instance, `xnvfile_get_string()` will accept the input descriptor for returning the written data as a null-terminated character string. On the other hand, `xnvfile_get_integer()` will attempt to return a long integer from the input data.

## Parameters

|              |                                                                                              |
|--------------|----------------------------------------------------------------------------------------------|
| <i>input</i> | A pointer to an input descriptor. It refers to an opaque data from the handler's standpoint. |
|--------------|----------------------------------------------------------------------------------------------|

## Returns

the number of bytes read from the input descriptor if the call succeeds. Otherwise, a negative error code. Return values from parsing helper routines are commonly passed back to the caller by the [store\(\) handler](#).

## Note

This handler is optional, and may be omitted for read-only vfiles.

The documentation for this struct was generated from the following file:

- `include/cobalt/kernel/vfile.h`

## 7.52 `xnfile_rev_tag` Struct Reference

Snapshot revision tag.

### Data Fields

- `int rev`  
*Current revision number.*

### 7.52.1 Detailed Description

Snapshot revision tag.

This structure defines a revision tag to be used with [snapshot-driven vfiles](#).

### 7.52.2 Field Documentation

#### 7.52.2.1 `rev`

```
int xnfile_rev_tag::rev
```

Current revision number.


The documentation for this struct was generated from the following file:

- `include/cobalt/kernel/vfile.h`

## 7.53 xnvfile\_snapshot Struct Reference

Snapshot vfile descriptor.

Collaboration diagram for xnvfile\_snapshot:



### 7.53.1 Detailed Description

Snapshot vfile descriptor.

This structure describes a snapshot-driven vfile. Reading from such a vfile involves a preliminary data collection phase under lock protection, and a subsequent formatting and output phase of the collected data records. Locking is done in a way that does not increase worst-case latency, regardless of the number of records to be collected for output.


The documentation for this struct was generated from the following file:

- `include/cobalt/kernel/vfile.h`

## 7.54 xnvfile\_snapshot\_iterator Struct Reference

Snapshot-driven vfile iterator.

Collaboration diagram for `xnvfile_snapshot_iterator`:



## Data Fields

- `int nrdata`  
*Number of collected records.*
- `caddr_t databuf`  
*Address of record buffer.*
- `struct seq_file * seq`  
*Backlink to the host sequential file supporting the vfile.*
- `struct xnvfile_snapshot * vfile`  
*Backlink to the vfile being read.*
- `void(* endfn)(struct xnvfile_snapshot_iterator *it, void *buf)`  
*Buffer release handler.*
- `char private [0]`  
*Start of private area.*

### 7.54.1 Detailed Description

Snapshot-driven vfile iterator.

This structure defines an iterator over a snapshot-driven vfile.

### 7.54.2 Field Documentation



7.54.2.1 `databuf`

```
caddr_t xnvfile_snapshot_iterator::databuf
```

Address of record buffer.

7.54.2.2 `endfn`

```
void(* xnvfile_snapshot_iterator::endfn) (struct xnvfile\_snapshot\_iterator *it, void *buf)
```

Buffer release handler.

7.54.2.3 `nrdata`

```
int xnvfile_snapshot_iterator::nrdata
```

Number of collected records.

7.54.2.4 `private`

```
char xnvfile_snapshot_iterator::private[0]
```

Start of private area.

Use `xnvfile_iterator_priv()` to address it.

7.54.2.5 `seq`

```
struct seq_file* xnvfile_snapshot_iterator::seq
```

Backlink to the host sequential file supporting the vfile.

7.54.2.6 `vfile`

```
struct xnvfile\_snapshot* xnvfile_snapshot_iterator::vfile
```

Backlink to the vfile being read.

The documentation for this struct was generated from the following file:

- `include/cobalt/kernel/vfile.h`

## 7.55 xnvfile\_snapshot\_ops Struct Reference

Snapshot vfile operation descriptor.

### Data Fields

- `int(* rewind)(struct xnvfile_snapshot_iterator *it)`
- `void *(* begin)(struct xnvfile_snapshot_iterator *it)`
- `void(* end)(struct xnvfile_snapshot_iterator *it, void *buf)`
- `int(* next)(struct xnvfile_snapshot_iterator *it, void *data)`
- `int(* show)(struct xnvfile_snapshot_iterator *it, void *data)`
- `ssize_t(* store)(struct xnvfile_input *input)`

### 7.55.1 Detailed Description

Snapshot vfile operation descriptor.

This structure describes the operations available with a snapshot-driven vfile. It defines handlers for returning a printable snapshot of some Xenomai object contents upon a user-space read request, and for updating this object upon a user-space write request.

### 7.55.2 Field Documentation

#### 7.55.2.1 begin

```
void*(* xnvfile_snapshot_ops::begin)(struct xnvfile_snapshot_iterator *it)
```

This handler should allocate the snapshot buffer to hold records during the data collection phase. When specified, all records collected via the `next()` handler" will be written to a cell from the memory area returned by `begin()`.

#### Parameters

|           |                                             |
|-----------|---------------------------------------------|
| <i>it</i> | A pointer to the current snapshot iterator. |
|-----------|---------------------------------------------|

#### Returns

A pointer to the record buffer, if the call succeeds. Otherwise:

- NULL in case of allocation error. This will abort the data collection, and return -ENOMEM to the reader.
- VFILE\_SEQ\_EMPTY, a special value indicating that no record will be output. In such a case, the `next()` handler will not be called, and the data collection will stop immediately. However, the `show()` handler will still be called once, with a NULL data pointer (i.e. header display request).

## Note

This handler is optional; if none is given, an internal allocation depending on the value returned by the `rewind()` handler can be obtained.

7.55.2.2 `end`

```
void(* xnvfile_snapshot_ops::end) (struct xnvfile\_snapshot\_iterator *it, void *buf)
```

This handler releases the memory buffer previously obtained from `begin()`. It is usually called after the snapshot data has been output by `show()`, but it may also be called before rewinding the vfile after a revision change, to release the dropped buffer.

## Parameters

|            |                                             |
|------------|---------------------------------------------|
| <i>it</i>  | A pointer to the current snapshot iterator. |
| <i>buf</i> | A pointer to the buffer to release.         |

## Note

This routine is optional and the pointer may be NULL. It is not needed upon internal buffer allocation; see the description of the `rewind()` handler".

7.55.2.3 `next`

```
int(* xnvfile_snapshot_ops::next) (struct xnvfile\_snapshot\_iterator *it, void *data)
```

This handler fetches the next record, as part of the snapshot data to be sent back to the reader via the `show()`.

## Parameters

|             |                                             |
|-------------|---------------------------------------------|
| <i>it</i>   | A pointer to the current snapshot iterator. |
| <i>data</i> | A pointer to the record to fill in.         |

## Returns

a strictly positive value, if the call succeeds and leaves a valid record into `data`, which should be passed to the `show()` handler) during the formatting and output phase. Otherwise:

- A negative error code. This will abort the data collection, and return this status to the reader.
- `VFILE_SEQ_SKIP`, a special value indicating that the current record should be skipped. In such a case, the `data` pointer is not advanced to the next position before the `next()` handler is called anew.

## Note

This handler is called with the vfile lock held. Before each invocation of this handler, the vfile core checks whether the revision tag has been touched, in which case the data collection is restarted from scratch. A data collection phase succeeds whenever all records can be fetched via the [next\(\) handler](#), while the revision tag remains unchanged, which indicates that a consistent snapshot of the object state was taken.

## 7.55.2.4 rewind

```
int(* xnvfile_snapshot_ops::rewind) (struct xnvfile\_snapshot\_iterator *it)
```

This handler (re-)initializes the data collection, moving the seek pointer at the first record. When the file revision tag is touched while collecting data, the current reading is aborted, all collected data dropped, and the vfile is eventually rewind.

## Parameters

|           |                                                                                                                         |
|-----------|-------------------------------------------------------------------------------------------------------------------------|
| <i>it</i> | A pointer to the current snapshot iterator. Two useful information can be retrieved from this iterator in this context: |
|-----------|-------------------------------------------------------------------------------------------------------------------------|

- *it->vfile* is a pointer to the descriptor of the virtual file being rewind.
- `xnvfile_iterator_priv(it)` returns a pointer to the private data area, available from the descriptor, which size is `vfile->privsz`. If the latter size is zero, the returned pointer is meaningless and should not be used.

## Returns

A negative error code aborts the data collection, and is passed back to the reader. Otherwise:

- a strictly positive value is interpreted as the total number of records which will be returned by the [next\(\) handler](#) during the data collection phase. If no [begin\(\) handler](#) is provided in the [operation descriptor](#), this value is used to allocate the snapshot buffer internally. The size of this buffer would then be `vfile->datasz * value`.
- zero leaves the allocation to the [begin\(\) handler](#) if present, or indicates that no record is to be output in case such handler is not given.

## Note

This handler is optional; a NULL value indicates that nothing needs to be done for rewinding the vfile. It is called with the vfile lock held.

## 7.55.2.5 show

```
int(* xnvfile_snapshot_ops::show) (struct xnvfile\_snapshot\_iterator *it, void *data)
```

This handler should format and output a record from the collected data.

`xnvfile_printf()`, `xnvfile_write()`, `xnvfile_puts()` and `xnvfile_putc()` are available to format and/or emit the output. All routines take the iterator argument *it* as their first parameter.

## Parameters

|             |                                                                                                                                                                                                                                                              |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>it</i>   | A pointer to the current snapshot iterator.                                                                                                                                                                                                                  |
| <i>data</i> | A pointer to the record to format then output. The first call to the handler is always passed a NULL <i>data</i> pointer; the show handler should test this special value to output any header that fits, prior to receiving more calls with actual records. |

## Returns

zero if the call succeeds, also indicating that the handler should be called for the next record if any. Otherwise:

- A negative error code. This will abort the output phase, and return this status to the reader.
- `VFILE_SEQ_SKIP`, a special value indicating that the current record should be skipped and will not be output.

7.55.2.6 `store`

```
ssize_t(* xnvfile_snapshot_ops::store) (struct xnvfile_input *input)
```

This handler receives data written to the vfile, likely for updating the associated Xenomai object's state, or triggering any other action which fits. This is the only handler which deals with the write-side of a vfile. It is called when writing to the `/proc` entry of the vfile from a user-space process.

The input data is described by a descriptor passed to the handler, which may be subsequently passed to parsing helper routines. For instance, `xnvfile_get_string()` will accept the input descriptor for returning the written data as a null-terminated character string. On the other hand, `xnvfile_get_integer()` will attempt to return a long integer from the input data.

## Parameters

|              |                                                                                              |
|--------------|----------------------------------------------------------------------------------------------|
| <i>input</i> | A pointer to an input descriptor. It refers to an opaque data from the handler's standpoint. |
|--------------|----------------------------------------------------------------------------------------------|

## Returns

the number of bytes read from the input descriptor if the call succeeds. Otherwise, a negative error code. Return values from parsing helper routines are commonly passed back to the caller by the `store()` handler.

## Note

This handler is optional, and may be omitted for read-only vfiles.

The documentation for this struct was generated from the following file:

- `include/cobalt/kernel/vfile.h`




## Chapter 8

# File Documentation

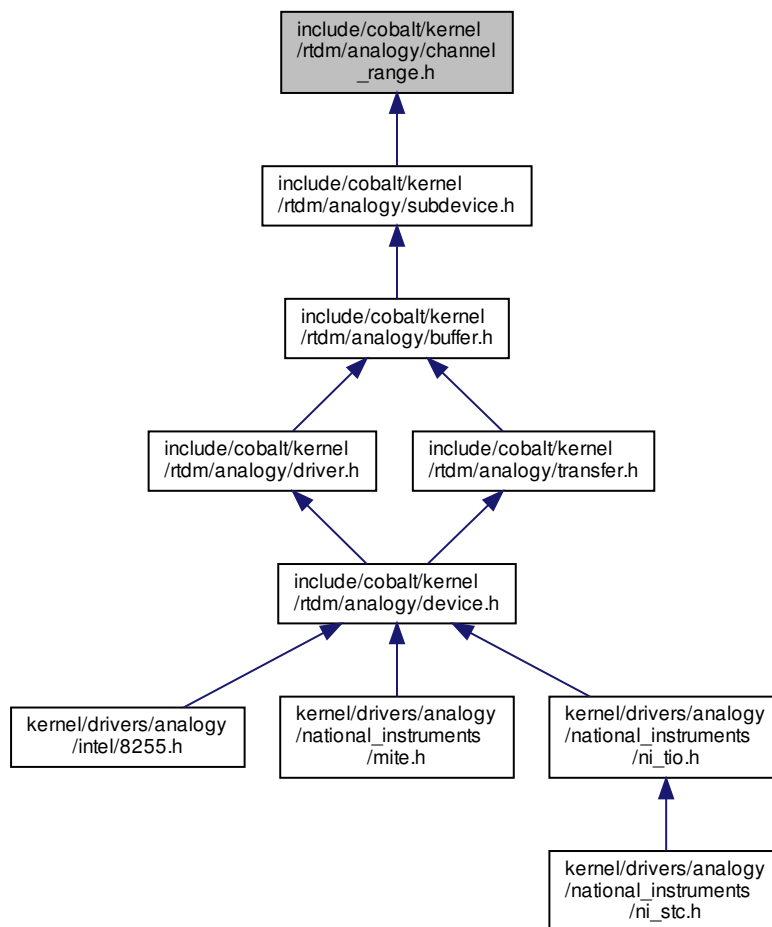
### 8.1 `include/cobalt/kernel/rtdm/analogy/channel_range.h` File Reference

Analogy for Linux, channel, range related features.

Include dependency graph for `channel_range.h`:



This graph shows which files directly or indirectly include this file:



## Data Structures

- struct [a4l\\_channel](#)  
*Structure describing some channel's characteristics.*
- struct [a4l\\_channels\\_desc](#)  
*Structure describing a channels set.*
- struct [a4l\\_range](#)  
*Structure describing a (unique) range.*

## Macros

- #define [A4L\\_CHAN\\_GLOBAL](#) 0x10  
*Internal use flag (must not be used by driver developer)*
- #define [A4L\\_RNG\\_GLOBAL](#) 0x8  
*Internal use flag (must not be used by driver developer)*
- #define [RANGE](#)(x, y)



- Macro to declare a (unique) range with no unit defined.
- #define `RANGE_V(x, y)`  
Macro to declare a (unique) range in Volt.
- #define `RANGE_mA(x, y)`  
Macro to declare a (unique) range in milliAmpere.
- #define `RANGE_ext(x, y)`  
Macro to declare a (unique) range in some external reference.
- #define `A4L_RNG_GLOBAL_RNGDESC 0`  
Constant to define a ranges descriptor as global (inter-channel)
- #define `A4L_RNG_PERCHAN_RNGDESC 1`  
Constant to define a ranges descriptor as specific for a channel.
- #define `RNG_GLOBAL(x)`  
Macro to declare a ranges global descriptor in one line.

### Channel reference

Flags to define the channel's reference

- #define `A4L_CHAN_AREF_GROUND 0x1`  
Ground reference.
- #define `A4L_CHAN_AREF_COMMON 0x2`  
Common reference.
- #define `A4L_CHAN_AREF_DIFF 0x4`  
Differential reference.
- #define `A4L_CHAN_AREF_OTHER 0x8`  
Misc reference.

### Channels declaration mode

Constant to define whether the channels in a descriptor are identical

- #define `A4L_CHAN_GLOBAL_CHANDESC 0`  
Global declaration, the set contains channels with similar characteristics.
- #define `A4L_CHAN_PERCHAN_CHANDESC 1`  
Per channel declaration, the decriptor gathers differents channels.

## 8.1.1 Detailed Description

Analogy for Linux, channel, range related features.

Copyright (C) 1997-2000 David A. Schleef [ds@schleef.org](mailto:ds@schleef.org) Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

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.

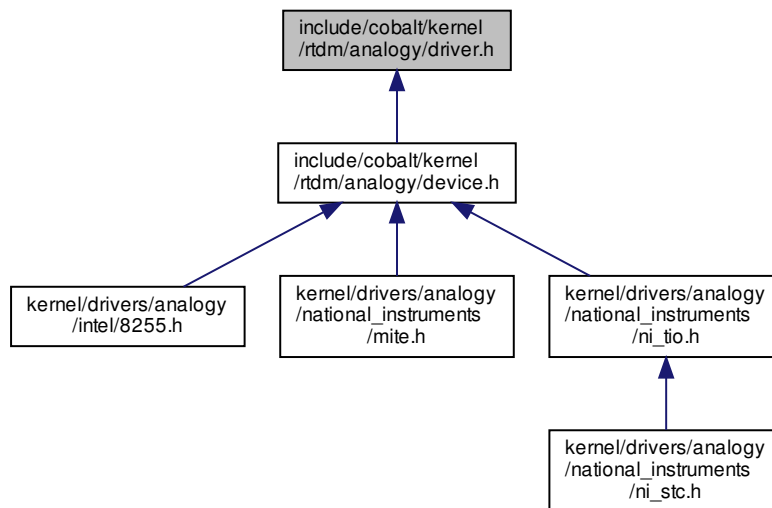
## 8.2 include/cobalt/kernel/rtdm/analogy/driver.h File Reference

Analogy for Linux, driver facilities.

Include dependency graph for driver.h:



This graph shows which files directly or indirectly include this file:



### Data Structures

- struct [a4l\\_driver](#)

*Structure containing driver declaration data.*

### 8.2.1 Detailed Description

Analogy for Linux, driver facilities.

Copyright (C) 1997-2000 David A. Schlee [ds@schlee.org](mailto:ds@schlee.org) Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

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.


## 8.3 include/cobalt/kernel/rtdm/driver.h File Reference

Real-Time Driver Model for Xenomai, driver API header.

Include dependency graph for driver.h:



This graph shows which files directly or indirectly include this file:



### Data Structures

- struct [rtdm\\_dev\\_context](#)  
*Device context.*
- struct [rtdm\\_profile\\_info](#)  
*RTDM profile information.*
- struct [rtdm\\_sm\\_ops](#)  
*RTDM state management handlers.*
- struct [rtdm\\_driver](#)  
*RTDM driver.*
- struct [rtdm\\_device](#)  
*RTDM device.*

## Macros

- #define [RTDM\\_MAX\\_MINOR](#) 4096  
*Maximum number of named devices per driver.*
- #define [RTDM\\_PROFILE\\_INFO](#)(\_\_name, \_\_id, \_\_subid, \_\_version)  
*Initializer for class profile information.*
- #define [cobalt\\_atomic\\_enter](#)(\_\_context)  
*Enter atomic section (dual kernel only)*
- #define [cobalt\\_atomic\\_leave](#)(\_\_context)  
*Leave atomic section (dual kernel only)*
- #define [RTDM\\_EXECUTE\\_ATOMICALY](#)(code\_block)  
  
*Execute code block atomically (DEPRECATED)*
- #define [RTDM\\_LOCK\\_UNLOCKED](#)(\_\_name) IPIPE\_SPIN\_LOCK\_UNLOCKED  
*Static lock initialisation.*
- #define [rtdm\\_lock\\_get\\_irqsave](#)(\_\_lock, \_\_context) ((\_\_context) = \_\_rtdm\_lock\_get\_irqsave(\_\_↵  
lock))  
  
*Acquire lock and disable preemption, by stalling the head domain.*
- #define [rtdm\\_lock\\_irqsave](#)(\_\_context) [splhigh](#)(\_\_context)  
  
*Disable preemption locally*
- #define [rtdm\\_lock\\_irqrestore](#)(\_\_context) [splexit](#)(\_\_context)  
  
*Restore preemption state*
- #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\\_FIXED\\_MINOR](#) 0x0002  
*Use fixed minor provided in the [rtdm\\_device](#) description for registering.*
- #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.*
- #define [RTDM\\_SECURE\\_DEVICE](#) 0x80000000  
*Flag indicating a secure variant of RTDM (not supported here)*

## RTDM\_IRQTYPE\_xxx

*Interrupt registrations flags*

- #define [RTDM\\_IRQTYPE\\_SHARED](#) XN\_IRQTYPE\_SHARED  
*Enable IRQ-sharing with other real-time drivers.*
- #define [RTDM\\_IRQTYPE\\_EDGE](#) XN\_IRQTYPE\_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_IRQ_NONE`  
*Unhandled interrupt.*
- #define `RTDM_IRQ_HANDLED` `XN_IRQ_HANDLED`  
*Denote handled interrupt.*
- #define `RTDM_IRQ_DISABLE` `XN_IRQ_DISABLE`  
*Request interrupt disabling on exit.*

**Task Priority Range**

Maximum and minimum task priorities

- #define `RTDM_TASK_LOWEST_PRIORITY` `0`
- #define `RTDM_TASK_HIGHEST_PRIORITY` `99`

**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 `ipipe_spinlock_t` `rtdm_lock_t`  
*Lock variable.*
- typedef `unsigned long` `rtdm_lockctx_t`  
*Variable to save the context while holding a lock.*
- 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.*

## Enumerations

**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` }

**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` }

## Functions

- static void \* [rtdm\\_fd\\_to\\_private](#) (struct [rtdm\\_fd](#) \*fd)
  - Locate the driver private area associated to a device context structure.*
- static struct [rtdm\\_fd](#) \* [rtdm\\_private\\_to\\_fd](#) (void \*dev\_private)
  - Locate a device file descriptor structure from its driver private area.*
- static bool [rtdm\\_fd\\_is\\_user](#) (struct [rtdm\\_fd](#) \*fd)
  - Tell whether the passed file descriptor belongs to an application.*
- static struct [rtdm\\_device](#) \* [rtdm\\_fd\\_device](#) (struct [rtdm\\_fd](#) \*fd)
  - Locate a device structure from a file descriptor.*
- int [rtdm\\_drv\\_set\\_sysclass](#) (struct [rtdm\\_driver](#) \*drv, struct class \*cls)
  - Set the kernel device class of a RTDM driver.*
- int [rtdm\\_dev\\_register](#) (struct [rtdm\\_device](#) \*device)
  - Register a RTDM device*
- void [rtdm\\_dev\\_unregister](#) (struct [rtdm\\_device](#) \*device)
  - Unregister a RTDM device*
- void [rtdm\\_toseq\\_init](#) ([rtdm\\_toseq\\_t](#) \*timeout\_seq, [nanosecs\\_rel\\_t](#) timeout)
  - Initialise a timeout sequence*
- static void [rtdm\\_lock\\_init](#) ([rtdm\\_lock\\_t](#) \*lock)
  - Dynamic lock initialisation*
- static void [rtdm\\_lock\\_get](#) ([rtdm\\_lock\\_t](#) \*lock)
  - Acquire lock from non-preemptible contexts*
- static void [rtdm\\_lock\\_put](#) ([rtdm\\_lock\\_t](#) \*lock)
  - Release lock without preemption restoration*
- static void [rtdm\\_lock\\_put\\_irqrestore](#) ([rtdm\\_lock\\_t](#) \*lock, [rtdm\\_lockctx\\_t](#) context)
  - Release lock and restore preemption state*
- 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*
- void [rtdm\\_schedule\\_nrt\\_work](#) (struct [work\\_struct](#) \*lostage\_work)
  - Put a work task in Linux non real-time global workqueue from primary mode.*
- 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_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_busy_sleep` (nanosecs\_rel\_t delay)

*Busy-wait a specified amount of time*

- void `rtdm_event_init` (rtdm\_event\_t \*event, unsigned long pending)

*Initialise an event*

- int `rtdm_event_wait` (rtdm\_event\_t \*event)

*Wait on event occurrence*

- int `rtdm_event_timedwait` (rtdm\_event\_t \*event, nanosecs\_rel\_t timeout, rtdm\_toseq\_t \*timeout\_↵\_seq)

*Wait on event occurrence with timeout*

- void `rtdm_event_signal` (rtdm\_event\_t \*event)

*Signal an event occurrence*

- void `rtdm_event_clear` (rtdm\_event\_t \*event)

*Clear event state*

- void `rtdm_event_pulse` (rtdm\_event\_t \*event)

*Signal an event occurrence to currently listening waiters*

- void `rtdm_event_destroy` (rtdm\_event\_t \*event)

*Destroy an event*

- void `rtdm_sem_init` (rtdm\_sem\_t \*sem, unsigned long value)

*Initialise a semaphore*

- int `rtdm_sem_down` (rtdm\_sem\_t \*sem)

*Decrement a semaphore*

- int `rtdm_sem_timeddown` (rtdm\_sem\_t \*sem, nanosecs\_rel\_t timeout, rtdm\_toseq\_t \*timeout\_↵\_seq)

*Decrement a semaphore with timeout*

- void `rtdm_sem_up` (rtdm\_sem\_t \*sem)

*Increment a semaphore*

- void `rtdm_sem_destroy` (rtdm\_sem\_t \*sem)

*Destroy a semaphore*

- void `rtdm_mutex_init` (rtdm\_mutex\_t \*mutex)

*Initialise a mutex*

- int `rtdm_mutex_lock` (rtdm\_mutex\_t \*mutex)

*Request a mutex*

- int `rtdm_mutex_timedlock` (rtdm\_mutex\_t \*mutex, nanosecs\_rel\_t timeout, rtdm\_toseq\_↵\_t \*timeout\_seq)

*Request a mutex with timeout*

- void `rtdm_mutex_unlock` (`rtdm_mutex_t *mutex`)

*Release a mutex*

- void `rtdm_mutex_destroy` (`rtdm_mutex_t *mutex`)

*Destroy a mutex*

### 8.3.1 Detailed Description

Real-Time Driver Model for Xenomai, driver API header.

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 Chantepedrix [gilles.chantepedrix@xenomai.org](mailto:gilles.chantepedrix@xenomai.org) Copyright (C) 2014 Philippe Gerum [rpm@xenomai.org](mailto:rpm@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.

### 8.3.2 Macro Definition Documentation

#### 8.3.2.1 RTDM\_PROFILE\_INFO

```
#define RTDM_PROFILE_INFO(
    __name,
    __id,
    __subid,
    __version )
```

#### Value:

```
{
    .name = ( # __name ),
    .class_id = (__id),
    .subclass_id = (__subid),
    .version = (__version),
    .magic = ~RTDM_CLASS_MAGIC,
    .owner = THIS_MODULE,
    .kdev_class = NULL,
}
```

Initializer for class profile information.

This macro must be used to fill in the [class profile information](#) field from a RTDM driver.



## Parameters

|                        |                                                                                 |
|------------------------|---------------------------------------------------------------------------------|
| <code>__name</code>    | Class name (unquoted).                                                          |
| <code>__id</code>      | Class major identification number ( <code>profile_version.class_id</code> ).    |
| <code>__subid</code>   | Class minor identification number ( <code>profile_version.subclass_id</code> ). |
| <code>__version</code> | Profile version number.                                                         |

## Note

See [Device Profiles](#).

## 8.3.3 Function Documentation

8.3.3.1 `rtdm_fd_device()`

```
static struct rtdm\_device* rtdm_fd_device (
    struct rtdm_fd * fd ) [inline], [static]
```

Locate a device structure from a file descriptor.

## Parameters

|    |           |                 |
|----|-----------|-----------------|
| in | <i>fd</i> | File descriptor |
|----|-----------|-----------------|

## Returns

The address of the device structure to which this file descriptor is attached.

8.3.3.2 `rtdm_fd_is_user()`

```
static bool rtdm_fd_is_user (
    struct rtdm_fd * fd ) [inline], [static]
```

Tell whether the passed file descriptor belongs to an application.

## Parameters

|    |           |                 |
|----|-----------|-----------------|
| in | <i>fd</i> | File descriptor |
|----|-----------|-----------------|

## Returns

true if passed file descriptor belongs to an application, false otherwise.

8.3.3.3 `rtm_fd_to_private()`

```
static void* rtm_fd_to_private (
    struct rtm_fd * fd ) [inline], [static]
```

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

## Parameters

|    |           |                                                                  |
|----|-----------|------------------------------------------------------------------|
| in | <i>fd</i> | File descriptor structure associated with opened device instance |
|----|-----------|------------------------------------------------------------------|

## Returns

The address of the private driver area associated to *file* descriptor.

8.3.3.4 `rtm_private_to_fd()`

```
static struct rtm_fd* rtm_private_to_fd (
    void * dev_private ) [inline], [static]
```

Locate a device file descriptor structure from its driver private area.

## Parameters

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

## Returns

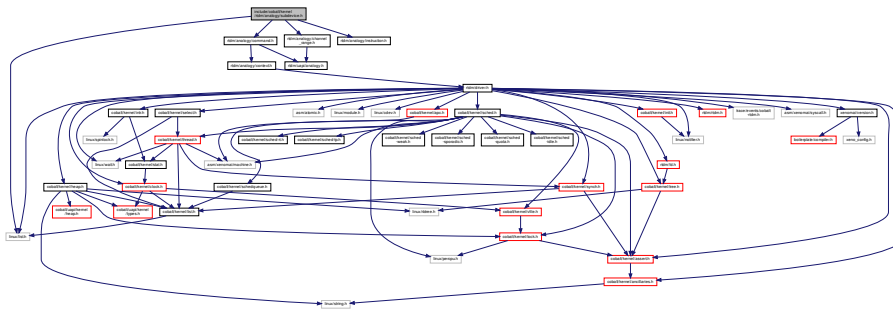
The address of the file descriptor structure defining *dev\_private*.

References `rtm_dev_context::dev_private`.

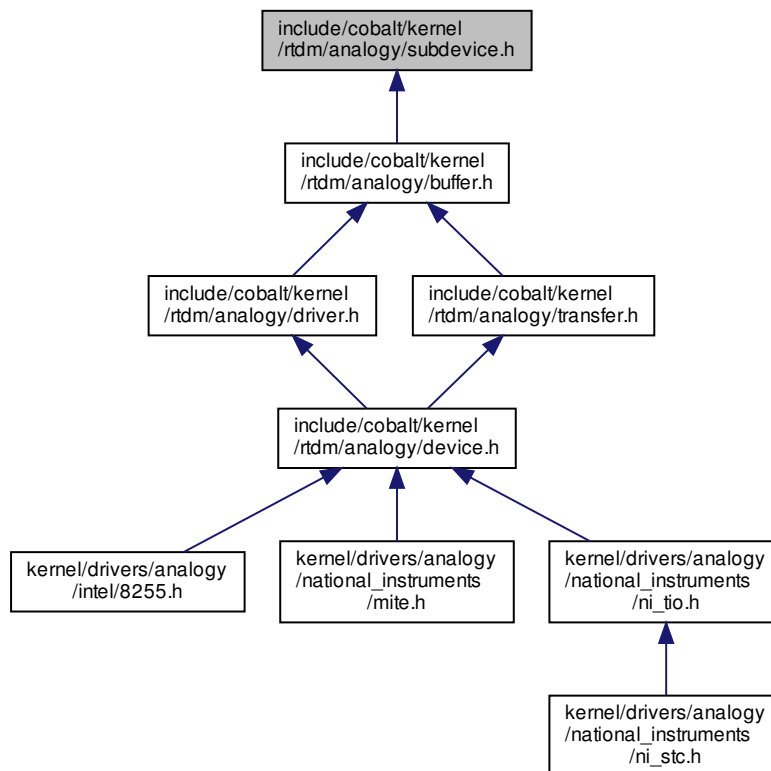
8.4 `include/cobalt/kernel/rtm/analogy/subdevice.h` File Reference

Analogy for Linux, subdevice related features.

Include dependency graph for subdevice.h:



This graph shows which files directly or indirectly include this file:



## Data Structures

- struct [a4l\\_subdevice](#)

*Structure describing the subdevice.*

### 8.4.1 Detailed Description

Analogy for Linux, subdevice related features.

Copyright (C) 1997-2000 David A. Schlee [ds@schleeef.org](mailto:ds@schleeef.org) Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

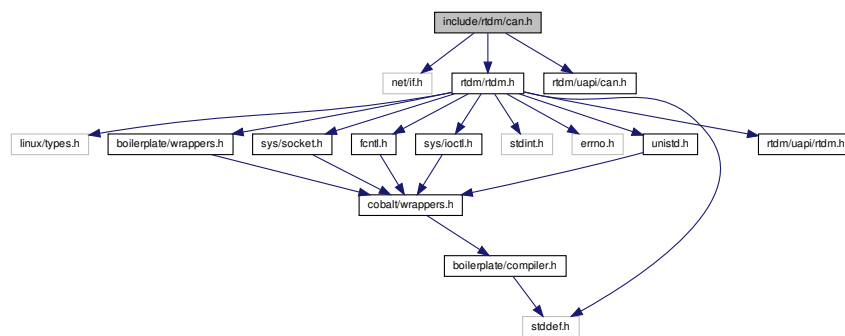
Xenomai is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

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.

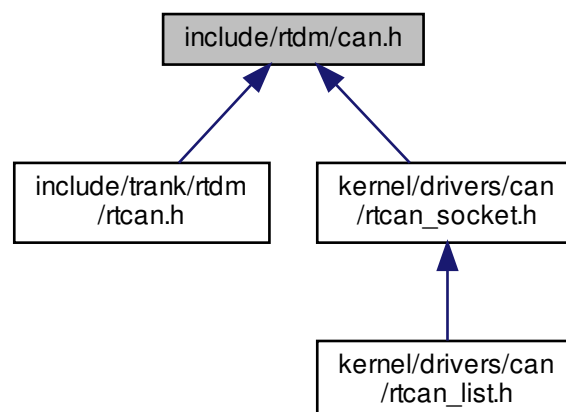
You should have received a copy of the GNU General Public License along with Xenomai; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 8.5 include/rtdm/can.h File Reference

Include dependency graph for can.h:



This graph shows which files directly or indirectly include this file:



### 8.5.1 Detailed Description

#### 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 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

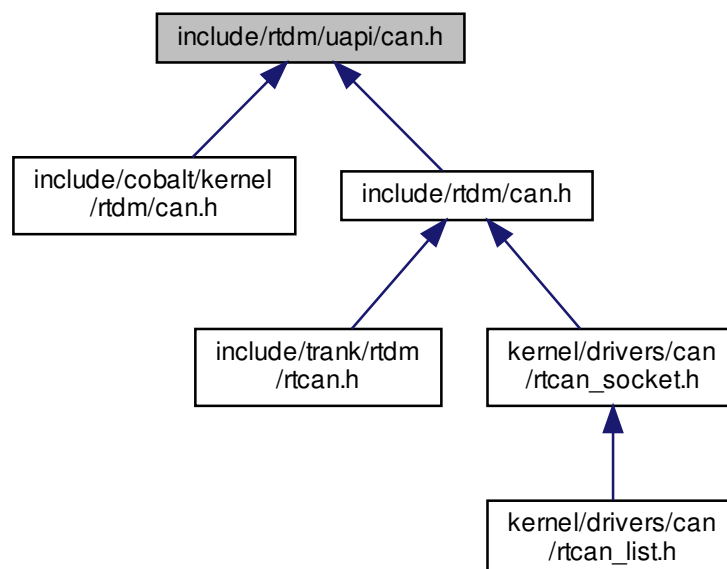
General Public License for more details.

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.

## 8.6 include/rtdm/uapi/can.h File Reference

Real-Time Driver Model for RT-Socket-CAN, CAN device profile header.

This graph shows which files directly or indirectly include this file:



## 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.*
- struct `can_ifreq`  
*CAN interface request descriptor.*

## 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

**Deprecated** *Passing struct ifreq as a request descriptor for CAN IOCTLs is still accepted for backward compatibility, however it is recommended to switch to struct [can\\_ifreq](#) at the first opportunity.*

- #define [SIOCGIFINDEX](#) defined\_by\_kernel\_header\_file  
  
*Get CAN interface index by name*
- #define [SIOCSCANBAUDRATE](#) \_IOW(RTIOC\_TYPE\_CAN, 0x01, struct can\_ifreq)  
  
*Set baud rate*
- #define [SIOCGCANBAUDRATE](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x02, struct can\_ifreq)  
  
*Get baud rate*
- #define [SIOCSCANCUSTOMBITTIME](#) \_IOW(RTIOC\_TYPE\_CAN, 0x03, struct can\_ifreq)  
  
*Set custom bit time parameter*
- #define [SIOCGCANCUSTOMBITTIME](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x04, struct can\_ifreq)  
  
*Get custom bit-time parameters*
- #define [SIOCSCANMODE](#) \_IOW(RTIOC\_TYPE\_CAN, 0x05, struct can\_ifreq)  
  
*Set operation mode of CAN controller*
- #define [SIOCGCANSTATE](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x06, struct can\_ifreq)

- Get current state of CAN controller*

  - #define [SIOCSCTRLMODE](#) \_IOW(RTIOC\_TYPE\_CAN, 0x07, struct can\_ifreq)
- Set special controller modes*

  - #define [SIOCGCANCTRLMODE](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x08, struct can\_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*

## 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\\_ACTIVE](#) = 0, [CAN\\_STATE\\_ERROR\\_WARNING](#)  
= 1, [CAN\\_STATE\\_BUS\\_WARNING](#) = 1,  
[CAN\\_STATE\\_ERROR\\_PASSIVE](#) = 2, [CAN\\_STATE\\_BUS\\_PASSIVE](#) = 2, [CAN\\_STATE\\_BUS\\_OFF](#),  
[CAN\\_STATE\\_SCANNING\\_BAUDRATE](#),  
[CAN\\_STATE\\_STOPPED](#), [CAN\\_STATE\\_SLEEPING](#) }

## 8.6.1 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.

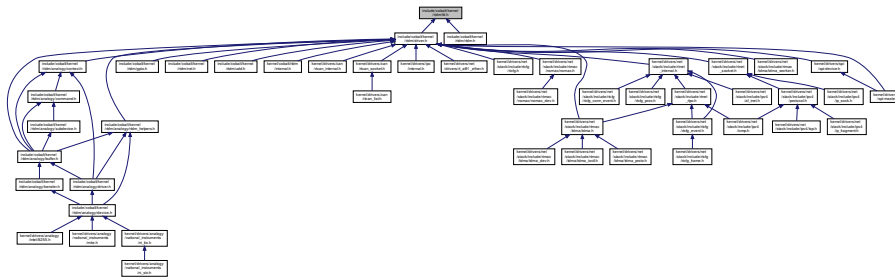
You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 8.7 include/cobalt/kernel/rtdm/fd.h File Reference

Include dependency graph for fd.h:



This graph shows which files directly or indirectly include this file:



### Data Structures

- struct `rtdm_fd_ops`  
*RTDM file operation descriptor.*

### Functions

- int `rtdm_open_handler` (struct `rtdm_fd` \*fd, int oflags)  
*Open handler for named devices.*
- int `rtdm_socket_handler` (struct `rtdm_fd` \*fd, int protocol)  
*Socket creation handler for protocol devices.*
- void `rtdm_close_handler` (struct `rtdm_fd` \*fd)  
*Close handler.*
- int `rtdm_ioctl_handler` (struct `rtdm_fd` \*fd, unsigned int request, void \_\_user \*arg)  
*IOCTL handler.*
- ssize\_t `rtdm_read_handler` (struct `rtdm_fd` \*fd, void \_\_user \*buf, size\_t size)  
*Read handler.*
- ssize\_t `rtdm_write_handler` (struct `rtdm_fd` \*fd, const void \_\_user \*buf, size\_t size)  
*Write handler.*
- ssize\_t `rtdm_recvmsg_handler` (struct `rtdm_fd` \*fd, struct `user_msghdr` \*msg, int flags)  
*Receive message handler.*
- ssize\_t `rtdm_sendmsg_handler` (struct `rtdm_fd` \*fd, const struct `user_msghdr` \*msg, int flags)  
*Transmit message handler.*

- int [rtdm\\_select\\_handler](#) (struct rtdm\_fd \*fd, struct xnselector \*selector, unsigned int type, unsigned int index)  
*Select handler.*
- int [rtdm\\_mmap\\_handler](#) (struct rtdm\_fd \*fd, struct vm\_area\_struct \*vma)  
*Memory mapping handler.*
- unsigned long [rtdm\\_get\\_unmapped\\_area\\_handler](#) (struct rtdm\_fd \*fd, unsigned long len, unsigned long pgoff, unsigned long flags)  
*Allocate mapping region in address space.*
- struct rtdm\_fd \* [rtdm\\_fd\\_get](#) (int ufd, unsigned int magic)  
  
*Retrieve and lock a RTDM file descriptor*
- int [rtdm\\_fd\\_lock](#) (struct rtdm\_fd \*fd)  
  
*Hold a reference on a RTDM file descriptor*
- void [rtdm\\_fd\\_put](#) (struct rtdm\_fd \*fd)  
  
*Release a RTDM file descriptor obtained via [rtdm\\_fd\\_get\(\)](#)*
- void [rtdm\\_fd\\_unlock](#) (struct rtdm\_fd \*fd)  
  
*Drop a reference on a RTDM file descriptor*
- int [rtdm\\_fd\\_select](#) (int ufd, struct xnselector \*selector, unsigned int type)  
  
*Bind a selector to specified event types of a given file descriptor*

## 8.7.1 Detailed Description

operation handlers

## 8.7.2 Function Documentation

### 8.7.2.1 rtdm\_fd\_get()

```
struct rtdm_fd* rtdm_fd_get (
    int ufd,
    unsigned int magic )
```

Retrieve and lock a RTDM file descriptor

Parameters

|    |              |                                  |
|----|--------------|----------------------------------|
| in | <i>ufd</i>   | User-side file descriptor        |
| in | <i>magic</i> | Magic word for lookup validation |

## Returns

Pointer to the RTDM file descriptor matching *ufd*. Otherwise:

- `ERR_PTR(-EADV)` if the use-space handle is either invalid, or not managed by RTDM.
- `ERR_PTR(-EBADF)` if the underlying device is being torned down at the time of the call.

## Note

The file descriptor returned must be later released by a call to `rtm_fd_put()`.

## Tags

`unrestricted`

8.7.2.2 `rtm_fd_lock()`

```
int rtm_fd_lock (  
    struct rtm_fd * fd )
```

Hold a reference on a RTDM file descriptor

## Parameters

|    |           |                        |
|----|-----------|------------------------|
| in | <i>fd</i> | Target file descriptor |
|----|-----------|------------------------|

## Note

`rtm_fd_lock()` increments the reference counter of *fd*. You only need to call this function in special scenarios, e.g. when keeping additional references to the file descriptor that have different lifetimes. Only use `rtm_fd_lock()` on descriptors that are currently locked via an earlier `rtm_fd_get()/rtm_fd_lock()` or while running a device operation handler.

## Tags

`unrestricted`

8.7.2.3 `rtm_fd_put()`

```
void rtm_fd_put (  
    struct rtm_fd * fd )
```

Release a RTDM file descriptor obtained via `rtm_fd_get()`

## Parameters

|    |           |                                 |
|----|-----------|---------------------------------|
| in | <i>fd</i> | RTDM file descriptor to release |
|----|-----------|---------------------------------|

## Note

Every call to [rtdm\\_fd\\_get\(\)](#) must be matched by a call to [rtdm\\_fd\\_put\(\)](#).

## Tags

[unrestricted](#)

## 8.7.2.4 rtdm\_fd\_select()

```
int rtdm_fd_select (
    int ufd,
    struct xnselector * selector,
    unsigned int type )
```

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

|        |                 |                                                        |
|--------|-----------------|--------------------------------------------------------|
| in     | <i>ufd</i>      | User-side file descriptor to bind to                   |
| in,out | <i>selector</i> | Selector object that shall be bound to the given event |
| in     | <i>type</i>     | Event type the caller is interested in                 |

## Returns

0 on success, otherwise:

- -EBADF is returned if the file descriptor *ufd* cannot be resolved.
- -EINVAL is returned if *type* is invalid.

## Tags

[task-unrestricted](#)

### 8.7.2.5 rtdm\_fd\_unlock()

```
void rtdm_fd_unlock (  
    struct rtdm_fd * fd )
```

Drop a reference on a RTDM file descriptor



Parameters

|    |           |                        |
|----|-----------|------------------------|
| in | <i>fd</i> | Target file descriptor |
|----|-----------|------------------------|

Note

Every call to `rtdm_fd_lock()` must be matched by a call to `rtdm_fd_unlock()`.


Tags

`unrestricted`


## 8.8 include/rtdm/uapi/gpiopwm.h File Reference

Real-Time Driver Model for Xenomai, pwm header.

Include dependency graph for gpiopwm.h:



This graph shows which files directly or indirectly include this file:



### 8.8.1 Detailed Description

Real-Time Driver Model for Xenomai, pwm header.

Note

Copyright (C) 2015 Jorge Ramirez [jro@xenomai.org](mailto:jro@xenomai.org)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

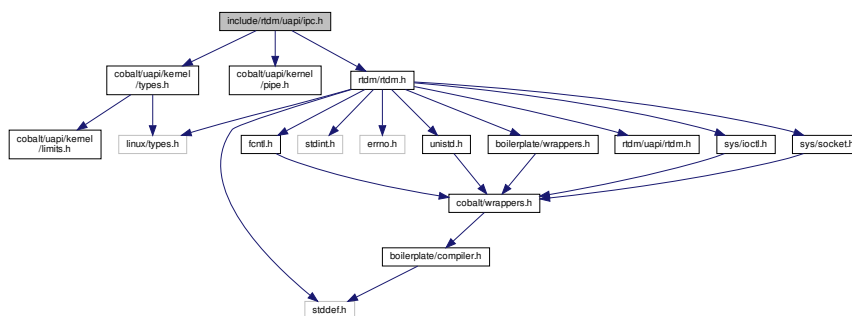
This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

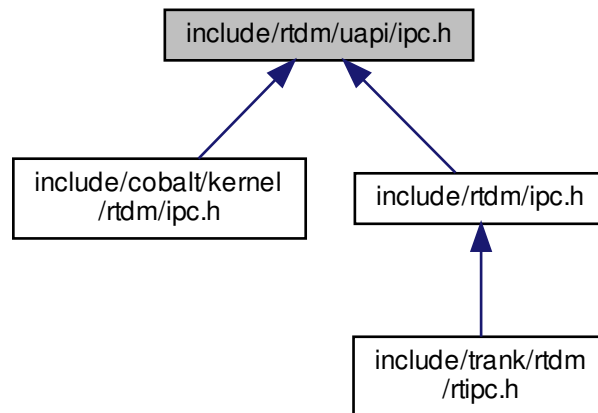
## 8.9 include/rtdm/uapi/ipc.h File Reference

This file is part of the Xenomai project.

Include dependency graph for ipc.h:



This graph shows which files directly or indirectly include this file:



## 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 [msghdr](#) \*msg, int flags)  
*Send a message on a RTIPC socket.*
- ssize\_t [recvmsg\\_\\_AF\\_RTIPC](#) (int sockfd, struct [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.*

### 8.9.1 Detailed Description

This file is part of the Xenomai project.

#### Note

Copyright (C) 2009 Philippe Gerum [rpm@xenomai.org](mailto:rpm@xenomai.org)

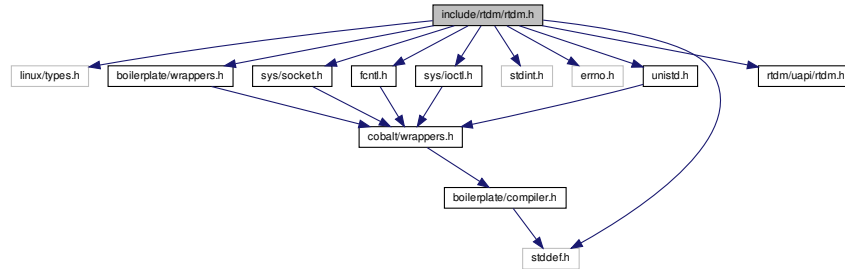
This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

## 8.10 include/rtdm/rtdm.h File Reference

Include dependency graph for rtdm.h:



This graph shows which files directly or indirectly include this file:



### 8.10.1 Detailed Description

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.

## 8.11 include/rtdm/uapi/rtdm.h File Reference

Real-Time Driver Model for Xenomai, user API header.

This graph shows which files directly or indirectly include this file:



## Data Structures

- struct [rtdm\\_device\\_info](#)  
*Device information.*

## Macros

### API Versioning

- #define [RTDM\\_API\\_VER](#) 9  
*Common user and driver API version.*
- #define [RTDM\\_API\\_MIN\\_COMPAT\\_VER](#) 9  
*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\\_COBALT](#) 8
- #define [RTDM\\_CLASS\\_UDD](#) 9
- #define [RTDM\\_CLASS\\_MEMORY](#) 10
- #define [RTDM\\_CLASS\\_GPIO](#) 11
- #define [RTDM\\_CLASS\\_SPI](#) 12
- #define [RTDM\\_CLASS\\_PWM](#) 13
- #define [RTDM\\_CLASS\\_MISC](#) 223
- #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 rtdm\_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.*

### 8.11.1 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)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

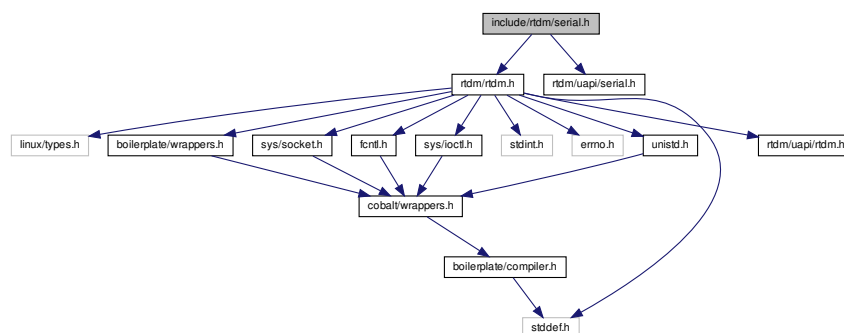
This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

## 8.12 include/rtdm/serial.h File Reference


Real-Time Driver Model for Xenomai, serial device profile header.

Include dependency graph for serial.h:





This graph shows which files directly or indirectly include this file:



### 8.12.1 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.

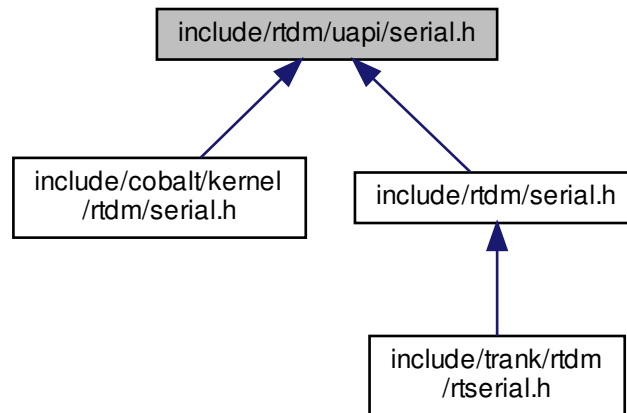
Xenomai is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

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.

## 8.13 include/rtdm/uapi/serial.h File Reference

Real-Time Driver Model for Xenomai, serial device profile header.

This graph shows which files directly or indirectly include this file:



## 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.*

### 8.13.1 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)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

### 8.13.2 Macro Definition Documentation

#### 8.13.2.1 RTSER\_RTIOC\_BREAK\_CTL

```
#define RTSER_RTIOC_BREAK_CTL _IOR(RTIOC_TYPE_SERIAL, 0x06, int)
```

Set or clear break on UART output line

Parameters

|    |     |                                          |
|----|-----|------------------------------------------|
| in | arg | RTSER_BREAK_SET or RTSER_BREAK_CLR (int) |
|----|-----|------------------------------------------|

Returns

0 on success, otherwise negative error code

Tags

[task-unrestricted](#)

Note

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

## 8.13.2.2 RTSER\_RTIOC\_GET\_CONFIG

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

Get serial device configuration

Parameters

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

Returns

0 on success, otherwise negative error code

Tags

[task-unrestricted](#)

## 8.13.2.3 RTSER\_RTIOC\_GET\_CONTROL

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

Get serial device's modem control register

Parameters

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

Returns

0 on success, otherwise negative error code

Tags

[task-unrestricted](#)

## 8.13.2.4 RTSER\_RTIOC\_GET\_STATUS

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

Get serial device status

## Parameters

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

## Returns

0 on success, otherwise negative error code

## Tags

[task-unrestricted](#)

## 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.

## 8.13.2.5 RTSER\_RTIOC\_SET\_CONFIG

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

## Set serial device configuration

## Parameters

|    |     |                                                                        |
|----|-----|------------------------------------------------------------------------|
| in | arg | 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.

## Tags

[task-unrestricted](#)

## 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.

## Examples:

[cross-link.c](#).



## 8.13.2.6 RTSER\_RTIOC\_SET\_CONTROL

```
#define RTSER_RTIOC_SET_CONTROL _IOW(RTIOC_TYPE_SERIAL, 0x04, int)
```

Set serial device's modem control register

Parameters

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

Returns

0 on success, otherwise negative error code

Tags

[task-unrestricted](#)

## 8.13.2.7 RTSER\_RTIOC\_WAIT\_EVENT

```
#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 | arg | 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.

Tags

[mode-unrestricted](#)

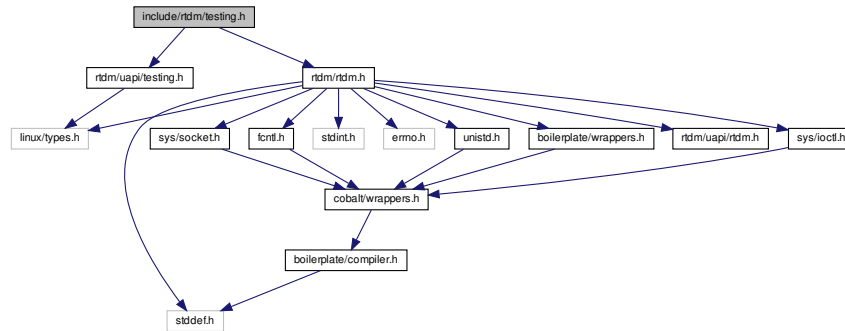
Examples:

[cross-link.c](#).


## 8.14 include/rtdm/testing.h File Reference

Real-Time Driver Model for Xenomai, testing device profile header.

Include dependency graph for testing.h:



This graph shows which files directly or indirectly include this file:



### 8.14.1 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.

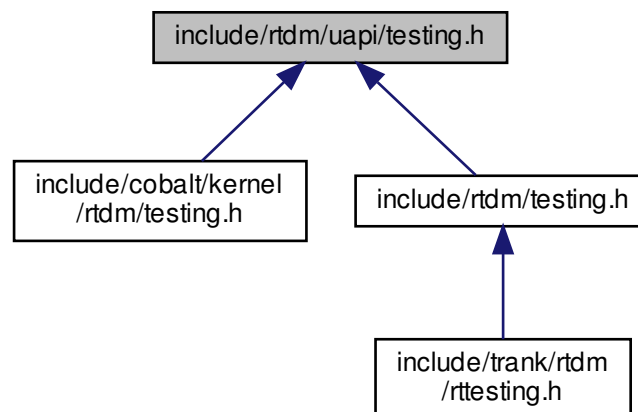
## 8.15 include/rtdm/uapi/testing.h File Reference

Real-Time Driver Model for Xenomai, testing device profile header.

Include dependency graph for testing.h:



This graph shows which files directly or indirectly include this file:



### 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\\_RTDMTEST](#) 3

- ```

      subclass name: "rtm"
    • #define RTDM_SUBCLASS_HEAPCHECK 4
      subclass name: "heapcheck"

```

## 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\_SWTEST\_SET\_TASKS\_COUNT** \_IOW(RTIOC\_TYPE\_TESTING, 0x30, \_\_u32)
- #define **RTTST\_RTIOC\_SWTEST\_SET\_CPU** \_IOW(RTIOC\_TYPE\_TESTING, 0x31, \_\_u32)
- #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, \_\_u32)
- #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, \_\_u32)
- #define **RTTST\_RTIOC\_RTDM\_DEFER\_CLOSE** \_IOW(RTIOC\_TYPE\_TESTING, 0x40, \_\_u32)
- #define **RTTST\_RTIOC\_RTDM\_ACTOR\_GET\_CPU** \_IOR(RTIOC\_TYPE\_TESTING, 0x41, \_\_u32)
- #define **RTTST\_RTIOC\_RTDM\_PING\_PRIMARY** \_IOR(RTIOC\_TYPE\_TESTING, 0x42, \_\_u32)
- #define **RTTST\_RTIOC\_RTDM\_PING\_SECONDARY** \_IOR(RTIOC\_TYPE\_TESTING, 0x43, \_\_u32)
- #define **RTTST\_RTIOC\_HEAP\_CHECK** \_IOR(RTIOC\_TYPE\_TESTING, 0x44, struct rttst\_heap\_parms)
- #define **RTTST\_RTIOC\_HEAP\_STAT\_COLLECT** \_IOR(RTIOC\_TYPE\_TESTING, 0x45, int)

### 8.15.1 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)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

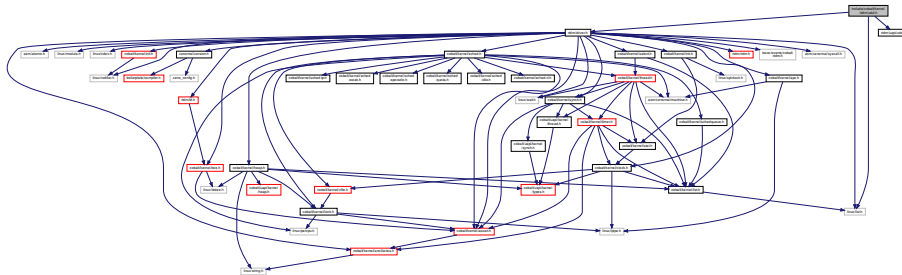
This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

## 8.16 include/cobalt/kernel/rtdm/udd.h File Reference

Copyright (C) 2014 Philippe Gerum [rpm@xenomai.org](mailto:rpm@xenomai.org)

Include dependency graph for udd.h:



### Data Structures

- struct [udd\\_memregion](#)
- struct [udd\\_device](#)
- struct [udd\\_device::udd\\_reserved](#)

*Reserved to the UDD core.*

### Macros

- #define [UDD\\_IRQ\\_NONE](#) 0  
*No IRQ managed.*
- #define [UDD\\_IRQ\\_CUSTOM](#) (-1)  
*IRQ directly managed from the mini-driver on top of the UDD core.*

#### Memory types for mapping

*Types of memory for mapping*

*The UDD core implements a default `->mmap()` handler which first attempts to hand over the request to the corresponding handler defined by the mini-driver. If not present, the UDD core establishes the mapping automatically, depending on the memory type defined for the region.*

- #define [UDD\\_MEM\\_NONE](#) 0  
*No memory region.*
- #define [UDD\\_MEM\\_PHYS](#) 1  
*Physical I/O memory region.*
- #define [UDD\\_MEM\\_LOGICAL](#) 2  
*Kernel logical memory region (e.g.*
- #define [UDD\\_MEM\\_VIRTUAL](#) 3  
*Virtual memory region with no direct physical mapping (e.g.*

## Functions

- int `udd_register_device` (struct `udd_device` \*udd)

*Register a UDD device*

- int `udd_unregister_device` (struct `udd_device` \*udd)

*Unregister a UDD device*

- struct `udd_device` \* `udd_get_device` (struct `rtm_fd` \*fd)

*RTDM file descriptor to target UDD device*

- void `udd_notify_event` (struct `udd_device` \*udd)

*Notify an IRQ event for an unmanaged interrupt*

- void `udd_enable_irq` (struct `udd_device` \*udd, `rtm_event_t` \*done)

*Enable the device IRQ line*

- void `udd_disable_irq` (struct `udd_device` \*udd, `rtm_event_t` \*done)

*Disable the device IRQ line*

### 8.16.1 Detailed Description

Copyright (C) 2014 Philippe Gerum [rpm@xenomai.org](mailto:rpm@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.


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.

## 8.17 include/rtdm/uapi/udd.h File Reference

This file is part of the Xenomai project.

This graph shows which files directly or indirectly include this file:



## Data Structures

- struct [udd\\_notify](#)  
*UDD event notification descriptor.*

## Macros

**UDD\_IOCTL***IOCTL requests*

- #define [UDD\\_RTIOC\\_IRQEN](#) \_IO(RTDM\_CLASS\_UDD, 0)  
*Enable the interrupt line.*
- #define [UDD\\_RTIOC\\_IRQDIS](#) \_IO(RTDM\_CLASS\_UDD, 1)  
*Disable the interrupt line.*
- #define [UDD\\_RTIOC\\_IRQSIG](#) \_IOW(RTDM\_CLASS\_UDD, 2, struct [udd\\_notify](#))  
*Enable/Disable signal notification upon interrupt event.*

## 8.17.1 Detailed Description

This file is part of the Xenomai project.

## Author

Copyright (C) 2014 Philippe Gerum [rpm@xenomai.org](mailto:rpm@xenomai.org)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.


This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.


## 8.18 include/rtdm/analogy.h File Reference

Analogy for Linux, library facilities.

Include dependency graph for analogy.h:



This graph shows which files directly or indirectly include this file:



## Data Structures

- struct [a4l\\_descriptor](#)  
Structure containing device-information useful to users.

## Macros

### **ANALOGY\_xxx\_DESC**

Constants used as argument so as to define the description depth to recover

- #define [A4L\\_BSC\\_DESC](#) 0x0  
*BSC stands for basic descriptor (device data)*
- #define [A4L\\_CPLX\\_DESC](#) 0x1  
*CPLX stands for complex descriptor (subdevice + channel + range data)*

## 8.18.1 Detailed Description

Analogy for Linux, library facilities.

### Note

Copyright (C) 1997-2000 David A. Schleef [ds@schleef.org](mailto:ds@schleef.org)  
Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This library 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 Lesser General Public License for more details.

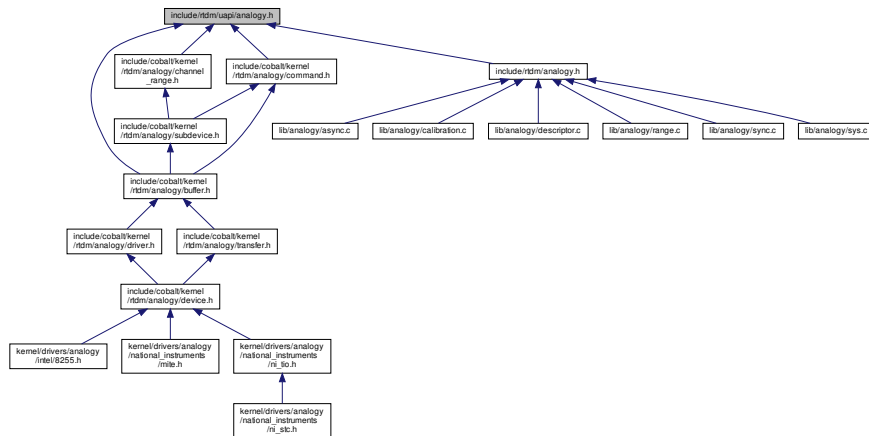
You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.



## 8.19 include/rtdm/uapi/analogy.h File Reference

Analogy for Linux, UAPI bits.

This graph shows which files directly or indirectly include this file:



## Data Structures

- struct [a4l\\_cmd\\_desc](#)  
Structure describing the asynchronous instruction.
- struct [a4l\\_instruction](#)  
Structure describing the synchronous instruction.
- struct [a4l\\_instruction\\_list](#)  
Structure describing the list of synchronous instructions.

## Macros

- #define [A4L\\_RNG\\_FACTOR](#) 1000000  
Constant for internal use only (must not be used by driver developer).
- #define [A4L\\_RNG\\_VOLT\\_UNIT](#) 0x0  
Volt unit range flag.
- #define [A4L\\_RNG\\_MAMP\\_UNIT](#) 0x1  
MilliAmpere unit range flag.
- #define [A4L\\_RNG\\_NO\\_UNIT](#) 0x2  
No unit range flag.
- #define [A4L\\_RNG\\_EXT\\_UNIT](#) 0x4  
External unit range flag.
- #define [A4L\\_RNG\\_UNIT\(x\)](#)  
Macro to retrieve the range unit from the range flags.
- #define [A4L\\_INSN\\_WAIT\\_MAX](#) 100000  
Maximal wait duration.

**ANALOGY\_CMD\_xxx**

Common command flags definitions

- #define `A4L_CMD_SIMUL` 0x1  
*Do not execute the command, just check it.*
- #define `A4L_CMD_BULK` 0x2  
*Perform data recovery / transmission in bulk mode.*
- #define `A4L_CMD_WRITE` 0x4  
*Perform a command which will write data to the device.*

## TRIG\_xxx

Command triggers flags definitions

- #define `TRIG_NONE` 0x00000001  
*Never trigger.*
- #define `TRIG_NOW` 0x00000002  
*Trigger now + N ns.*
- #define `TRIG_FOLLOW` 0x00000004  
*Trigger on next lower level trig.*
- #define `TRIG_TIME` 0x00000008  
*Trigger at time N ns.*
- #define `TRIG_TIMER` 0x00000010  
*Trigger at rate N ns.*
- #define `TRIG_COUNT` 0x00000020  
*Trigger when count reaches N.*
- #define `TRIG_EXT` 0x00000040  
*Trigger on external signal N.*
- #define `TRIG_INT` 0x00000080  
*Trigger on analogy-internal signal N.*
- #define `TRIG_OTHER` 0x00000100  
*Driver defined trigger.*
- #define `TRIG_WAKE_EOS` 0x0020  
*Wake up on end-of-scan.*
- #define `TRIG_ROUND_MASK` 0x00030000  
*Trigger not implemented yet.*
- #define `TRIG_ROUND_NEAREST` 0x00000000  
*Trigger not implemented yet.*
- #define `TRIG_ROUND_DOWN` 0x00010000  
*Trigger not implemented yet.*
- #define `TRIG_ROUND_UP` 0x00020000  
*Trigger not implemented yet.*
- #define `TRIG_ROUND_UP_NEXT` 0x00030000  
*Trigger not implemented yet.*

## Channel macros

Specific precompilation macros and constants useful for the channels descriptors tab located in the command structure

- #define `CHAN(a)` ((a) & 0xffff)  
*Channel indication macro.*
- #define `RNG(a)` (((a) & 0xff) << 16)  
*Range definition macro.*
- #define `AREF(a)` (((a) & 0x03) << 24)  
*Reference definition macro.*
- #define `FLAGS(a)` ((a) & CR\_FLAGS\_MASK)  
*Flags definition macro.*
- #define `PACK(a, b, c)` (a | `RNG(b)` | `AREF(c)`)  
*Channel + range + reference definition macro.*
- #define `PACK_FLAGS(a, b, c, d)` (`PACK(a, b, c)` | `FLAGS(d)`)  
*Channel + range + reference + flags definition macro.*

- #define `AREF_GROUND` 0x00  
*Analog reference is analog ground.*
- #define `AREF_COMMON` 0x01  
*Analog reference is analog common.*
- #define `AREF_DIFF` 0x02  
*Analog reference is differential.*
- #define `AREF_OTHER` 0x03  
*Analog reference is undefined.*

### Subdevices types

Flags to define the subdevice type

- #define `A4L_SUBD_UNUSED` (`A4L_SUBD_MASK_SPECIAL|0x1`)  
*Unused subdevice.*
- #define `A4L_SUBD_AI` (`A4L_SUBD_MASK_READ|0x2`)  
*Analog input subdevice.*
- #define `A4L_SUBD_AO` (`A4L_SUBD_MASK_WRITE|0x4`)  
*Analog output subdevice.*
- #define `A4L_SUBD_DI` (`A4L_SUBD_MASK_READ|0x8`)  
*Digital input subdevice.*
- #define `A4L_SUBD_DO` (`A4L_SUBD_MASK_WRITE|0x10`)  
*Digital output subdevice.*
- #define `A4L_SUBD_DIO` (`A4L_SUBD_MASK_SPECIAL|0x20`)  
*Digital input/output subdevice.*
- #define `A4L_SUBD_COUNTER` (`A4L_SUBD_MASK_SPECIAL|0x40`)  
*Counter subdevice.*
- #define `A4L_SUBD_TIMER` (`A4L_SUBD_MASK_SPECIAL|0x80`)  
*Timer subdevice.*
- #define `A4L_SUBD_MEMORY` (`A4L_SUBD_MASK_SPECIAL|0x100`)  
*Memory, EEPROM, DPRAM.*
- #define `A4L_SUBD_CALIB` (`A4L_SUBD_MASK_SPECIAL|0x200`)  
*Calibration subdevice DACs.*
- #define `A4L_SUBD_PROC` (`A4L_SUBD_MASK_SPECIAL|0x400`)  
*Processor, DSP.*
- #define `A4L_SUBD_SERIAL` (`A4L_SUBD_MASK_SPECIAL|0x800`)  
*Serial IO subdevice.*
- #define `A4L_SUBD_TYPES`  
*Mask which gathers all the types.*

### Subdevice features

Flags to define the subdevice's capabilities

- #define `A4L_SUBD_CMD` 0x1000  
*The subdevice can handle command (i.e it can perform asynchronous acquisition)*
- #define `A4L_SUBD_MMAP` 0x8000  
*The subdevice support mmap operations (technically, any driver can do it; however, the developer might want that his driver must be accessed through read / write.*

### Subdevice status

Flags to define the subdevice's status

- #define `A4L_SUBD_BUSY_NR` 0  
*The subdevice is busy, a synchronous or an asynchronous acquisition is occurring.*
- #define `A4L_SUBD_BUSY` (`1 << A4L_SUBD_BUSY_NR`)  
*The subdevice is busy, a synchronous or an asynchronous acquisition is occurring.*
- #define `A4L_SUBD_CLEAN_NR` 1

The subdevice is about to be cleaned in the middle of the detach procedure.

- #define **A4L\_SUBD\_CLEAN** (1 << **A4L\_SUBD\_CLEAN\_NR**)

The subdevice is busy, a synchronous or an asynchronous acquisition is occurring.

## Instruction type

Flags to define the type of instruction

- #define **A4L\_INSN\_READ** (0 | **A4L\_INSN\_MASK\_READ**)  
Read instruction.
- #define **A4L\_INSN\_WRITE** (1 | **A4L\_INSN\_MASK\_WRITE**)  
Write instruction.
- #define **A4L\_INSN\_BITS**  
"Bits" instruction
- #define **A4L\_INSN\_CONFIG**  
Configuration instruction.
- #define **A4L\_INSN\_GTOD**  
Get time instruction.
- #define **A4L\_INSN\_WAIT**  
Wait instruction.
- #define **A4L\_INSN\_INTTRIG**  
Trigger instruction (to start asynchronous acquisition)

## Configuration instruction type

Values to define the type of configuration instruction

- #define **A4L\_INSN\_CONFIG\_DIO\_INPUT** 0
- #define **A4L\_INSN\_CONFIG\_DIO\_OUTPUT** 1
- #define **A4L\_INSN\_CONFIG\_DIO\_OPENDRAIN** 2
- #define **A4L\_INSN\_CONFIG\_ANALOG\_TRIG** 16
- #define **A4L\_INSN\_CONFIG\_ALT\_SOURCE** 20
- #define **A4L\_INSN\_CONFIG\_DIGITAL\_TRIG** 21
- #define **A4L\_INSN\_CONFIG\_BLOCK\_SIZE** 22
- #define **A4L\_INSN\_CONFIG\_TIMER\_1** 23
- #define **A4L\_INSN\_CONFIG\_FILTER** 24
- #define **A4L\_INSN\_CONFIG\_CHANGE\_NOTIFY** 25
- #define **A4L\_INSN\_CONFIG\_SERIAL\_CLOCK** 26
- #define **A4L\_INSN\_CONFIG\_BIDIRECTIONAL\_DATA** 27
- #define **A4L\_INSN\_CONFIG\_DIO\_QUERY** 28
- #define **A4L\_INSN\_CONFIG\_PWM\_OUTPUT** 29
- #define **A4L\_INSN\_CONFIG\_GET\_PWM\_OUTPUT** 30
- #define **A4L\_INSN\_CONFIG\_ARM** 31
- #define **A4L\_INSN\_CONFIG\_DISARM** 32
- #define **A4L\_INSN\_CONFIG\_GET\_COUNTER\_STATUS** 33
- #define **A4L\_INSN\_CONFIG\_RESET** 34
- #define **A4L\_INSN\_CONFIG\_GPCT\_SINGLE\_PULSE\_GENERATOR** 1001 /\* Use CTR as single pulsegenerator \*/
- #define **A4L\_INSN\_CONFIG\_GPCT\_PULSE\_TRAIN\_GENERATOR** 1002 /\* Use CTR as pulsetraingenerator \*/
- #define **A4L\_INSN\_CONFIG\_GPCT\_QUADRATURE\_ENCODER** 1003 /\* Use the counter as encoder \*/
- #define **A4L\_INSN\_CONFIG\_SET\_GATE\_SRC** 2001 /\* Set gate source \*/
- #define **A4L\_INSN\_CONFIG\_GET\_GATE\_SRC** 2002 /\* Get gate source \*/
- #define **A4L\_INSN\_CONFIG\_SET\_CLOCK\_SRC** 2003 /\* Set master clock source \*/
- #define **A4L\_INSN\_CONFIG\_GET\_CLOCK\_SRC** 2004 /\* Get master clock source \*/
- #define **A4L\_INSN\_CONFIG\_SET\_OTHER\_SRC** 2005 /\* Set other source \*/
- #define **A4L\_INSN\_CONFIG\_SET\_COUNTER\_MODE** 4097
- #define **A4L\_INSN\_CONFIG\_SET\_ROUTING** 4099
- #define **A4L\_INSN\_CONFIG\_GET\_ROUTING** 4109

### Counter status bits

Status bits for `INSN_CONFIG_GET_COUNTER_STATUS`

- #define **A4L\_COUNTER\_ARMED** 0x1
- #define **A4L\_COUNTER\_COUNTING** 0x2
- #define **A4L\_COUNTER\_TERMINAL\_COUNT** 0x4

### IO direction

Values to define the IO polarity

- #define **A4L\_INPUT** 0
- #define **A4L\_OUTPUT** 1
- #define **A4L\_OPENDRAIN** 2

### Events types

Values to define the Analogy events. They might used to send some specific events through the instruction interface.

- #define **A4L\_EV\_START** 0x00040000
- #define **A4L\_EV\_SCAN\_BEGIN** 0x00080000
- #define **A4L\_EV\_CONVERT** 0x00100000
- #define **A4L\_EV\_SCAN\_END** 0x00200000
- #define **A4L\_EV\_STOP** 0x00400000

## 8.19.1 Detailed Description

Analogy for Linux, UAPI bits.

Note

Copyright (C) 1997-2000 David A. Schleef [ds@schleef.org](mailto:ds@schleef.org)  
Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

## 8.19.2 Macro Definition Documentation

### 8.19.2.1 A4L\_RNG\_FACTOR


```
#define A4L_RNG_FACTOR 1000000
```

Constant for internal use only (must not be used by driver developer).

## 8.20 lib/analogy/calibration.h File Reference

Analogy for Linux, internal calibration declarations.

This graph shows which files directly or indirectly include this file:



### 8.20.1 Detailed Description

Analogy for Linux, internal calibration declarations.

Note

Copyright (C) 2014 Jorge A Ramirez-Ortiz [jro@xenomai.org](mailto:jro@xenomai.org)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

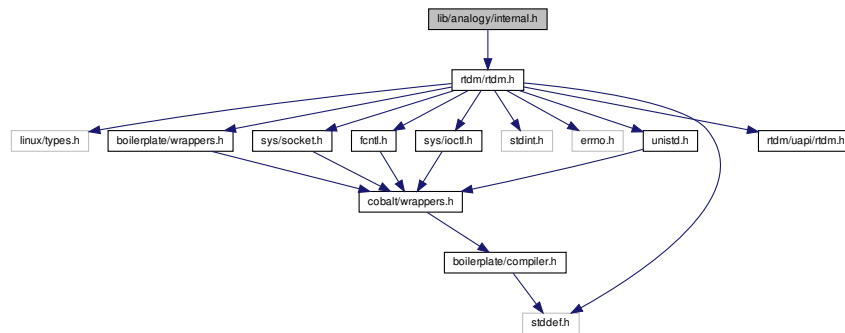
This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.


## 8.21 lib/analogy/internal.h File Reference

Analogy for Linux, internal declarations.

Include dependency graph for internal.h:



This graph shows which files directly or indirectly include this file:



### 8.21.1 Detailed Description

Analogy for Linux, internal declarations.

Note

Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.


This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

## 8.22 lib/analogy/async.c File Reference

Analogy for Linux, command, transfer, etc.

Include dependency graph for `async.c`:



### Functions

- `int a4l_snd_command (a4l_desc_t *dsc, a4l_cmd_t *cmd)`  
*Send a command to an Analogy device.*
- `int a4l_snd_cancel (a4l_desc_t *dsc, unsigned int idx_subd)`  
*Cancel an asynchronous acquisition.*
- `int a4l_set_bufsize (a4l_desc_t *dsc, unsigned int idx_subd, unsigned long size)`  
*Change the size of the asynchronous buffer.*
- `int a4l_get_bufsize (a4l_desc_t *dsc, unsigned int idx_subd, unsigned long *size)`  
*Get the size of the asynchronous buffer.*
- `int a4l_mark_bufrw (a4l_desc_t *dsc, unsigned int idx_subd, unsigned long cur, unsigned long *new)`  
*Update the asynchronous buffer state.*
- `int a4l_poll (a4l_desc_t *dsc, unsigned int idx_subd, unsigned long ms_timeout)`  
*Get the available data count.*
- `int a4l_mmap (a4l_desc_t *dsc, unsigned int idx_subd, unsigned long size, void **ptr)`  
*Map the asynchronous ring-buffer into a user-space.*
- `int a4l_async_read (a4l_desc_t *dsc, void *buf, size_t nbyte, unsigned long ms_timeout)`  
*Perform asynchronous read operation on the analog input subdevice.*
- `int a4l_async_write (a4l_desc_t *dsc, void *buf, size_t nbyte, unsigned long ms_timeout)`  
*Perform asynchronous write operation on the analog input subdevice.*

### 8.22.1 Detailed Description

Analogy for Linux, command, transfer, etc.

related features







## 8.24.1 Detailed Description

Analogy for Linux, descriptor related features.

Note

Copyright (C) 1997-2000 David A. Schlee [ds@schleef.org](mailto:ds@schleef.org)  
 Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.


This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

## 8.25 lib/analogy/range.c File Reference

Analogy for Linux, range related features.

Include dependency graph for range.c:



## Functions

- int [a4l\\_sizeof\\_chan](#) (a4l\_chinfo\_t \*chan)  
 Get the size in memory of an acquired element.
- int [a4l\\_sizeof\\_subd](#) (a4l\_sbinfo\_t \*subd)  
 Get the size in memory of a digital acquired element.
- int [a4l\\_find\\_range](#) (a4l\_desc\_t \*dsc, unsigned int idx\_subd, unsigned int idx\_chan, unsigned long unit, double min, double max, a4l\_rnginfo\_t \*\*rng)  
 Find the most suitable range.
- int [a4l\\_rawtoul](#) (a4l\_chinfo\_t \*chan, unsigned long \*dst, void \*src, int cnt)  
 Unpack raw data (from the driver) into unsigned long values.
- int [a4l\\_rawtof](#) (a4l\_chinfo\_t \*chan, a4l\_rnginfo\_t \*rng, float \*dst, void \*src, int cnt)  
 Convert raw data (from the driver) to float-typed samples.

- int [a4l\\_rawtod](#) (a4l\_chinfo\_t \*chan, a4l\_rnginfo\_t \*rng, double \*dst, void \*src, int cnt)  
*Convert raw data (from the driver) to double-typed samples.*
- int [a4l\\_ultoraw](#) (a4l\_chinfo\_t \*chan, void \*dst, unsigned long \*src, int cnt)  
*Pack unsigned long values into raw data (for the driver)*
- int [a4l\\_ftoraw](#) (a4l\_chinfo\_t \*chan, a4l\_rnginfo\_t \*rng, void \*dst, float \*src, int cnt)  
*Convert float-typed samples to raw data (for the driver)*
- int [a4l\\_dtoraw](#) (a4l\_chinfo\_t \*chan, a4l\_rnginfo\_t \*rng, void \*dst, double \*src, int cnt)  
*Convert double-typed samples to raw data (for the driver)*

### 8.25.1 Detailed Description

Analogy for Linux, range related features.

Note

Copyright (C) 1997-2000 David A. Schleef [ds@schleef.org](mailto:ds@schleef.org)

Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.


This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

## 8.26 lib/analogy/root\_leaf.h File Reference

Analogy for Linux, root / leaf system.

This graph shows which files directly or indirectly include this file:



### 8.26.1 Detailed Description

Analogy for Linux, root / leaf system.

Note

Copyright (C) 1997-2000 David A. Schleef [ds@schleef.org](mailto:ds@schleef.org)  
 Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.


This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

## 8.27 lib/analogy/sync.c File Reference

Analogy for Linux, instruction related features.

Include dependency graph for sync.c:



### Functions

- int [a4l\\_snd\\_insnlist](#) ([a4l\\_desc\\_t](#) \*dsc, [a4l\\_insnlst\\_t](#) \*arg)  
*Perform a list of synchronous acquisition misc operations.*
- int [a4l\\_snd\\_insn](#) ([a4l\\_desc\\_t](#) \*dsc, [a4l\\_insn\\_t](#) \*arg)  
*Perform a synchronous acquisition misc operation.*
- int [a4l\\_sync\\_write](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned int chan\_desc, unsigned int ns\_delay, void \*buf, [size\\_t](#) nbyte)  
*Perform a synchronous acquisition write operation.*
- int [a4l\\_sync\\_read](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned int chan\_desc, unsigned int ns\_delay, void \*buf, [size\\_t](#) nbyte)  
*Perform a synchronous acquisition read operation.*
- int [a4l\\_sync\\_dio](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, void \*mask, void \*buf)  
*Perform a synchronous acquisition digital acquisition.*
- int [a4l\\_config\\_subd](#) ([a4l\\_desc\\_t](#) \*dsc, unsigned int idx\_subd, unsigned int type,...)  
*Configure a subdevice.*

### 8.27.1 Detailed Description

Analogy for Linux, instruction related features.

Note

Copyright (C) 1997-2000 David A. Schleef [ds@schleef.org](mailto:ds@schleef.org)  
 Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.


This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

## 8.28 lib/analogy/sys.c File Reference

Analogy for Linux, descriptor related features.

Include dependency graph for sys.c:



### Functions

- int [a4l\\_sys\\_open](#) (const char \*fname)  
*Open an Analogy device.*
- int [a4l\\_sys\\_close](#) (int fd)  
*Close an Analogy device.*
- int [a4l\\_sys\\_read](#) (int fd, void \*buf, size\_t nbyte)  
*Read from an Analogy device.*
- int [a4l\\_sys\\_write](#) (int fd, void \*buf, size\_t nbyte)  
*Write to an Analogy device.*
- int [a4l\\_sys\\_attach](#) (int fd, a4l\_Inkdesc\_t \*arg)  
*Attach an Analogy device to a driver.*
- int [a4l\\_sys\\_detach](#) (int fd)  
*Detach an Analogy device from a driver.*
- int [a4l\\_sys\\_bufcfg](#) (int fd, unsigned int idx\_subd, unsigned long size)  
*Configure the buffer size.*

### 8.28.1 Detailed Description

Analogy for Linux, descriptor related features.

Note

Copyright (C) 1997-2000 David A. Schleef [ds@schleef.org](mailto:ds@schleef.org)

Copyright (C) 2008 Alexis Berlemont [alexis.berlemont@free.fr](mailto:alexis.berlemont@free.fr)

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This library 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 Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.





# Chapter 9

## Example Documentation

### 9.1 bufp-label.c

```
/*
 * Copyright (C) 2009 Philippe Gerum <rpm@xenomai.org>.
 *
 * This library is free software; you can redistribute it and/or
 * modify it under the terms of the GNU Lesser General Public
 * License as published by the Free Software Foundation; either
 * version 2 of the License, or (at your option) any later version.
 *
 * This library 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
 * Lesser General Public License for more details.
 *
 * You should have received a copy of the GNU Lesser General Public
 * License along with this library; if not, write to the Free Software
 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.
 *
 *
 * 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).
 */
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <pthread.h>
#include <errno.h>
#include <rtm/ipc.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, IPPROTO_BUF);
    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_BUF, BUF_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, BUF_PORT_LABEL);
    ret = setsockopt(s, SOL_BUF, BUF_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/bufp directory once the socket
     * is bound.
     *
     * saddr.sipc_port specifies the port number to use. If -1 is
     * passed, the BUF 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");
        }
        printf("%s: received %d bytes, \"%s\"\n",
               __FUNCTION__, ret, buf);
    }

    return NULL;
}

static void *client(void *arg)
{
    struct rtipc_port_label plabel;
    struct sockaddr_ipc svaddr;
    int ret, s, n = 0, len;
    struct timespec ts;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPPROTO_BUF);
    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
     * BUF does not try to register this label for the client
     * port as well (like the server thread did).
     */

```

```

strcpy(plabel.label, BUFP_PORT_LABEL);
ret = setsockopt(s, SOL_BUFPP, BUFP_LABEL,
                &plabel, sizeof(plabel));
if (ret)
    fail("setsockopt");

memset(&svsaddr, 0, sizeof(svsaddr));
svsaddr.sipc_family = AF_RTIPC;
svsaddr.sipc_port = -1; /* Tell BUFP to search by label. */
ret = connect(s, (struct sockaddr *)&svsaddr, sizeof(svsaddr));
if (ret)
    fail("connect");

for (;;) {
    len = strlen(msg[n]);
    ret = write(s, msg[n], len);
    if (ret < 0) {
        close(s);
        fail("write");
    }
    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;
}

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 set;
    int sig;

    sigemptyset(&set);
    sigaddset(&set, SIGINT);
    sigaddset(&set, SIGTERM);
    sigaddset(&set, SIGHUP);
    pthread_sigmask(SIG_BLOCK, &set, NULL);

    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");

    sigwait(&set, &sig);
    pthread_cancel(svtid);
    pthread_cancel(cltid);
    pthread_join(svtid, NULL);
    pthread_join(cltid, NULL);

    return 0;
}

```

## 9.2 bufp-readwrite.c

```
/*
```

```

* Copyright (C) 2009 Philippe Gerum <rpm@xenomai.org>.
*
* This library is free software; you can redistribute it and/or
* modify it under the terms of the GNU Lesser General Public
* License as published by the Free Software Foundation; either
* version 2 of the License, or (at your option) any later version.
*
* This library 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
* Lesser General Public License for more details.
*
* You should have received a copy of the GNU Lesser General Public
* License along with this library; if not, write to the Free Software
* Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.
*
*
* 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).
*/
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <pthread.h>
#include <errno.h>
#include <rtdm/ipc.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, IPPROTO_BUF);
    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_BUF, 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");
        }
        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");
        }
        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;
}

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 set;
    int sig;

    sigemptyset(&set);
    sigaddset(&set, SIGINT);
    sigaddset(&set, SIGTERM);
    sigaddset(&set, SIGHUP);
    pthread_sigmask(SIG_BLOCK, &set, NULL);

    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");

        sigwait(&set, &sig);
        pthread_cancel(svtid);
        pthread_cancel(cltid);
        pthread_join(svtid, NULL);
        pthread_join(cltid, NULL);

    return 0;
}

```

## 9.3 can-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 <memory.h>
#include <netinet/in.h>
#include <net/if.h>
#include <sys/ioctl.h>
#include <rtdm/can.h>
#include <xenomai/init.h>

#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;
};

void application_usage(void)
{
    fprintf(stderr, "usage: %s [options] <tx-can-interface> <rx-can-interface>:\n",
        get_program_name());
    fprintf(stderr,
        " -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");
}

static 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, t;

    /* Pre-fill CAN frame */
    frame.can_id = can_id;
    frame.can_dlc = sizeof(*rtt_time);

    pthread_setname_np(pthread_self(), "rtcan_rtt_transmitter");
    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);
        t = (long long)time.tv_sec * NSEC_PER_SEC + time.tv_nsec;
        memcpy(rtt_time, &t, sizeof(t));

        /* 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++;
    }
}

static 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, t;
    struct rtt_stat rtt_stat = {0, 100000000000000000LL, -100000000000000000LL,
        0, 0, 0};

    pthread_setname_np(pthread_self(), "rtcan_rtt_receiver");
    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 {
        memcpy(&t, rtt_time, sizeof(t));
        clock_gettime(CLOCK_MONOTONIC, &time);
        if (rxcount > 0) {
            rtt_stat.rtt = ((long long)time.tv_sec * 1000000000LL +
                time.tv_nsec - t);
            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;
        }
    }
}

static void catch_signal(int sig)
{
    mq_close(mq);
    close(rxsock);
    close(txsock);
}

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 can_ifreq ifr;
    int ret, opt;

    struct option long_options[] = {
        { "id", required_argument, 0, 'i' },
        { "cycle", required_argument, 0, 'c' },
        { "repeater", no_argument, 0, 'r' },
        { 0, 0, 0, 0 },
    };

    while ((opt = getopt_long(argc, argv, "ri: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);
                exit(-1);
        }
    }

    printf("%d %d\n", optind, argc);
    if (optind + 2 != argc) {
        xenomai_usage();
        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;
}

namecpy(ifr.ifr_name, rxdev);
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 = 0x3ff;
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;
    }

    namecpy(ifr.ifr_name, txdev);
    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);

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);
ret = pthread_create(&rxthread, &thattr, &receiver, NULL);
if (ret) {
    fprintf(stderr, "%s: pthread_create(receiver) failed\n",
            strerror(-ret));
    goto failure3;
}

if (!repeater) {
    /* Create transmitter 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 %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! */
close(rxsock);
close(txsock);
pthread_join(txthread, NULL);
pthread_cancel(rxthread);
pthread_join(rxthread, NULL);

return 0;

failure4:
pthread_cancel(rxthread);
pthread_join(rxthread, NULL);
failure3:
mq_close(mq);
failure2:
close(txsock);
failure1:
close(rxsock);

return 1;
}

```

## 9.4 cross-link.c

```

/*
 * cross-link.c
 *
 * Userspace test program (Xenomai alchemy 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 <alchemy/task.h>
#include <alchemy/timer.h>
#include <rtdm/serial.h>

#define MAIN_PREFIX "main : "
#define WTASK_PREFIX "write_task: "
#define RTASK_PREFIX "read_task: "

#define WRITE_FILE "/dev/rtdm/rtser0"
#define READ_FILE "/dev/rtdm/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 = 10000000011lu;
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 = 1000000000, /* 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 = 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(errno));
        break;
    }
} while (err == -EAGAIN && i < 10);

return err;
}

static 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;
    }
}

static void catch_signal(int sig)
{
    cleanup_all();
    printf(MAIN_PREFIX "exit\n");
    return;
}

static void write_task_proc(void *arg)
{
    int err;
    RTIME write_time;
    ssize_t sz = sizeof(RTIME);
    int 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 = write(write_fd, &write_time, sz);
        if (written < 0) {
            printf(WTASK_PREFIX "error on write, %s\n",
                   strerror(errno));
            break;
        } else if (written != sz) {
            printf(WTASK_PREFIX "only %d / %zd 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");
}

static 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);
    int rd = 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 = ioctl(read_fd, RTSER_RTIOC_WAIT_EVENT, &rx_event);
        if (err) {
            printf(WTASK_PREFIX
                "error on RTSER_RTIOC_WAIT_EVENT, %s\n",
                strerror(errno));
            if (err == -ETIMEDOUT)
                continue;
            break;
        }

        irq_time = rx_event.rxpend_timestamp;
        rd = read(read_fd, &write_time, sz);
        if (rd == 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 (rd < 0) {
            printf(WTASK_PREFIX "error on read, code %s\n",
                strerror(errno));
            break;
        } else {
            printf(WTASK_PREFIX "only %d / %zd byte received \n",
                rd, sz);
            break;
        }
    }

    if ((read_state & STATE_FILE_OPENED) &&
        close_file(read_fd, READ_FILE " (read)") == 0)
        read_state &= ~STATE_FILE_OPENED;

    printf(WTASK_PREFIX "exit\n");
}

int main(int argc, char* argv[])
{
    int err = 0;

    signal(SIGTERM, catch_signal);
    signal(SIGINT, catch_signal);

    /* open rtser0 */
    write_fd = open(WRITE_FILE, 0);
    if (write_fd < 0) {
        printf(MAIN_PREFIX "can't open %s (write), %s\n", WRITE_FILE,
            strerror(errno));
        goto error;
    }
    write_state |= STATE_FILE_OPENED;
    printf(MAIN_PREFIX "write-file opened\n");

    /* writing write-config */
    err = ioctl(write_fd, RTSER_RTIOC_SET_CONFIG, &write_config);
    if (err) {
        printf(MAIN_PREFIX "error while RTSER_RTIOC_SET_CONFIG, %s\n",

```

```

        strerror(errno));
    goto error;
}
printf(MAIN_PREFIX "write-config written\n");

/* open rtser1 */
read_fd = open( READ_FILE, 0 );
if (read_fd < 0) {
    printf(MAIN_PREFIX "can't open %s (read), %s\n", READ_FILE,
        strerror(errno));
    goto error;
}
read_state |= STATE_FILE_OPENED;
printf(MAIN_PREFIX "read-file opened\n");

/* writing read-config */
err = ioctl(read_fd, RTSER_RTIOC_SET_CONFIG, &read_config);
if (err) {
    printf(MAIN_PREFIX "error while ioctl, %s\n",
        strerror(errno));
    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;
}

for (;;)
    pause();

return 0;

error:
cleanup_all();
return err;
}

```

## 9.5 iddp-label.c

```

/*
 * Copyright (C) 2009 Philippe Gerum <rpm@xenomai.org>.
 *
 * This library is free software; you can redistribute it and/or
 * modify it under the terms of the GNU Lesser General Public
 * License as published by the Free Software Foundation; either
 * version 2 of the License, or (at your option) any later version.

```

```

*
* This library 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
* Lesser General Public License for more details.
*
* You should have received a copy of the GNU Lesser General Public
* License along with this library; if not, write to the Free Software
* Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.
*
*
* 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.
*/
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <pthread.h>
#include <errno.h>
#include <rtdm/ipc.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");
    }
    printf("%s: received %d bytes, \"%s\" from port %d\n",
          __FUNCTION__, 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");
        }
        printf("%s: sent %d bytes, \"%s\"\n",
              __FUNCTION__, ret, msg[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;
}

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 set;
    int sig;

    sigemptyset(&set);
    sigaddset(&set, SIGINT);
    sigaddset(&set, SIGTERM);
    sigaddset(&set, SIGHUP);
    pthread_sigmask(SIG_BLOCK, &set, NULL);

    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(&svid, &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");

    sigwait(&set, &sig);
    pthread_cancel(svid);
    pthread_cancel(cltid);
    pthread_join(svid, NULL);
    pthread_join(cltid, NULL);

    return 0;
}

```

## 9.6 iddp-sendrecv.c

```

/*
 * Copyright (C) 2009 Philippe Gerum <rpm@xenomai.org>.
 *
 * This library is free software; you can redistribute it and/or
 * modify it under the terms of the GNU Lesser General Public
 * License as published by the Free Software Foundation; either
 * version 2 of the License, or (at your option) any later version.
 *
 * This library 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
 * Lesser General Public License for more details.
 *
 * You should have received a copy of the GNU Lesser General Public
 * License along with this library; if not, write to the Free Software
 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.
 *
 *
 * 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.
*/
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <pthread.h>
#include <errno.h>
#include <rtdm/ipc.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, IPPROTO_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");
        }
        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, IPPROTO_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");
        }
        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;
}

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 set;
    int sig;

    sigemptyset(&set);
    sigaddset(&set, SIGINT);
    sigaddset(&set, SIGTERM);
    sigaddset(&set, SIGHUP);
    pthread_sigmask(SIG_BLOCK, &set, NULL);

    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");

    sigwait(&set, &sig);
    pthread_cancel(svtid);
    pthread_cancel(cltid);
    pthread_join(svtid, NULL);
    pthread_join(cltid, NULL);

    return 0;
}

```

## 9.7 rtcانconfig.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 <boilerplate/ancillaries.h>

#include <rtm/can.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",
            prg);
}

static can_baudrate_t string_to_baudrate(char *str)
{
    can_baudrate_t baudrate;
    if (sscanf(str, "%i", &baudrate) != 1)
        return -1;
    return baudrate;
}

static 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;
}

static 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[IFNAMSIZ];
    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  can_ifreq ifr;
    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;
    }

    namecpy(ifname, argv[optind]);
    namecpy(ifr.ifr_name, ifname);

    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 = 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 = ioctl(can_fd, SIOCGIFINDEX, &ifr);
if (ret) {
    fprintf(stderr, "Can't get interface index for %s, errno = %d\n", ifname, errno);
    return ret;
}

if (new_baudrate != -1) {
    if (verbose)
        printf("baudrate: %d\n", new_baudrate);
    ifr.ifr_ifru.baudrate = new_baudrate;
    ret = ioctl(can_fd, SIOCSCANBAUDRATE, &ifr);
    if (ret) {
        goto abort;
    }
}

if (bittime_count) {
    bittime = &ifr.ifr_ifru.bittime;
    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 = ioctl(can_fd, SIOCSCANCUSTOMBITTIME, &ifr);
    if (ret) {
        goto abort;
    }
}

if (set_ctrlmode != 0) {
    ifr.ifr_ifru.ctrlmode = new_ctrlmode;
    if (verbose)
        printf("ctrlmode: %#x\n", new_ctrlmode);
    ret = ioctl(can_fd, SIOCSCANCTRLMODE, &ifr);
    if (ret) {
        goto abort;
    }
}

if (new_mode != -1) {
    ifr.ifr_ifru.mode = new_mode;
    ret = ioctl(can_fd, SIOCSCANMODE, &ifr);
    if (ret) {
        goto abort;
    }
}

close(can_fd);
return 0;

abort:
close(can_fd);
return ret;

```

```
}

```

## 9.8 rtcanrecv.c

```
#include <stdio.h>
#include <stdlib.h>
#include <signal.h>
#include <unistd.h>
#include <time.h>
#include <errno.h>
#include <getopt.h>

#include <alchemy/task.h>
#include <boilerplate/ancillaries.h>

#include <rtm/can.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",
            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;

static 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;
}

static void cleanup(void)
{
    int ret;

    if (verbose)
        printf("Cleaning up...\n");

    if (s >= 0) {
        ret = close(s);
        s = -1;
        if (ret) {
            fprintf(stderr, "close: %s\n", strerror(errno));
        }
        exit(EXIT_SUCCESS);
    }
}

static void cleanup_and_exit(int sig)
{
    if (verbose)
        printf("Signal %d received\n", sig);
    cleanup();
}

```

```

    exit(0);
}

static 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 = recvmsg(s, &msg, 0);
        } else
            ret = recvfrom(s, (void *)&frame, sizeof(can_frame_t), 0,
                (struct sockaddr *)&addr, &addrlen);

        if (ret < 0) {
            switch (errno) {
                case ETIMEDOUT:
                    if (verbose)
                        printf("recv: timed out\n");
                    continue;
                case EBADF:
                    if (verbose)
                        printf("recv: aborted because socket was closed\n");
                    break;
                default:
                    fprintf(stderr, "recv: %s\n", strerror(errno));
            }
            break;
        }

        if (print && (count % print) == 0) {
            printf("#d: (%d) ", count, addr.can_ifindex);
            if (with_timestamp && msg.msg_controllen) {
                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 can_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 },
};

signal(SIGTERM, cleanup_and_exit);
signal(SIGINT, cleanup_and_exit);

while ((opt = getopt_long(argc, argv, "hve:ft: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 = socket(PF_CAN, SOCK_RAW, CAN_RAW);
if (ret < 0) {
    fprintf(stderr, "socket: %s\n", strerror(errno));
    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]);
}

```

```

    namecpy(ifr.ifr_name, argv[optind]);
    if (verbose)
        printf("s=%d, ifr_name=%s\n", s, ifr.ifr_name);

    ret = ioctl(s, SIOCGIFINDEX, &ifr);
    if (ret < 0) {
        fprintf(stderr, "ioctl GET_IFINDEX: %s\n", strerror(errno));
        goto failure;
    }
}

if (err_mask) {
    ret = setsockopt(s, SOL_CAN_RAW, CAN_RAW_ERR_FILTER,
                    &err_mask, sizeof(err_mask));
    if (ret < 0) {
        fprintf(stderr, "setsockopt: %s\n", strerror(errno));
        goto failure;
    }
    if (verbose)
        printf("Using err_mask=%#x\n", err_mask);
}

if (filter_count) {
    ret = setsockopt(s, SOL_CAN_RAW, CAN_RAW_FILTER,
                    &rcv_filter, filter_count *
                    sizeof(struct can_filter));
    if (ret < 0) {
        fprintf(stderr, "setsockopt: %s\n", strerror(errno));
        goto failure;
    }
}

rcv_addr.can_family = AF_CAN;
rcv_addr.can_ifindex = ifr.ifr_ifindex;
ret = bind(s, (struct sockaddr *)&rcv_addr,
           sizeof(struct sockaddr_can));
if (ret < 0) {
    fprintf(stderr, "bind: %s\n", strerror(errno));
    goto failure;
}

if (timeout) {
    if (verbose)
        printf("Timeout: %lld ns\n", (long long)timeout);
    ret = ioctl(s, RTCAN_RTIOC_RCV_TIMEOUT, &timeout);
    if (ret) {
        fprintf(stderr, "ioctl RCV_TIMEOUT: %s\n", strerror(errno));
        goto failure;
    }
}

if (with_timestamp) {
    ret = ioctl(s, RTCAN_RTIOC_TAKE_TIMESTAMP,
               RTCAN_TAKE_TIMESTAMPS);
    if (ret) {
        fprintf(stderr, "ioctl TAKE_TIMESTAMP: %s\n", strerror(errno));
        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;
}

```

## 9.9 rtcansend.c

```

#include <stdio.h>
#include <stdlib.h>
#include <signal.h>
#include <unistd.h>

```

```

#include <time.h>
#include <errno.h>
#include <getopt.h>

#include <boilerplate/ancillaries.h>
#include <alchemy/task.h>
#include <alchemy/timer.h>

#include <rtdm/can.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;

static void cleanup(void)
{
    int ret;

    if (verbose)
        printf("Cleaning up...\n");

    usleep(100000);

    if (s >= 0) {
        ret = close(s);
        s = -1;
        if (ret) {
            fprintf(stderr, "close: %s\n", strerror(errno));
        }
        exit(EXIT_SUCCESS);
    }
}

static void cleanup_and_exit(int sig)
{
    if (verbose)
        printf("Signal %d received\n", sig);
    cleanup();
    exit(0);
}

static 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 definiton of a receive filter list */
        if (use_send)
            ret = send(s, (void *)&frame, sizeof(can_frame_t), 0);
        else
            ret = sendto(s, (void *)&frame, sizeof(can_frame_t), 0,
                        (struct sockaddr *)&to_addr, sizeof(to_addr));
        if (ret < 0) {
            switch (errno) {

```

```

        case ETIMEDOUT:
            if (verbose)
                printf("send(to): timed out\n");
            break;
        case EBADF:
            if (verbose)
                printf("send(to): aborted because socket was closed\n");
            break;
        default:
            fprintf(stderr, "send: %s\n", strerror(errno));
            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 can_ifreq ifr;
    char name[32];

    struct option long_options[] = {
        { "help", no_argument, 0, 'h' },
        { "identfier", 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 },
    };

    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 = socket(PF_CAN, SOCK_RAW, CAN_RAW);
if (ret < 0) {
    fprintf(stderr, "socket: %s\n", strerror(errno));
    return -1;
}
s = ret;

if (loopback >= 0) {
    ret = setsockopt(s, SOL_CAN_RAW, CAN_RAW_LOOPBACK,
                    &loopback, sizeof(loopback));

    if (ret < 0) {
        fprintf(stderr, "setsockopt: %s\n", strerror(errno));
        goto failure;
    }
    if (verbose)
        printf("Using loopback=%d\n", loopback);
}

strcpy(ifr.ifr_name, argv[optind]);
if (verbose)
    printf("s=%d, ifr_name=%s\n", s, ifr.ifr_name);

ret = ioctl(s, SIOCGIFINDEX, &ifr);
if (ret < 0) {
    fprintf(stderr, "ioctl: %s\n", strerror(errno));
    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 definition of a default receive filter list */
    ret = setsockopt(s, SOL_CAN_RAW, CAN_RAW_FILTER, NULL, 0);
    if (ret < 0) {
        fprintf(stderr, "setsockopt: %s\n", strerror(errno));
        goto failure;
    }

    ret = bind(s, (struct sockaddr *)&to_addr, sizeof(to_addr));
    if (ret < 0) {
        fprintf(stderr, "bind: %s\n", strerror(errno));
        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 = ioctl(s, RTCAN_RTIOC_SND_TIMEOUT, &timeout);
    if (ret) {
        fprintf(stderr, "ioctl SND_TIMEOUT: %s\n", strerror(errno));
        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;
}

```

## 9.10 xddp-echo.c

```

/*
 * Copyright (C) 2009 Philippe Gerum <rpm@xenomai.org>.
 *
 * This library is free software; you can redistribute it and/or
 * modify it under the terms of the GNU Lesser General Public
 * License as published by the Free Software Foundation; either
 * version 2 of the License, or (at your option) any later version.
 *
 * This library 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
 * Lesser General Public License for more details.
 *
 * You should have received a copy of the GNU Lesser General Public
 * License along with this library; if not, write to the Free Software
 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.
 *
 *
 * 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()               +--+
*/
#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 <rtdm/ipc.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 poolasz;
    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.
     */
    poolasz = 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");

    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");

    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;
}

int main(int argc, char **argv)
{
    struct sched_param rtparam = { .sched_priority = 42 };

```



```

pthread_attr_t rtattr, regattr;
sigset_t set;
int sig;

sigemptyset(&set);
sigaddset(&set, SIGINT);
sigaddset(&set, SIGTERM);
sigaddset(&set, SIGHUP);
pthread_sigmask(SIG_BLOCK, &set, NULL);

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");

sigwait(&set, &sig);
pthread_cancel(rt);
pthread_cancel(nrt);
pthread_join(rt, NULL);
pthread_join(nrt, NULL);

return 0;
}

```

## 9.11 xddp-label.c

```

/*
 * Copyright (C) 2009 Philippe Gerum <rpm@xenomai.org>.
 *
 * This library is free software; you can redistribute it and/or
 * modify it under the terms of the GNU Lesser General Public
 * License as published by the Free Software Foundation; either
 * version 2 of the License, or (at your option) any later version.
 *
 * This library 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
 * Lesser General Public License for more details.
 *
 * You should have received a copy of the GNU Lesser General Public
 * License along with this library; if not, write to the Free Software
 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.
 *
 *
 * 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/rtpc/xddp/xddp-demo |
* => read traffic from RT domain via read()   |
* => mirror traffic to RT domain via write()  +---+
*/
#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 <rtdm/ipc.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");

    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");

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");

    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 *regular_thread(void *arg)
{
    char buf[128], *devname;
    int fd, ret;

    if (asprintf(&devname,
        "/proc/xenomai/registry/rtpic/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;
}

int main(int argc, char **argv)
{
    struct sched_param rtparam = { .sched_priority = 42 };
    pthread_attr_t rtattr, regattr;
    sigset_t set;
    int sig;

    sigemptyset(&set);
    sigaddset(&set, SIGINT);
    sigaddset(&set, SIGTERM);
    sigaddset(&set, SIGHUP);
    pthread_sigmask(SIG_BLOCK, &set, NULL);

    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");

    sigwait(&set, &sig);
    pthread_cancel(rt1);
    pthread_cancel(rt2);
    pthread_cancel(nrt);
    pthread_join(rt1, NULL);
    pthread_join(rt2, NULL);
    pthread_join(nrt, NULL);

    return 0;
}

```

## 9.12 xddp-stream.c

```

/*
 * Copyright (C) 2009 Philippe Gerum <rpm@xenomai.org>.
 *
 * This library is free software; you can redistribute it and/or
 * modify it under the terms of the GNU Lesser General Public
 * License as published by the Free Software Foundation; either
 * version 2 of the License, or (at your option) any later version.
 *
 * This library 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
 * Lesser General Public License for more details.
 *
 * You should have received a copy of the GNU Lesser General Public
 * License along with this library; if not, write to the Free Software
 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.
 *
 *
 * 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() +---+
*/
#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 <rtm/ipc.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 streamsiz;
    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.
     */
    streamsiz = 1024; /* bytes */
    ret = setsockopt(s, SOL_XDDP, XDDP_BUFSZ,
                    &streamsiz, sizeof(streamsiz));
    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");
    }

    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");

    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;
}

int main(int argc, char **argv)
{
    struct sched_param rtparam = { .sched_priority = 42 };
    pthread_attr_t rtattr, regattr;
    sigset_t set;
    int sig;

    sigemptyset(&set);
    sigaddset(&set, SIGINT);
    sigaddset(&set, SIGTERM);
    sigaddset(&set, SIGHUP);

```

```
pthread_sigmask(SIG_BLOCK, &set, NULL);

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");

sigwait(&set, &sig);
pthread_cancel(rt);
pthread_cancel(nrt);
pthread_join(rt, NULL);
pthread_join(nrt, NULL);

return 0;
}
```



# Index

- `__attribute__`
  - Mutual exclusion, [352](#)
- `__xntimer_migrate`
  - Timer services, [279](#)
- `A4L_RNG_FACTOR`
  - `uapi/analogy.h`, [697](#)
- `a4l_add_subd`
  - Subdevice management services, [301](#)
- `a4l_alloc_subd`
  - Subdevice management services, [301](#)
- `a4l_async_read`
  - Asynchronous acquisition API, [410](#)
- `a4l_async_write`
  - Asynchronous acquisition API, [411](#)
- `a4l_buf_commit_absget`
  - Buffer management services, [304](#)
- `a4l_buf_commit_absput`
  - Buffer management services, [305](#)
- `a4l_buf_commit_get`
  - Buffer management services, [305](#)
- `a4l_buf_commit_put`
  - Buffer management services, [306](#)
- `a4l_buf_count`
  - Buffer management services, [306](#)
- `a4l_buf_evt`
  - Buffer management services, [308](#)
- `a4l_buf_get`
  - Buffer management services, [308](#)
- `a4l_buf_prepare_absget`
  - Buffer management services, [309](#)
- `a4l_buf_prepare_absput`
  - Buffer management services, [309](#)
- `a4l_buf_prepare_get`
  - Buffer management services, [310](#)
- `a4l_buf_prepare_put`
  - Buffer management services, [311](#)
- `a4l_buf_put`
  - Buffer management services, [311](#)
- `a4l_channel`, [579](#)
  - flags, [579](#)
  - nb\_bits, [579](#)
- `a4l_channels_desc`, [580](#)
  - chans, [580](#)
  - length, [580](#)
  - mode, [581](#)
- `a4l_close`
  - Descriptor API, [417](#)
- `a4l_cmd_desc`, [581](#)
  - data\_len, [582](#)
  - idx\_subd, [582](#)
- `a4l_config_subd`
  - Synchronous acquisition API, [436](#)
- `a4l_dcaltoraw`
  - Software calibration API, [412](#)
- `a4l_descriptor`, [582](#)
  - board\_name, [583](#)
  - driver\_name, [583](#)
  - fd, [583](#)
  - idx\_read\_subd, [583](#)
  - idx\_write\_subd, [583](#)
  - magic, [584](#)
  - nb\_subd, [584](#)
  - sbddata, [584](#)
  - sbsize, [584](#)
- `a4l_driver`, [584](#)
- `a4l_dtoraw`
  - Range / conversion API, [424](#)
- `a4l_fill_desc`
  - Descriptor API, [418](#)
- `a4l_find_range`
  - Range / conversion API, [425](#)
- `a4l_free_irq`
  - Interrupt management services, [313](#)
- `a4l_ftoraw`
  - Range / conversion API, [426](#)
- `a4l_get_bufsize`
  - Asynchronous acquisition API, [404](#)
- `a4l_get_chan`
  - Buffer management services, [312](#)
- `a4l_get_chinfo`
  - Descriptor API, [418](#)
- `a4l_get_cmd`
  - Buffer management services, [312](#)
- `a4l_get_irq`
  - Interrupt management services, [314](#)
- `a4l_get_rnginfo`
  - Descriptor API, [419](#)
- `a4l_get_softcal_converter`
  - Software calibration API, [413](#)
- `a4l_get_subd`
  - Subdevice management services, [302](#)
- `a4l_get_subdinfo`
  - Descriptor API, [419](#)
- `a4l_get_time`
  - Misc services, [316](#)
- `a4l_instruction`, [585](#)
  - idx\_subd, [586](#)
- `a4l_instruction_list`, [586](#)

- a4l\_mark\_bufrw
  - Asynchronous acquisition API, 404
- a4l\_math\_mean
  - Math API, 421
- a4l\_math\_polyfit
  - Math API, 422
- a4l\_math\_stddev
  - Math API, 422
- a4l\_math\_stddev\_of\_mean
  - Math API, 423
- a4l\_mmap
  - Asynchronous acquisition API, 405
- a4l\_open
  - Descriptor API, 420
- a4l\_poll
  - Asynchronous acquisition API, 406
- a4l\_range, 587
  - flags, 587
  - max, 587
  - min, 587
- a4l\_rawtod
  - Range / conversion API, 426
- a4l\_rawtodcal
  - Software calibration API, 413
- a4l\_rawtof
  - Range / conversion API, 427
- a4l\_rawtoul
  - Range / conversion API, 427
- a4l\_read\_calibration\_file
  - Software calibration API, 414
- a4l\_register\_drv
  - Driver management services, 297
- a4l\_request\_irq
  - Interrupt management services, 314
- a4l\_set\_bufsize
  - Asynchronous acquisition API, 407
- a4l\_sizeof\_chan
  - Range / conversion API, 428
- a4l\_sizeof\_subd
  - Range / conversion API, 428
- a4l\_snd\_cancel
  - Asynchronous acquisition API, 407
- a4l\_snd\_command
  - Asynchronous acquisition API, 408
- a4l\_snd\_insn
  - Synchronous acquisition API, 433
- a4l\_snd\_insnlist
  - Synchronous acquisition API, 433
- a4l\_subdevice, 588
- a4l\_sync\_dio
  - Synchronous acquisition API, 437
- a4l\_sync\_read
  - Synchronous acquisition API, 437
- a4l\_sync\_write
  - Synchronous acquisition API, 438
- a4l\_sys\_attach
  - Attach / detach Syscall API, 445
- a4l\_sys\_bufcfg
  - Attach / detach Syscall API, 446
- a4l\_sys\_close
  - Basic Syscall API, 442
- a4l\_sys\_desc
  - Descriptor Syscall API, 415
- a4l\_sys\_detach
  - Attach / detach Syscall API, 446
- a4l\_sys\_open
  - Basic Syscall API, 442
- a4l\_sys\_read
  - Basic Syscall API, 443
- a4l\_sys\_write
  - Basic Syscall API, 443
- a4l\_ultoraw
  - Range / conversion API, 429
- a4l\_unregister\_drv
  - Driver management services, 298
- addr
  - udd\_memregion, 624
- Alarm services, 448
  - rt\_alarm\_create, 449
  - rt\_alarm\_delete, 449
  - rt\_alarm\_inquire, 450
  - rt\_alarm\_start, 451
  - rt\_alarm\_stop, 451
- Alchemy API, 491
- Analogy framework, 295
- Analogy user API, 440
- Asynchronous acquisition API, 402, 410
  - a4l\_async\_read, 410
  - a4l\_async\_write, 411
  - a4l\_get\_bufsize, 404
  - a4l\_mark\_bufrw, 404
  - a4l\_mmap, 405
  - a4l\_poll, 406
  - a4l\_set\_bufsize, 407
  - a4l\_snd\_cancel, 407
  - a4l\_snd\_command, 408
- Asynchronous Procedure Calls, 98
  - xnipc\_alloc, 98
  - xnipc\_free, 99
  - xnipc\_schedule, 100
- atomic\_t, 589
- Attach / detach Syscall API, 445
  - a4l\_sys\_attach, 445
  - a4l\_sys\_bufcfg, 446
  - a4l\_sys\_detach, 446
- B\_PRIO
  - Buffer services, 454
- BUFP\_BUFSZ
  - Real-time IPC, 81
- BUFP\_LABEL
  - Real-time IPC, 82
- base\_minor
  - rtdm\_driver, 607
- Basic Syscall API, 442
  - a4l\_sys\_close, 442
  - a4l\_sys\_open, 442

- a4l\_sys\_read, [443](#)
- a4l\_sys\_write, [443](#)
- begin
  - xnvmfile\_regular\_ops, [631](#)
  - xnvmfile\_snapshot\_ops, [638](#)
- Big dual kernel lock, [27](#)
  - cobalt\_atomic\_enter, [27](#)
  - cobalt\_atomic\_leave, [28](#)
  - RTDM\_EXECUTE\_ATOMICALY, [28](#)
- bind\_\_AF\_RTIPC
  - Real-time IPC, [90](#)
- board\_name
  - a4l\_descriptor, [583](#)
- Buffer descriptor, [102](#)
  - xnbufd\_copy\_from\_kmem, [104](#)
  - xnbufd\_copy\_to\_kmem, [105](#)
  - xnbufd\_invalidate, [106](#)
  - xnbufd\_map\_kread, [107](#)
  - xnbufd\_map\_kwrite, [107](#)
  - xnbufd\_map\_uread, [108](#)
  - xnbufd\_map\_uwrite, [108](#)
  - xnbufd\_reset, [109](#)
  - xnbufd\_unmap\_kread, [109](#)
  - xnbufd\_unmap\_kwrite, [109](#)
  - xnbufd\_unmap\_uread, [110](#)
  - xnbufd\_unmap\_uwrite, [110](#)
- Buffer management services, [303](#)
  - a4l\_buf\_commit\_absget, [304](#)
  - a4l\_buf\_commit\_absput, [305](#)
  - a4l\_buf\_commit\_get, [305](#)
  - a4l\_buf\_commit\_put, [306](#)
  - a4l\_buf\_count, [306](#)
  - a4l\_buf\_evt, [308](#)
  - a4l\_buf\_get, [308](#)
  - a4l\_buf\_prepare\_absget, [309](#)
  - a4l\_buf\_prepare\_absput, [309](#)
  - a4l\_buf\_prepare\_get, [310](#)
  - a4l\_buf\_prepare\_put, [311](#)
  - a4l\_buf\_put, [311](#)
  - a4l\_get\_chan, [312](#)
  - a4l\_get\_cmd, [312](#)
- Buffer services, [453](#)
  - B\_PRIO, [454](#)
  - rt\_buffer\_bind, [454](#)
  - rt\_buffer\_clear, [455](#)
  - rt\_buffer\_create, [456](#)
  - rt\_buffer\_delete, [457](#)
  - rt\_buffer\_inquire, [457](#)
  - rt\_buffer\_read, [458](#)
  - rt\_buffer\_read\_timed, [459](#)
  - rt\_buffer\_read\_until, [460](#)
  - rt\_buffer\_unbind, [461](#)
  - rt\_buffer\_write, [461](#)
  - rt\_buffer\_write\_timed, [462](#)
  - rt\_buffer\_write\_until, [463](#)
- CAN Devices, [47](#)
  - CAN\_BITTIME\_TYPE, [70](#)
  - CAN\_CTRLMODE\_3\_SAMPLES, [57](#)
  - CAN\_CTRLMODE\_LISTENONLY, [58](#)
  - CAN\_CTRLMODE\_LOOPBACK, [58](#)
  - CAN\_ERR\_LOSTARB\_UNSPEC, [58](#)
  - CAN\_MODE, [70](#)
  - CAN\_RAW\_ERR\_FILTER, [58](#)
  - CAN\_RAW\_FILTER, [59](#)
  - CAN\_RAW\_LOOPBACK, [60](#)
  - CAN\_RAW\_RECV\_OWN\_MSGS, [60](#)
  - CAN\_STATE, [70](#)
  - can\_filter\_t, [69](#)
  - can\_frame\_t, [69](#)
  - RTCAN\_RTIOC\_RCV\_TIMEOUT, [61](#)
  - RTCAN\_RTIOC\_SND\_TIMEOUT, [61](#)
  - RTCAN\_RTIOC\_TAKE\_TIMESTAMP, [62](#)
  - SIOCGCANBAUDRATE, [63](#)
  - SIOCGCANCTRLMODE, [63](#)
  - SIOCGCANCUSTOMBITTIME, [64](#)
  - SIOCGCANSTATE, [64](#)
  - SIOCGIFINDEX, [65](#)
  - SIOCSCANBAUDRATE, [65](#)
  - SIOCSCANCTRLMODE, [66](#)
  - SIOCSCANCUSTOMBITTIME, [67](#)
  - SIOCSCANMODE, [68](#)
  - SOL\_CAN\_RAW, [69](#)
- CAN\_BITTIME\_TYPE
  - CAN Devices, [70](#)
- CAN\_CTRLMODE\_3\_SAMPLES
  - CAN Devices, [57](#)
- CAN\_CTRLMODE\_LISTENONLY
  - CAN Devices, [58](#)
- CAN\_CTRLMODE\_LOOPBACK
  - CAN Devices, [58](#)
- CAN\_ERR\_LOSTARB\_UNSPEC
  - CAN Devices, [58](#)
- CAN\_MODE
  - CAN Devices, [70](#)
- CAN\_RAW\_ERR\_FILTER
  - CAN Devices, [58](#)
- CAN\_RAW\_FILTER
  - CAN Devices, [59](#)
- CAN\_RAW\_LOOPBACK
  - CAN Devices, [60](#)
- CAN\_RAW\_RECV\_OWN\_MSGS
  - CAN Devices, [60](#)
- CAN\_STATE
  - CAN Devices, [70](#)
- COMPAT\_\_rt\_alarm\_create
  - Transition Kit, [567](#)
- COMPAT\_\_rt\_event\_clear
  - Transition Kit, [568](#)
- COMPAT\_\_rt\_event\_create
  - Transition Kit, [569](#)
- COMPAT\_\_rt\_event\_signal
  - Transition Kit, [570](#)
- COMPAT\_\_rt\_pipe\_create
  - Transition Kit, [571](#)
- COMPAT\_\_rt\_task\_create
  - Transition Kit, [572](#)

- COMPAT\_\_rt\_task\_set\_periodic
  - Transition Kit, [573](#)
- can\_bittime, [590](#)
- can\_bittime\_btr, [591](#)
- can\_bittime\_std, [591](#)
- can\_filter, [592](#)
  - can\_id, [592](#)
  - can\_mask, [593](#)
- can\_filter\_t
  - CAN Devices, [69](#)
- can\_frame, [593](#)
  - can\_id, [594](#)
- can\_frame\_t
  - CAN Devices, [69](#)
- can\_id
  - can\_filter, [592](#)
  - can\_frame, [594](#)
- can\_ifindex
  - sockaddr\_can, [617](#)
- can\_ifreq, [594](#)
- can\_mask
  - can\_filter, [593](#)
- Channels and ranges, [23](#)
- chans
  - a4l\_channels\_desc, [580](#)
- Clock Services, [176](#)
  - rtdm\_clock\_read, [176](#)
  - rtdm\_clock\_read\_monotonic, [176](#)
- Clock services, [112](#)
  - xnclock\_adjust, [112](#)
  - xnclock\_deregister, [113](#)
  - xnclock\_register, [113](#)
  - xnclock\_tick, [114](#)
- clock\_getres
  - Clocks and timers, [318](#)
- clock\_gettime
  - Clocks and timers, [319](#)
- clock\_nanosleep
  - Clocks and timers, [320](#)
- clock\_settime
  - Clocks and timers, [321](#)
- Clocks and timers, [317](#)
  - clock\_getres, [318](#)
  - clock\_gettime, [319](#)
  - clock\_nanosleep, [320](#)
  - clock\_settime, [321](#)
  - nanosleep, [322](#)
  - timer\_create, [322](#)
  - timer\_delete, [323](#)
  - timer\_getoverrun, [324](#)
  - timer\_gettime, [325](#)
  - timer\_settime, [326](#)
- close
  - rtdm\_fd\_ops, [609](#)
  - udd\_device, [620](#)
- close\_\_AF\_RTIPC
  - Real-time IPC, [92](#)
- Cobalt, [119](#)
- Cobalt kernel, [120](#)
- cobalt\_atomic\_enter
  - Big dual kernel lock, [27](#)
- cobalt\_atomic\_leave
  - Big dual kernel lock, [28](#)
- cobalt\_vfroot
  - Virtual file services, [294](#)
- Condition variable services, [464](#)
  - rt\_cond\_bind, [465](#)
  - rt\_cond\_broadcast, [466](#)
  - rt\_cond\_create, [466](#)
  - rt\_cond\_delete, [467](#)
  - rt\_cond\_inquire, [468](#)
  - rt\_cond\_signal, [468](#)
  - rt\_cond\_unbind, [469](#)
  - rt\_cond\_wait, [469](#)
  - rt\_cond\_wait\_timed, [470](#)
  - rt\_cond\_wait\_until, [471](#)
- Condition variables, [327](#)
  - pthread\_cond\_broadcast, [328](#)
  - pthread\_cond\_destroy, [329](#)
  - pthread\_cond\_init, [329](#)
  - pthread\_cond\_signal, [330](#)
  - pthread\_cond\_timedwait, [331](#)
  - pthread\_cond\_wait, [332](#)
  - pthread\_condattr\_destroy, [333](#)
  - pthread\_condattr\_getclock, [333](#)
  - pthread\_condattr\_getpshared, [334](#)
  - pthread\_condattr\_init, [335](#)
  - pthread\_condattr\_setclock, [335](#)
  - pthread\_condattr\_setpshared, [336](#)
- connect\_\_AF\_RTIPC
  - Real-time IPC, [92](#)
- context\_size
  - rtdm\_driver, [607](#)
- cpu
  - xnsched, [626](#)
- curr
  - xnsched, [626](#)
- data\_len
  - a4l\_cmd\_desc, [582](#)
- databuf
  - xnvfile\_snapshot\_iterator, [636](#)
- date
  - rt\_timer\_info, [602](#)
- Debugging services, [115](#)
- Descriptor API, [417](#)
  - a4l\_close, [417](#)
  - a4l\_fill\_desc, [418](#)
  - a4l\_get\_chinfo, [418](#)
  - a4l\_get\_rnginfo, [419](#)
  - a4l\_get\_subdinfo, [419](#)
  - a4l\_open, [420](#)
- Descriptor Syscall API, [415](#)
  - a4l\_sys\_desc, [415](#)
- device
  - rtdm\_dev\_context, [603](#)
- Device Profiles, [162](#)

- RTIOC\_DEVICE\_INFO, [164](#)
- RTIOC\_PURGE, [164](#)
- Device Registration Services, [165](#)
- RTDM\_DEVICE\_TYPE\_MASK, [166](#)
- RTDM\_EXCLUSIVE, [166](#)
- RTDM\_FIXED\_MINOR, [167](#)
- RTDM\_MAX\_MINOR, [167](#)
- RTDM\_NAMED\_DEVICE, [167](#)
- RTDM\_PROTOCOL\_DEVICE, [167](#)
- rtdm\_close\_handler, [167](#)
- rtdm\_dev\_register, [168](#)
- rtdm\_dev\_unregister, [168](#)
- rtdm\_drv\_set\_sysclass, [169](#)
- rtdm\_get\_unmapped\_area\_handler, [170](#)
- rtdm\_ioctl\_handler, [170](#)
- rtdm\_mmap\_handler, [171](#)
- rtdm\_open\_handler, [171](#)
- rtdm\_read\_handler, [172](#)
- rtdm\_recvmsg\_handler, [173](#)
- rtdm\_select\_handler, [173](#)
- rtdm\_sendmsg\_handler, [174](#)
- rtdm\_socket\_handler, [174](#)
- rtdm\_write\_handler, [175](#)
- device\_count
  - rtdm\_driver, [608](#)
- device\_flags
  - rtdm\_driver, [608](#)
  - udd\_device, [620](#)
- device\_subclass
  - udd\_device, [620](#)
- driver
  - rtdm\_device, [605](#)
- Driver API, [296](#)
- Driver management services, [297](#)
  - a4l\_register\_drv, [297](#)
  - a4l\_unregister\_drv, [298](#)
- Driver programming interface, [142](#)
- Driver to driver services, [143](#)
  - rtdm\_accept, [144](#)
  - rtdm\_bind, [145](#)
  - rtdm\_close, [146](#)
  - rtdm\_connect, [147](#)
  - rtdm\_getpeername, [148](#)
  - rtdm\_getsockname, [148](#)
  - rtdm\_getsockopt, [149](#)
  - rtdm\_ioctl, [150](#)
  - rtdm\_listen, [151](#)
  - rtdm\_open, [152](#)
  - rtdm\_read, [152](#)
  - rtdm\_recv, [153](#)
  - rtdm\_recvfrom, [154](#)
  - rtdm\_recvmsg, [155](#)
  - rtdm\_send, [155](#)
  - rtdm\_sendmsg, [156](#)
  - rtdm\_sendto, [157](#)
  - rtdm\_setsockopt, [158](#)
  - rtdm\_shutdown, [159](#)
  - rtdm\_socket, [159](#)
  - rtdm\_write, [160](#)
- driver.h
  - RTDM\_PROFILE\_INFO, [652](#)
  - rtdm\_fd\_device, [653](#)
  - rtdm\_fd\_is\_user, [653](#)
  - rtdm\_fd\_to\_private, [654](#)
  - rtdm\_private\_to\_fd, [654](#)
- driver\_name
  - a4l\_descriptor, [583](#)
- Dynamic memory allocation services, [116](#)
  - xnheap\_alloc, [116](#)
  - xnheap\_destroy, [117](#)
  - xnheap\_free, [117](#)
  - xnheap\_set\_name, [118](#)
- EV\_ANY
  - Event flag group services, [473](#)
- EV\_PRIO
  - Event flag group services, [473](#)
- end
  - xnfile\_regular\_ops, [631](#)
  - xnfile\_snapshot\_ops, [639](#)
- endfn
  - xnfile\_snapshot\_iterator, [637](#)
- Event flag group services, [472](#)
  - EV\_ANY, [473](#)
  - EV\_PRIO, [473](#)
  - rt\_event\_bind, [474](#)
  - rt\_event\_clear, [474](#)
  - rt\_event\_create, [475](#)
  - rt\_event\_delete, [476](#)
  - rt\_event\_inquire, [477](#)
  - rt\_event\_signal, [477](#)
  - rt\_event\_unbind, [478](#)
  - rt\_event\_wait, [478](#)
  - rt\_event\_wait\_timed, [479](#)
  - rt\_event\_wait\_until, [480](#)
- Event Services, [208](#)
  - rtdm\_event\_clear, [208](#)
  - rtdm\_event\_destroy, [210](#)
  - rtdm\_event\_init, [210](#)
  - rtdm\_event\_pulse, [210](#)
  - rtdm\_event\_select, [211](#)
  - rtdm\_event\_signal, [212](#)
  - rtdm\_event\_timedwait, [212](#)
  - rtdm\_event\_wait, [213](#)
- fd
  - a4l\_descriptor, [583](#)
- fd.h
  - rtdm\_fd\_get, [665](#)
  - rtdm\_fd\_lock, [666](#)
  - rtdm\_fd\_put, [666](#)
  - rtdm\_fd\_select, [667](#)
  - rtdm\_fd\_unlock, [667](#)
- flags
  - a4l\_channel, [579](#)
  - a4l\_range, [587](#)

- get
  - xnvmfile\_lock\_ops, 628
- get\_unmapped\_area
  - rtm\_fd\_ops, 610
- getpeername\_\_AF\_RTIPC
  - Real-time IPC, 93
- getsockname\_\_AF\_RTIPC
  - Real-time IPC, 93
- getsockopt\_\_AF\_RTIPC
  - Real-time IPC, 94
- H\_PRIO
  - Heap management services, 483
- Heap management services, 482
  - H\_PRIO, 483
  - rt\_heap\_alloc, 483
  - rt\_heap\_alloc\_timed, 483
  - rt\_heap\_alloc\_until, 485
  - rt\_heap\_bind, 486
  - rt\_heap\_create, 487
  - rt\_heap\_delete, 488
  - rt\_heap\_free, 489
  - rt\_heap\_inquire, 489
  - rt\_heap\_unbind, 490
- heapsize
  - RT\_HEAP\_INFO, 598
- htimer
  - xnsched, 627
- IDDP\_LABEL
  - Real-time IPC, 83
- IDDP\_POOLSZ
  - Real-time IPC, 84
- idx\_read\_subd
  - a4l\_descriptor, 583
- idx\_subd
  - a4l\_cmd\_desc, 582
  - a4l\_instruction, 586
- idx\_write\_subd
  - a4l\_descriptor, 583
- In-kernel arithmetics, 101
  - xnarch\_generic\_full\_divmod64, 101
- include/cobalt/kernel/rtm/analogy/channel\_↔  
range.h, 643
- include/cobalt/kernel/rtm/analogy/driver.h, 646
- include/cobalt/kernel/rtm/analogy/subdevice.h,  
654
- include/cobalt/kernel/rtm/driver.h, 647
- include/cobalt/kernel/rtm/fd.h, 664
- include/cobalt/kernel/rtm/udd.h, 689
- include/rtm/analogy.h, 691
- include/rtm/can.h, 656
- include/rtm/rtm.h, 674
- include/rtm/serial.h, 676
- include/rtm/testing.h, 686
- include/rtm/uapi/analogy.h, 693
- include/rtm/uapi/can.h, 657
- include/rtm/uapi/gpiopwm.h, 669
- include/rtm/uapi/ipc.h, 670
- include/rtm/uapi/rtm.h, 674
- include/rtm/uapi/serial.h, 677
- include/rtm/uapi/testing.h, 687
- include/rtm/uapi/udd.h, 690
- inesting
  - xnsched, 627
- interrupt
  - udd\_device, 621
- Interrupt management, 123
  - xnintr\_affinity, 123
  - xnintr\_attach, 125
  - xnintr\_destroy, 125
  - xnintr\_detach, 126
  - xnintr\_disable, 126
  - xnintr\_enable, 127
  - xnintr\_init, 127
- Interrupt Management Services, 223
  - rtm\_irq\_disable, 225
  - rtm\_irq\_enable, 225
  - rtm\_irq\_free, 226
  - rtm\_irq\_get\_arg, 224
  - rtm\_irq\_handler\_t, 224
  - rtm\_irq\_request, 226
- Interrupt management services, 313
  - a4l\_free\_irq, 313
  - a4l\_get\_irq, 314
  - a4l\_request\_irq, 314
- ioctl
  - udd\_device, 621
- ioctl\_nrt
  - rtm\_fd\_ops, 610
- ioctl\_rt
  - rtm\_fd\_ops, 610
- irq
  - udd\_device, 621
- label
  - rtm\_device, 605
  - rtipc\_port\_label, 614
- len
  - udd\_memregion, 624
- length
  - a4l\_channels\_desc, 580
- Level 0 API, 441
- Level 1 API, 430
- Level 2 API, 435
- lflags
  - xnsched, 627
- lib/analogy/async.c, 700
- lib/analogy/calibration.c, 701
- lib/analogy/calibration.h, 698
- lib/analogy/descriptor.c, 702
- lib/analogy/internal.h, 699
- lib/analogy/range.c, 703
- lib/analogy/root\_leaf.h, 704
- lib/analogy/sync.c, 705
- lib/analogy/sys.c, 706
- Lightweight key-to-object mapping service, 132
  - xnmap\_create, 133

- xnmap\_delete, 133
- xnmap\_enter, 134
- xnmap\_fetch, 134
- xnmap\_fetch\_noccheck, 135
- xnmap\_remove, 136
- Locking services, 130
  - splexit, 130
  - splhigh, 131
  - spltest, 131
- macb\_dma\_desc, 595
- macb\_tx\_skb, 595
- magic
  - a4l\_descriptor, 584
- Math API, 421
  - a4l\_math\_mean, 421
  - a4l\_math\_polyfit, 422
  - a4l\_math\_stddev, 422
  - a4l\_math\_stddev\_of\_mean, 423
- max
  - a4l\_range, 587
- mem\_regions
  - udd\_device, 622
- Message pipe services, 501
  - P\_MINOR\_AUTO, 502
  - P\_URGENT, 502
  - rt\_pipe\_bind, 502
  - rt\_pipe\_create, 503
  - rt\_pipe\_delete, 504
  - rt\_pipe\_read, 505
  - rt\_pipe\_read\_timed, 505
  - rt\_pipe\_read\_until, 506
  - rt\_pipe\_stream, 507
  - rt\_pipe\_unbind, 508
  - rt\_pipe\_write, 508
- Message queue services, 510
  - Q\_PRIO, 511
  - rt\_queue\_alloc, 511
  - rt\_queue\_bind, 512
  - rt\_queue\_create, 513
  - rt\_queue\_delete, 514
  - rt\_queue\_flush, 515
  - rt\_queue\_free, 515
  - rt\_queue\_inquire, 516
  - rt\_queue\_read, 516
  - rt\_queue\_read\_timed, 517
  - rt\_queue\_read\_until, 518
  - rt\_queue\_receive, 519
  - rt\_queue\_receive\_timed, 519
  - rt\_queue\_receive\_until, 520
  - rt\_queue\_send, 521
  - rt\_queue\_unbind, 522
- Message queues, 340
  - mq\_close, 341
  - mq\_getattr, 341
  - mq\_notify, 342
  - mq\_open, 343
  - mq\_receive, 345
  - mq\_send, 345
  - mq\_setattr, 346
  - mq\_timedreceive, 347
  - mq\_timedsend, 348
  - mq\_unlink, 349
- min
  - a4l\_range, 587
- minor
  - rtdm\_device, 605
- Misc services, 316
  - a4l\_get\_time, 316
- mmap
  - rtdm\_fd\_ops, 610
  - udd\_device, 622
- mode
  - a4l\_channels\_desc, 581
- mq\_close
  - Message queues, 341
- mq\_getattr
  - Message queues, 341
- mq\_notify
  - Message queues, 342
- mq\_open
  - Message queues, 343
- mq\_receive
  - Message queues, 345
- mq\_send
  - Message queues, 345
- mq\_setattr
  - Message queues, 346
- mq\_timedreceive
  - Message queues, 347
- mq\_timedsend
  - Message queues, 348
- mq\_unlink
  - Message queues, 349
- Mutex services, 219, 493
  - rt\_mutex\_acquire, 494
  - rt\_mutex\_acquire\_timed, 494
  - rt\_mutex\_acquire\_until, 495
  - rt\_mutex\_bind, 496
  - rt\_mutex\_create, 497
  - rt\_mutex\_delete, 497
  - rt\_mutex\_inquire, 498
  - rt\_mutex\_release, 499
  - rt\_mutex\_unbind, 499
  - rtdm\_mutex\_destroy, 219
  - rtdm\_mutex\_init, 220
  - rtdm\_mutex\_lock, 220
  - rtdm\_mutex\_timedlock, 221
  - rtdm\_mutex\_unlock, 221
- Mutual exclusion, 351
  - \_\_attribute\_\_, 352
  - pthread\_mutex\_destroy, 353
  - pthread\_mutex\_getprioceiling, 353
  - pthread\_mutex\_init, 354
  - pthread\_mutex\_lock, 355
  - pthread\_mutex\_setprioceiling, 356
  - pthread\_mutex\_timedlock, 357

- pthread\_mutex\_trylock, 358
- pthread\_mutex\_unlock, 358
- pthread\_mutexattr\_destroy, 359
- pthread\_mutexattr\_getprotocol, 360
- pthread\_mutexattr\_getpshared, 360
- pthread\_mutexattr\_gettype, 361
- pthread\_mutexattr\_init, 362
- pthread\_mutexattr\_setprotocol, 363
- pthread\_mutexattr\_setpshared, 364
- pthread\_mutexattr\_settype, 364
- nanosecs\_abs\_t
  - RTDM, 73
- nanosecs\_rel\_t
  - RTDM, 74
- nanosleep
  - Clocks and timers, 322
- nb\_bits
  - a4l\_channel, 579
- nb\_subd
  - a4l\_descriptor, 584
- next
  - xnvfile\_regular\_ops, 632
  - xnvfile\_snapshot\_ops, 639
- Non-Real-Time Signalling Services, 228
  - rtdm\_nrtsig\_destroy, 229
  - rtdm\_nrtsig\_handler\_t, 228
  - rtdm\_nrtsig\_init, 229
  - rtdm\_nrtsig\_pend, 230
  - rtdm\_schedule\_nrt\_work, 230
- nrddata
  - xnvfile\_snapshot\_iterator, 637
- open
  - rtdm\_fd\_ops, 610
  - udd\_device, 622
- owner
  - RT\_MUTEX\_INFO, 599
- P\_MINOR\_AUTO
  - Message pipe services, 502
- P\_URGENT
  - Message pipe services, 502
- POSIX interface, 338
- pSOS@ emulator, 566
- pid
  - udd\_notify, 625
- pos
  - xnvfile\_regular\_iterator, 629
- private
  - xnvfile\_regular\_iterator, 630
  - xnvfile\_snapshot\_iterator, 637
- Process scheduling, 366
  - sched\_get\_priority\_max, 367
  - sched\_get\_priority\_max\_ex, 367
  - sched\_get\_priority\_min, 368
  - sched\_get\_priority\_min\_ex, 369
  - sched\_getconfig\_np, 370
  - sched\_getscheduler, 371
  - sched\_getscheduler\_ex, 371
  - sched\_setconfig\_np, 372
  - sched\_setscheduler, 374
  - sched\_setscheduler\_ex, 375
  - sched\_yield, 376
- profile\_info
  - rtdm\_driver, 608
- program\_htick\_shot
  - Timer services, 279
- pthread\_cond\_broadcast
  - Condition variables, 328
- pthread\_cond\_destroy
  - Condition variables, 329
- pthread\_cond\_init
  - Condition variables, 329
- pthread\_cond\_signal
  - Condition variables, 330
- pthread\_cond\_timedwait
  - Condition variables, 331
- pthread\_cond\_wait
  - Condition variables, 332
- pthread\_condattr\_destroy
  - Condition variables, 333
- pthread\_condattr\_getclock
  - Condition variables, 333
- pthread\_condattr\_getpshared
  - Condition variables, 334
- pthread\_condattr\_init
  - Condition variables, 335
- pthread\_condattr\_setclock
  - Condition variables, 335
- pthread\_condattr\_setpshared
  - Condition variables, 336
- pthread\_create
  - Thread management, 386
- pthread\_getschedparam
  - Scheduling management, 392
- pthread\_getschedparam\_ex
  - Scheduling management, 393
- pthread\_join
  - Thread management, 388
- pthread\_kill
  - Thread management, 388
- pthread\_make\_periodic\_np
  - Transition Kit, 575
- pthread\_mutex\_destroy
  - Mutual exclusion, 353
- pthread\_mutex\_getprioceiling
  - Mutual exclusion, 353
- pthread\_mutex\_init
  - Mutual exclusion, 354
- pthread\_mutex\_lock
  - Mutual exclusion, 355
- pthread\_mutex\_setprioceiling
  - Mutual exclusion, 356
- pthread\_mutex\_timedlock
  - Mutual exclusion, 357
- pthread\_mutex\_trylock



- Mutual exclusion, [358](#)
- pthread\_mutex\_unlock
  - Mutual exclusion, [358](#)
- pthread\_mutexattr\_destroy
  - Mutual exclusion, [359](#)
- pthread\_mutexattr\_getprotocol
  - Mutual exclusion, [360](#)
- pthread\_mutexattr\_getpshared
  - Mutual exclusion, [360](#)
- pthread\_mutexattr\_gettype
  - Mutual exclusion, [361](#)
- pthread\_mutexattr\_init
  - Mutual exclusion, [362](#)
- pthread\_mutexattr\_setprotocol
  - Mutual exclusion, [363](#)
- pthread\_mutexattr\_setpshared
  - Mutual exclusion, [364](#)
- pthread\_mutexattr\_settype
  - Mutual exclusion, [364](#)
- pthread\_setmode\_np
  - Thread management, [389](#)
- pthread\_setname\_np
  - Thread management, [391](#)
- pthread\_setschedparam
  - Scheduling management, [394](#)
- pthread\_setschedparam\_ex
  - Scheduling management, [395](#)
- pthread\_wait\_np
  - Transition Kit, [576](#)
- pthread\_yield
  - Scheduling management, [396](#)
- put
  - xnvfile\_lock\_ops, [629](#)
- Q\_PRIO
  - Message queue services, [511](#)
- RT\_ALARM\_INFO, [596](#)
- RT\_BUFFER\_INFO, [596](#)
- RT\_COND\_INFO, [597](#)
- RT\_EVENT\_INFO, [597](#)
- RT\_HEAP\_INFO, [598](#)
  - heapsize, [598](#)
  - usablemem, [598](#)
  - usedmem, [599](#)
- RT\_MUTEX\_INFO, [599](#)
  - owner, [599](#)
- RT\_QUEUE\_INFO, [600](#)
- RT\_SEM\_INFO, [600](#)
- RT\_TASK\_INFO, [601](#)
- RT\_TIMER\_INFO
  - Timer management services, [562](#)
- RTCAN\_RTIOC\_RCV\_TIMEOUT
  - CAN Devices, [61](#)
- RTCAN\_RTIOC\_SND\_TIMEOUT
  - CAN Devices, [61](#)
- RTCAN\_RTIOC\_TAKE\_TIMESTAMP
  - CAN Devices, [62](#)
- RTDM User API, [75](#)
- RTDM\_DEVICE\_TYPE\_MASK
  - Device Registration Services, [166](#)
- RTDM\_EXCLUSIVE
  - Device Registration Services, [166](#)
- RTDM\_EXECUTE\_ATOMICALLY
  - Big dual kernel lock, [28](#)
- RTDM\_FIXED\_MINOR
  - Device Registration Services, [167](#)
- RTDM\_MAX\_MINOR
  - Device Registration Services, [167](#)
- RTDM\_NAMED\_DEVICE
  - Device Registration Services, [167](#)
- RTDM\_PROFILE\_INFO
  - driver.h, [652](#)
- RTDM\_PROTOCOL\_DEVICE
  - Device Registration Services, [167](#)
- RTDM\_TIMEOUT\_INFINITE
  - RTDM, [73](#)
- RTDM\_TIMEOUT\_NONE
  - RTDM, [73](#)
- RTDM, [72](#)
  - nanosecs\_abs\_t, [73](#)
  - nanosecs\_rel\_t, [74](#)
  - RTDM\_TIMEOUT\_INFINITE, [73](#)
  - RTDM\_TIMEOUT\_NONE, [73](#)
- RTIOC\_DEVICE\_INFO
  - Device Profiles, [164](#)
- RTIOC\_PURGE
  - Device Profiles, [164](#)
- RTSER\_RTIOC\_BREAK\_CTL
  - rtdm/uapi/serial.h, [682](#)
- RTSER\_RTIOC\_GET\_CONFIG
  - rtdm/uapi/serial.h, [682](#)
- RTSER\_RTIOC\_GET\_CONTROL
  - rtdm/uapi/serial.h, [683](#)
- RTSER\_RTIOC\_GET\_STATUS
  - rtdm/uapi/serial.h, [683](#)
- RTSER\_RTIOC\_SET\_CONFIG
  - rtdm/uapi/serial.h, [684](#)
- RTSER\_RTIOC\_SET\_CONTROL
  - rtdm/uapi/serial.h, [684](#)
- RTSER\_RTIOC\_WAIT\_EVENT
  - rtdm/uapi/serial.h, [685](#)
- Range / conversion API, [424](#)
  - a4l\_dtoraw, [424](#)
  - a4l\_find\_range, [425](#)
  - a4l\_ftoraw, [426](#)
  - a4l\_rawtod, [426](#)
  - a4l\_rawtof, [427](#)
  - a4l\_rawtoul, [427](#)
  - a4l\_sizeof\_chan, [428](#)
  - a4l\_sizeof\_subd, [428](#)
  - a4l\_ultoraw, [429](#)
- read\_nrt
  - rtdm\_fd\_ops, [610](#)
- read\_rt
  - rtdm\_fd\_ops, [611](#)
- Real-time IPC, [79](#)

- BUFP\_BUFSZ, [81](#)
- BUFP\_LABEL, [82](#)
- bind\_\_AF\_RTIPC, [90](#)
- close\_\_AF\_RTIPC, [92](#)
- connect\_\_AF\_RTIPC, [92](#)
- getpeername\_\_AF\_RTIPC, [93](#)
- getsockname\_\_AF\_RTIPC, [93](#)
- getsockopt\_\_AF\_RTIPC, [94](#)
- IDDP\_LABEL, [83](#)
- IDDP\_POOLSZ, [84](#)
- recvmsg\_\_AF\_RTIPC, [94](#)
- SO\_RCVTIMEO, [84](#)
- SO\_SNDTIMEO, [85](#)
- sendmsg\_\_AF\_RTIPC, [95](#)
- setsockopt\_\_AF\_RTIPC, [96](#)
- socket\_\_AF\_RTIPC, [96](#)
- XDDP\_BUFSZ, [85](#)
- XDDP\_EVTDOWN, [86](#)
- XDDP\_EVTIN, [86](#)
- XDDP\_EVTNOBUF, [86](#)
- XDDP\_EVTOUT, [87](#)
- XDDP\_LABEL, [87](#)
- XDDP\_MONITOR, [88](#)
- XDDP\_POOLSZ, [88](#)
- recvmsg\_\_AF\_RTIPC
  - Real-time IPC, [94](#)
- recvmsg\_nrt
  - rtdm\_fd\_ops, [611](#)
- recvmsg\_rt
  - rtdm\_fd\_ops, [611](#)
- Registry services, [137](#)
  - xnregistry\_bind, [137](#)
  - xnregistry\_enter, [138](#)
  - xnregistry\_lookup, [139](#)
  - xnregistry\_remove, [140](#)
  - xnregistry\_unlink, [141](#)
- resched
  - xnsched, [627](#)
- rev
  - xnfile\_rev\_tag, [634](#)
- rewind
  - xnfile\_regular\_ops, [632](#)
  - xnfile\_snapshot\_ops, [640](#)
- rrbtimer
  - xnsched, [627](#)
- rt
  - xnsched, [627](#)
- rt\_alarm\_create
  - Alarm services, [449](#)
- rt\_alarm\_delete
  - Alarm services, [449](#)
- rt\_alarm\_inquire
  - Alarm services, [450](#)
- rt\_alarm\_start
  - Alarm services, [451](#)
- rt\_alarm\_stop
  - Alarm services, [451](#)
- rt\_alarm\_wait
  - Transition Kit, [577](#)
- rt\_buffer\_bind
  - Buffer services, [454](#)
- rt\_buffer\_clear
  - Buffer services, [455](#)
- rt\_buffer\_create
  - Buffer services, [456](#)
- rt\_buffer\_delete
  - Buffer services, [457](#)
- rt\_buffer\_inquire
  - Buffer services, [457](#)
- rt\_buffer\_read
  - Buffer services, [458](#)
- rt\_buffer\_read\_timed
  - Buffer services, [459](#)
- rt\_buffer\_read\_until
  - Buffer services, [460](#)
- rt\_buffer\_unbind
  - Buffer services, [461](#)
- rt\_buffer\_write
  - Buffer services, [461](#)
- rt\_buffer\_write\_timed
  - Buffer services, [462](#)
- rt\_buffer\_write\_until
  - Buffer services, [463](#)
- rt\_cond\_bind
  - Condition variable services, [465](#)
- rt\_cond\_broadcast
  - Condition variable services, [466](#)
- rt\_cond\_create
  - Condition variable services, [466](#)
- rt\_cond\_delete
  - Condition variable services, [467](#)
- rt\_cond\_inquire
  - Condition variable services, [468](#)
- rt\_cond\_signal
  - Condition variable services, [468](#)
- rt\_cond\_unbind
  - Condition variable services, [469](#)
- rt\_cond\_wait
  - Condition variable services, [469](#)
- rt\_cond\_wait\_timed
  - Condition variable services, [470](#)
- rt\_cond\_wait\_until
  - Condition variable services, [471](#)
- rt\_event\_bind
  - Event flag group services, [474](#)
- rt\_event\_clear
  - Event flag group services, [474](#)
- rt\_event\_create
  - Event flag group services, [475](#)
- rt\_event\_delete
  - Event flag group services, [476](#)
- rt\_event\_inquire
  - Event flag group services, [477](#)
- rt\_event\_signal
  - Event flag group services, [477](#)
- rt\_event\_unbind

- Event flag group services, [478](#)
- rt\_event\_wait
  - Event flag group services, [478](#)
- rt\_event\_wait\_timed
  - Event flag group services, [479](#)
- rt\_event\_wait\_until
  - Event flag group services, [480](#)
- rt\_heap\_alloc
  - Heap management services, [483](#)
- rt\_heap\_alloc\_timed
  - Heap management services, [483](#)
- rt\_heap\_alloc\_until
  - Heap management services, [485](#)
- rt\_heap\_bind
  - Heap management services, [486](#)
- rt\_heap\_create
  - Heap management services, [487](#)
- rt\_heap\_delete
  - Heap management services, [488](#)
- rt\_heap\_free
  - Heap management services, [489](#)
- rt\_heap\_inquire
  - Heap management services, [489](#)
- rt\_heap\_unbind
  - Heap management services, [490](#)
- rt\_mutex\_acquire
  - Mutex services, [494](#)
- rt\_mutex\_acquire\_timed
  - Mutex services, [494](#)
- rt\_mutex\_acquire\_until
  - Mutex services, [495](#)
- rt\_mutex\_bind
  - Mutex services, [496](#)
- rt\_mutex\_create
  - Mutex services, [497](#)
- rt\_mutex\_delete
  - Mutex services, [497](#)
- rt\_mutex\_inquire
  - Mutex services, [498](#)
- rt\_mutex\_release
  - Mutex services, [499](#)
- rt\_mutex\_unbind
  - Mutex services, [499](#)
- rt\_pipe\_bind
  - Message pipe services, [502](#)
- rt\_pipe\_create
  - Message pipe services, [503](#)
- rt\_pipe\_delete
  - Message pipe services, [504](#)
- rt\_pipe\_read
  - Message pipe services, [505](#)
- rt\_pipe\_read\_timed
  - Message pipe services, [505](#)
- rt\_pipe\_read\_until
  - Message pipe services, [506](#)
- rt\_pipe\_stream
  - Message pipe services, [507](#)
- rt\_pipe\_unbind
  - Message pipe services, [508](#)
- rt\_pipe\_write
  - Message pipe services, [508](#)
- rt\_queue\_alloc
  - Message queue services, [511](#)
- rt\_queue\_bind
  - Message queue services, [512](#)
- rt\_queue\_create
  - Message queue services, [513](#)
- rt\_queue\_delete
  - Message queue services, [514](#)
- rt\_queue\_flush
  - Message queue services, [515](#)
- rt\_queue\_free
  - Message queue services, [515](#)
- rt\_queue\_inquire
  - Message queue services, [516](#)
- rt\_queue\_read
  - Message queue services, [516](#)
- rt\_queue\_read\_timed
  - Message queue services, [517](#)
- rt\_queue\_read\_until
  - Message queue services, [518](#)
- rt\_queue\_receive
  - Message queue services, [519](#)
- rt\_queue\_receive\_timed
  - Message queue services, [519](#)
- rt\_queue\_receive\_until
  - Message queue services, [520](#)
- rt\_queue\_send
  - Message queue services, [521](#)
- rt\_queue\_unbind
  - Message queue services, [522](#)
- rt\_sem\_bind
  - Semaphore services, [524](#)
- rt\_sem\_broadcast
  - Semaphore services, [525](#)
- rt\_sem\_create
  - Semaphore services, [526](#)
- rt\_sem\_delete
  - Semaphore services, [527](#)
- rt\_sem\_inquire
  - Semaphore services, [527](#)
- rt\_sem\_p
  - Semaphore services, [528](#)
- rt\_sem\_p\_timed
  - Semaphore services, [528](#)
- rt\_sem\_p\_until
  - Semaphore services, [529](#)
- rt\_sem\_unbind
  - Semaphore services, [530](#)
- rt\_sem\_v
  - Semaphore services, [530](#)
- rt\_task\_bind
  - Task management services, [535](#)
- rt\_task\_create
  - Task management services, [536](#)
- rt\_task\_delete

- Task management services, [537](#)
- rt\_task\_inquire
  - Task management services, [538](#)
- rt\_task\_join
  - Task management services, [539](#)
- rt\_task\_receive
  - Task management services, [540](#)
- rt\_task\_receive\_timed
  - Task management services, [540](#)
- rt\_task\_receive\_until
  - Task management services, [541](#)
- rt\_task\_reply
  - Task management services, [542](#)
- rt\_task\_resume
  - Task management services, [543](#)
- rt\_task\_same
  - Task management services, [544](#)
- rt\_task\_self
  - Task management services, [544](#)
- rt\_task\_send
  - Task management services, [545](#)
- rt\_task\_send\_timed
  - Task management services, [545](#)
- rt\_task\_send\_until
  - Task management services, [547](#)
- rt\_task\_set\_affinity
  - Task management services, [548](#)
- rt\_task\_set\_mode
  - Task management services, [548](#)
- rt\_task\_set\_periodic
  - Task management services, [549](#)
- rt\_task\_set\_priority
  - Task management services, [550](#)
- rt\_task\_shadow
  - Task management services, [551](#)
- rt\_task\_sleep
  - Task management services, [553](#)
- rt\_task\_sleep\_until
  - Task management services, [554](#)
- rt\_task\_slice
  - Task management services, [555](#)
- rt\_task\_spawn
  - Task management services, [555](#)
- rt\_task\_start
  - Task management services, [556](#)
- rt\_task\_suspend
  - Task management services, [557](#)
- rt\_task\_unbind
  - Task management services, [558](#)
- rt\_task\_unblock
  - Task management services, [558](#)
- rt\_task\_wait\_period
  - Task management services, [559](#)
- rt\_task\_yield
  - Task management services, [560](#)
- rt\_timer\_info, [602](#)
  - date, [602](#)
- rt\_timer\_inquire
  - Timer management services, [562](#)
- rt\_timer\_ns2ticks
  - Timer management services, [562](#)
- rt\_timer\_read
  - Timer management services, [563](#)
- rt\_timer\_spin
  - Timer management services, [563](#)
- rt\_timer\_ticks2ns
  - Timer management services, [564](#)
- rtdm/uapi/serial.h
  - RTSER\_RTIOC\_BREAK\_CTL, [682](#)
  - RTSER\_RTIOC\_GET\_CONFIG, [682](#)
  - RTSER\_RTIOC\_GET\_CONTROL, [683](#)
  - RTSER\_RTIOC\_GET\_STATUS, [683](#)
  - RTSER\_RTIOC\_SET\_CONFIG, [684](#)
  - RTSER\_RTIOC\_SET\_CONTROL, [684](#)
  - RTSER\_RTIOC\_WAIT\_EVENT, [685](#)
- rtdm\_accept
  - Driver to driver services, [144](#)
- rtdm\_available
  - Utility Services, [232](#)
- rtdm\_bind
  - Driver to driver services, [145](#)
- rtdm\_clock\_read
  - Clock Services, [176](#)
- rtdm\_clock\_read\_monotonic
  - Clock Services, [176](#)
- rtdm\_close
  - Driver to driver services, [146](#)
- rtdm\_close\_handler
  - Device Registration Services, [167](#)
- rtdm\_connect
  - Driver to driver services, [147](#)
- rtdm\_copy\_from\_user
  - Utility Services, [232](#)
- rtdm\_copy\_to\_user
  - Utility Services, [233](#)
- rtdm\_dev\_context, [603](#)
  - device, [603](#)
- rtdm\_dev\_register
  - Device Registration Services, [168](#)
- rtdm\_dev\_unregister
  - Device Registration Services, [168](#)
- rtdm\_device, [604](#)
  - driver, [605](#)
  - label, [605](#)
  - minor, [605](#)
- rtdm\_device\_info, [606](#)
- rtdm\_driver, [606](#)
  - base\_minor, [607](#)
  - context\_size, [607](#)
  - device\_count, [608](#)
  - device\_flags, [608](#)
  - profile\_info, [608](#)
- rtdm\_drv\_set\_sysclass
  - Device Registration Services, [169](#)
- rtdm\_event\_clear
  - Event Services, [208](#)

- rtm\_event\_destroy
  - Event Services, 210
- rtm\_event\_init
  - Event Services, 210
- rtm\_event\_pulse
  - Event Services, 210
- rtm\_event\_select
  - Event Services, 211
- rtm\_event\_signal
  - Event Services, 212
- rtm\_event\_timedwait
  - Event Services, 212
- rtm\_event\_wait
  - Event Services, 213
- rtm\_fd\_device
  - driver.h, 653
- rtm\_fd\_get
  - fd.h, 665
- rtm\_fd\_is\_user
  - driver.h, 653
- rtm\_fd\_lock
  - fd.h, 666
- rtm\_fd\_ops, 608
  - close, 609
  - get\_unmapped\_area, 610
  - ioctl\_nrt, 610
  - ioctl\_rt, 610
  - mmap, 610
  - open, 610
  - read\_nrt, 610
  - read\_rt, 611
  - recvmsg\_nrt, 611
  - recvmsg\_rt, 611
  - select, 611
  - sendmsg\_nrt, 611
  - sendmsg\_rt, 611
  - socket, 612
  - write\_nrt, 612
  - write\_rt, 612
- rtm\_fd\_put
  - fd.h, 666
- rtm\_fd\_select
  - fd.h, 667
- rtm\_fd\_to\_private
  - driver.h, 654
- rtm\_fd\_unlock
  - fd.h, 667
- rtm\_for\_each\_waiter
  - Synchronisation Services, 196
- rtm\_for\_each\_waiter\_safe
  - Synchronisation Services, 197
- rtm\_free
  - Utility Services, 234
- rtm\_get\_unmapped\_area\_handler
  - Device Registration Services, 170
- rtm\_getpeername
  - Driver to driver services, 148
- rtm\_getsockname
  - Driver to driver services, 148
- rtm\_getsockopt
  - Driver to driver services, 149
- rtm\_gpio\_readout, 612
- rtm\_in\_rt\_context
  - Utility Services, 234
- rtm\_ioctl
  - Driver to driver services, 150
- rtm\_ioctl\_handler
  - Device Registration Services, 170
- rtm\_iomap\_to\_user
  - Utility Services, 234
- rtm\_irq\_disable
  - Interrupt Management Services, 225
- rtm\_irq\_enable
  - Interrupt Management Services, 225
- rtm\_irq\_free
  - Interrupt Management Services, 226
- rtm\_irq\_get\_arg
  - Interrupt Management Services, 224
- rtm\_irq\_handler\_t
  - Interrupt Management Services, 224
- rtm\_irq\_request
  - Interrupt Management Services, 226
- rtm\_listen
  - Driver to driver services, 151
- rtm\_lock\_get
  - Spinlock with preemption deactivation, 32
- rtm\_lock\_get\_irqsave
  - Spinlock with preemption deactivation, 31
- rtm\_lock\_init
  - Spinlock with preemption deactivation, 32
- rtm\_lock\_irqrestore
  - Spinlock with preemption deactivation, 31
- rtm\_lock\_irqsave
  - Spinlock with preemption deactivation, 31
- rtm\_lock\_put
  - Spinlock with preemption deactivation, 33
- rtm\_lock\_put\_irqrestore
  - Spinlock with preemption deactivation, 33
- rtm\_malloc
  - Utility Services, 235
- rtm\_mmap\_handler
  - Device Registration Services, 171
- rtm\_mmap\_iomem
  - Utility Services, 236
- rtm\_mmap\_kmem
  - Utility Services, 237
- rtm\_mmap\_to\_user
  - Utility Services, 237
- rtm\_mmap\_vmem
  - Utility Services, 238
- rtm\_munmap
  - Utility Services, 239
- rtm\_mutex\_destroy
  - Mutex services, 219
- rtm\_mutex\_init
  - Mutex services, 220

- rtdm\_mutex\_lock
  - Mutex services, [220](#)
- rtdm\_mutex\_timedlock
  - Mutex services, [221](#)
- rtdm\_mutex\_unlock
  - Mutex services, [221](#)
- rtdm\_nrtsig\_destroy
  - Non-Real-Time Signalling Services, [229](#)
- rtdm\_nrtsig\_handler\_t
  - Non-Real-Time Signalling Services, [228](#)
- rtdm\_nrtsig\_init
  - Non-Real-Time Signalling Services, [229](#)
- rtdm\_nrtsig\_pend
  - Non-Real-Time Signalling Services, [230](#)
- rtdm\_open
  - Driver to driver services, [152](#)
- rtdm\_open\_handler
  - Device Registration Services, [171](#)
- rtdm\_printk
  - Utility Services, [239](#)
- rtdm\_printk\_ratelimited
  - Utility Services, [240](#)
- rtdm\_private\_to\_fd
  - driver.h, [654](#)
- rtdm\_profile\_info, [613](#)
- rtdm\_read
  - Driver to driver services, [152](#)
- rtdm\_read\_handler
  - Device Registration Services, [172](#)
- rtdm\_read\_user\_ok
  - Utility Services, [240](#)
- rtdm\_rcv
  - Driver to driver services, [153](#)
- rtdm\_rcvfrom
  - Driver to driver services, [154](#)
- rtdm\_rcvmsg
  - Driver to driver services, [155](#)
- rtdm\_rcvmsg\_handler
  - Device Registration Services, [173](#)
- rtdm\_rt\_capable
  - Utility Services, [241](#)
- rtdm\_rw\_user\_ok
  - Utility Services, [242](#)
- rtdm\_safe\_copy\_from\_user
  - Utility Services, [242](#)
- rtdm\_safe\_copy\_to\_user
  - Utility Services, [243](#)
- rtdm\_schedule\_nrt\_work
  - Non-Real-Time Signalling Services, [230](#)
- rtdm\_select\_handler
  - Device Registration Services, [173](#)
- rtdm\_selecttype
  - Synchronisation Services, [196](#)
- rtdm\_sem\_destroy
  - Semaphore Services, [214](#)
- rtdm\_sem\_down
  - Semaphore Services, [215](#)
- rtdm\_sem\_init
  - Semaphore Services, [215](#)
- rtdm\_sem\_select
  - Semaphore Services, [216](#)
- rtdm\_sem\_timeddown
  - Semaphore Services, [216](#)
- rtdm\_sem\_up
  - Semaphore Services, [217](#)
- rtdm\_send
  - Driver to driver services, [155](#)
- rtdm\_sendmsg
  - Driver to driver services, [156](#)
- rtdm\_sendmsg\_handler
  - Device Registration Services, [174](#)
- rtdm\_sendto
  - Driver to driver services, [157](#)
- rtdm\_setsockopt
  - Driver to driver services, [158](#)
- rtdm\_shutdown
  - Driver to driver services, [159](#)
- rtdm\_sm\_ops, [613](#)
- rtdm\_socket
  - Driver to driver services, [159](#)
- rtdm\_socket\_handler
  - Device Registration Services, [174](#)
- rtdm\_spi\_config, [614](#)
- rtdm\_strncpy\_from\_user
  - Utility Services, [243](#)
- rtdm\_task\_busy\_sleep
  - Task Services, [180](#)
- rtdm\_task\_busy\_wait
  - Task Services, [180](#)
- rtdm\_task\_current
  - Task Services, [181](#)
- rtdm\_task\_destroy
  - Task Services, [181](#)
- rtdm\_task\_init
  - Task Services, [182](#)
- rtdm\_task\_join
  - Task Services, [182](#)
- rtdm\_task\_proc\_t
  - Task Services, [179](#)
- rtdm\_task\_set\_period
  - Task Services, [183](#)
- rtdm\_task\_set\_priority
  - Task Services, [183](#)
- rtdm\_task\_should\_stop
  - Task Services, [184](#)
- rtdm\_task\_sleep
  - Task Services, [184](#)
- rtdm\_task\_sleep\_abs
  - Task Services, [185](#)
- rtdm\_task\_sleep\_until
  - Task Services, [185](#)
- rtdm\_task\_unblock
  - Task Services, [186](#)
- rtdm\_task\_wait\_period
  - Task Services, [186](#)
- rtdm\_timedwait

- Synchronisation Services, 197
- rtdm\_timedwait\_condition
  - Synchronisation Services, 198
- rtdm\_timedwait\_condition\_locked
  - Synchronisation Services, 199
- rtdm\_timedwait\_locked
  - Synchronisation Services, 200
- rtdm\_timer\_destroy
  - Timer Services, 190
- rtdm\_timer\_handler\_t
  - Timer Services, 190
- rtdm\_timer\_init
  - Timer Services, 191
- rtdm\_timer\_mode
  - Timer Services, 190
- rtdm\_timer\_start
  - Timer Services, 191
- rtdm\_timer\_start\_in\_handler
  - Timer Services, 192
- rtdm\_timer\_stop
  - Timer Services, 192
- rtdm\_timer\_stop\_in\_handler
  - Timer Services, 193
- rtdm\_toseq\_init
  - Synchronisation Services, 200
- rtdm\_wait
  - Synchronisation Services, 201
- rtdm\_wait\_complete
  - Task Services, 187
- rtdm\_wait\_condition
  - Synchronisation Services, 202
- rtdm\_wait\_condition\_locked
  - Synchronisation Services, 202
- rtdm\_wait\_is\_completed
  - Task Services, 187
- rtdm\_wait\_locked
  - Synchronisation Services, 203
- rtdm\_wait\_prepare
  - Task Services, 188
- rtdm\_waitqueue\_broadcast
  - Synchronisation Services, 204
- rtdm\_waitqueue\_destroy
  - Synchronisation Services, 204
- rtdm\_waitqueue\_flush
  - Synchronisation Services, 205
- rtdm\_waitqueue\_init
  - Synchronisation Services, 205
- rtdm\_waitqueue\_lock
  - Synchronisation Services, 206
- rtdm\_waitqueue\_signal
  - Synchronisation Services, 206
- rtdm\_waitqueue\_unlock
  - Synchronisation Services, 207
- rtdm\_waitqueue\_wakeup
  - Synchronisation Services, 207
- rtdm\_write
  - Driver to driver services, 160
- rtdm\_write\_handler
  - Device Registration Services, 175
- rtipc\_port\_label, 614
  - label, 614
- rtser\_config, 615
- rtser\_event, 616
- rtser\_status, 616
- S\_PRIO
  - Semaphore services, 524
- SCHED\_QUOTA scheduling policy, 245
- SIOCGCANBAUDRATE
  - CAN Devices, 63
- SIOCGCANCTRLMODE
  - CAN Devices, 63
- SIOCGCANCUSTOMBITTIME
  - CAN Devices, 64
- SIOCGCANSTATE
  - CAN Devices, 64
- SIOCGIFINDEX
  - CAN Devices, 65
- SIOCSCANBAUDRATE
  - CAN Devices, 65
- SIOCSCANCTRLMODE
  - CAN Devices, 66
- SIOCSCANCUSTOMBITTIME
  - CAN Devices, 67
- SIOCSCANMODE
  - CAN Devices, 68
- SO\_RCVTIMEO
  - Real-time IPC, 84
- SO\_SNDTIMEO
  - Real-time IPC, 85
- SOL\_CAN\_RAW
  - CAN Devices, 69
- sbddata
  - a4l\_descriptor, 584
- sbsize
  - a4l\_descriptor, 584
- sched\_get\_priority\_max
  - Process scheduling, 367
- sched\_get\_priority\_max\_ex
  - Process scheduling, 367
- sched\_get\_priority\_min
  - Process scheduling, 368
- sched\_get\_priority\_min\_ex
  - Process scheduling, 369
- sched\_getconfig\_np
  - Process scheduling, 370
- sched\_getscheduler
  - Process scheduling, 371
- sched\_getscheduler\_ex
  - Process scheduling, 371
- sched\_setconfig\_np
  - Process scheduling, 372
- sched\_setscheduler
  - Process scheduling, 374
- sched\_setscheduler\_ex
  - Process scheduling, 375
- sched\_yield

- Process scheduling, [376](#)
- Scheduling management, [392](#)
  - [pthread\\_getschedparam](#), [392](#)
  - [pthread\\_getschedparam\\_ex](#), [393](#)
  - [pthread\\_setschedparam](#), [394](#)
  - [pthread\\_setschedparam\\_ex](#), [395](#)
  - [pthread\\_yield](#), [396](#)
- select
  - [rtdm\\_fd\\_ops](#), [611](#)
- [sem\\_close](#)
  - Semaphores, [379](#)
- [sem\\_destroy](#)
  - Semaphores, [379](#)
- [sem\\_init](#)
  - Semaphores, [380](#)
- [sem\\_post](#)
  - Semaphores, [381](#)
- [sem\\_timedwait](#)
  - Semaphores, [382](#)
- [sem\\_trywait](#)
  - Semaphores, [383](#)
- [sem\\_unlink](#)
  - Semaphores, [384](#)
- [sem\\_wait](#)
  - Semaphores, [385](#)
- Semaphore Services, [214](#)
  - [rtdm\\_sem\\_destroy](#), [214](#)
  - [rtdm\\_sem\\_down](#), [215](#)
  - [rtdm\\_sem\\_init](#), [215](#)
  - [rtdm\\_sem\\_select](#), [216](#)
  - [rtdm\\_sem\\_timeddown](#), [216](#)
  - [rtdm\\_sem\\_up](#), [217](#)
- Semaphore services, [523](#)
  - [rt\\_sem\\_bind](#), [524](#)
  - [rt\\_sem\\_broadcast](#), [525](#)
  - [rt\\_sem\\_create](#), [526](#)
  - [rt\\_sem\\_delete](#), [527](#)
  - [rt\\_sem\\_inquire](#), [527](#)
  - [rt\\_sem\\_p](#), [528](#)
  - [rt\\_sem\\_p\\_timed](#), [528](#)
  - [rt\\_sem\\_p\\_until](#), [529](#)
  - [rt\\_sem\\_unbind](#), [530](#)
  - [rt\\_sem\\_v](#), [530](#)
  - [S\\_PRIO](#), [524](#)
- Semaphores, [378](#)
  - [sem\\_close](#), [379](#)
  - [sem\\_destroy](#), [379](#)
  - [sem\\_init](#), [380](#)
  - [sem\\_post](#), [381](#)
  - [sem\\_timedwait](#), [382](#)
  - [sem\\_trywait](#), [383](#)
  - [sem\\_unlink](#), [384](#)
  - [sem\\_wait](#), [385](#)
- [sendmsg\\_\\_AF\\_RTIPC](#)
  - Real-time IPC, [95](#)
- [sendmsg\\_nrt](#)
  - [rtdm\\_fd\\_ops](#), [611](#)
- [sendmsg\\_rt](#)
  - [rtdm\\_fd\\_ops](#), [611](#)
- seq
  - [xnvfile\\_regular\\_iterator](#), [630](#)
  - [xnvfile\\_snapshot\\_iterator](#), [637](#)
- Serial Devices, [76](#)
- [setsockopt\\_\\_AF\\_RTIPC](#)
  - Real-time IPC, [96](#)
- show
  - [xnvfile\\_regular\\_ops](#), [633](#)
  - [xnvfile\\_snapshot\\_ops](#), [640](#)
- sig
  - [udd\\_signotify](#), [625](#)
- [sipc\\_port](#)
  - [sockaddr\\_ipc](#), [618](#)
- Smokey API, [398](#)
- [sockaddr\\_can](#), [617](#)
  - [can\\_ifindex](#), [617](#)
- [sockaddr\\_ipc](#), [618](#)
  - [sipc\\_port](#), [618](#)
- socket
  - [rtdm\\_fd\\_ops](#), [612](#)
- [socket\\_\\_AF\\_RTIPC](#)
  - Real-time IPC, [96](#)
- Software calibration API, [412](#)
  - [a4l\\_dcaltoraw](#), [412](#)
  - [a4l\\_get\\_softcal\\_converter](#), [413](#)
  - [a4l\\_rawtodcal](#), [413](#)
  - [a4l\\_read\\_calibration\\_file](#), [414](#)
- Spinlock with preemption deactivation, [30](#)
  - [rtdm\\_lock\\_get](#), [32](#)
  - [rtdm\\_lock\\_get\\_irqsave](#), [31](#)
  - [rtdm\\_lock\\_init](#), [32](#)
  - [rtdm\\_lock\\_irqrestore](#), [31](#)
  - [rtdm\\_lock\\_irqsave](#), [31](#)
  - [rtdm\\_lock\\_put](#), [33](#)
  - [rtdm\\_lock\\_put\\_irqrestore](#), [33](#)
- splexit
  - Locking services, [130](#)
- splhigh
  - Locking services, [131](#)
- spltest
  - Locking services, [131](#)
- status
  - [xnsched](#), [627](#)
- store
  - [xnvfile\\_regular\\_ops](#), [633](#)
  - [xnvfile\\_snapshot\\_ops](#), [641](#)
- Subdevice management services, [299](#)
  - [a4l\\_add\\_subd](#), [301](#)
  - [a4l\\_alloc\\_subd](#), [301](#)
  - [a4l\\_get\\_subd](#), [302](#)
- [switch\\_htick\\_mode](#)
  - Timer services, [280](#)
- Synchronisation Services, [194](#)
  - [rtdm\\_for\\_each\\_waiter](#), [196](#)
  - [rtdm\\_for\\_each\\_waiter\\_safe](#), [197](#)
  - [rtdm\\_selecttype](#), [196](#)
  - [rtdm\\_timedwait](#), [197](#)



- rt dm\_timedwait\_condition, 198
- rt dm\_timedwait\_condition\_locked, 199
- rt dm\_timedwait\_locked, 200
- rt dm\_toseq\_init, 200
- rt dm\_wait, 201
- rt dm\_wait\_condition, 202
- rt dm\_wait\_condition\_locked, 202
- rt dm\_wait\_locked, 203
- rt dm\_waitqueue\_broadcast, 204
- rt dm\_waitqueue\_destroy, 204
- rt dm\_waitqueue\_flush, 205
- rt dm\_waitqueue\_init, 205
- rt dm\_waitqueue\_lock, 206
- rt dm\_waitqueue\_signal, 206
- rt dm\_waitqueue\_unlock, 207
- rt dm\_waitqueue\_wakeup, 207
- Synchronous acquisition API, 431, 436
  - a4l\_config\_subd, 436
  - a4l\_snd\_insn, 433
  - a4l\_snd\_insnlist, 433
  - a4l\_sync\_dio, 437
  - a4l\_sync\_read, 437
  - a4l\_sync\_write, 438
- Synchronous I/O multiplexing, 248
  - xnselect, 249
  - xnselect\_bind, 250
  - xnselect\_destroy, 250
  - xnselect\_init, 251
  - xnselect\_signal, 251
  - xnselector\_destroy, 252
  - xnselector\_init, 252
- T\_LOCK
  - Task management services, 534
- T\_LOPRIO
  - Task management services, 534
- T\_WARNSW
  - Task management services, 534
- Task management services, 532
  - rt\_task\_bind, 535
  - rt\_task\_create, 536
  - rt\_task\_delete, 537
  - rt\_task\_inquire, 538
  - rt\_task\_join, 539
  - rt\_task\_receive, 540
  - rt\_task\_receive\_timed, 540
  - rt\_task\_receive\_until, 541
  - rt\_task\_reply, 542
  - rt\_task\_resume, 543
  - rt\_task\_same, 544
  - rt\_task\_self, 544
  - rt\_task\_send, 545
  - rt\_task\_send\_timed, 545
  - rt\_task\_send\_until, 547
  - rt\_task\_set\_affinity, 548
  - rt\_task\_set\_mode, 548
  - rt\_task\_set\_periodic, 549
  - rt\_task\_set\_priority, 550
  - rt\_task\_shadow, 551
  - rt\_task\_sleep, 553
  - rt\_task\_sleep\_until, 554
  - rt\_task\_slice, 555
  - rt\_task\_spawn, 555
  - rt\_task\_start, 556
  - rt\_task\_suspend, 557
  - rt\_task\_unbind, 558
  - rt\_task\_unblock, 558
  - rt\_task\_wait\_period, 559
  - rt\_task\_yield, 560
  - T\_LOCK, 534
  - T\_LOPRIO, 534
  - T\_WARNSW, 534
- Task Services, 178
  - rt dm\_task\_busy\_sleep, 180
  - rt dm\_task\_busy\_wait, 180
  - rt dm\_task\_current, 181
  - rt dm\_task\_destroy, 181
  - rt dm\_task\_init, 182
  - rt dm\_task\_join, 182
  - rt dm\_task\_proc\_t, 179
  - rt dm\_task\_set\_period, 183
  - rt dm\_task\_set\_priority, 183
  - rt dm\_task\_should\_stop, 184
  - rt dm\_task\_sleep, 184
  - rt dm\_task\_sleep\_abs, 185
  - rt dm\_task\_sleep\_until, 185
  - rt dm\_task\_unblock, 186
  - rt dm\_task\_wait\_period, 186
  - rt dm\_wait\_complete, 187
  - rt dm\_wait\_is\_completed, 187
  - rt dm\_wait\_prepare, 188
- Testing Devices, 78
- Thread information flags, 46
- Thread management, 386
  - pthread\_create, 386
  - pthread\_join, 388
  - pthread\_kill, 388
  - pthread\_setmode\_np, 389
  - pthread\_setname\_np, 391
- Thread runtime statistics, 42
- Thread scheduling control, 246
  - xnsched\_rotate, 246
  - xnsched\_run, 247
- Thread services, 262
  - xnthread\_cancel, 263
  - xnthread\_current, 264
  - xnthread\_from\_task, 264
  - xnthread\_harden, 265
  - xnthread\_init, 265
  - xnthread\_join, 267
  - xnthread\_map, 267
  - xnthread\_relax, 268
  - xnthread\_resume, 269
  - xnthread\_set\_mode, 269
  - xnthread\_set\_periodic, 270
  - xnthread\_set\_schedparam, 271
  - xnthread\_set\_slice, 272

- xnthread\_start, 273
- xnthread\_suspend, 274
- xnthread\_test\_cancel, 275
- xnthread\_unblock, 275
- xnthread\_wait\_period, 276
- Thread state flags, 43
  - XNHELD, 44
  - XNMIGRATE, 44
  - XNPEND, 44
  - XNREADY, 44
  - XNSUSP, 45
  - XNTRAPLB, 45
- Thread synchronization services, 253
  - xnsynch\_acquire, 254
  - xnsynch\_destroy, 254
  - xnsynch\_flush, 255
  - xnsynch\_init, 256
  - xnsynch\_init\_protect, 257
  - xnsynch\_peek\_pendq, 257
  - xnsynch\_release, 258
  - xnsynch\_sleep\_on, 259
  - xnsynch\_try\_acquire, 259
  - xnsynch\_wakeup\_one\_sleeper, 260
  - xnsynch\_wakeup\_this\_sleeper, 261
- Timer management services, 561
  - RT\_TIMER\_INFO, 562
  - rt\_timer\_inquire, 562
  - rt\_timer\_ns2ticks, 562
  - rt\_timer\_read, 563
  - rt\_timer\_spin, 563
  - rt\_timer\_ticks2ns, 564
- Timer Services, 189
  - rtdm\_timer\_destroy, 190
  - rtdm\_timer\_handler\_t, 190
  - rtdm\_timer\_init, 191
  - rtdm\_timer\_mode, 190
  - rtdm\_timer\_start, 191
  - rtdm\_timer\_start\_in\_handler, 192
  - rtdm\_timer\_stop, 192
  - rtdm\_timer\_stop\_in\_handler, 193
- Timer services, 278
  - \_\_xntimer\_migrate, 279
  - program\_htick\_shot, 279
  - switch\_htick\_mode, 280
  - xntimer\_destroy, 281
  - xntimer\_get\_date, 281
  - xntimer\_get\_overruns, 282
  - xntimer\_get\_timeout, 282
  - xntimer\_grab\_hardware, 283
  - xntimer\_init, 283
  - xntimer\_interval, 284
  - xntimer\_release\_hardware, 285
  - xntimer\_start, 285
  - xntimer\_stop, 286
- timer\_create
  - Clocks and timers, 322
- timer\_delete
  - Clocks and timers, 323
- timer\_getoverrun
  - Clocks and timers, 324
- timer\_gettime
  - Clocks and timers, 325
- timer\_settime
  - Clocks and timers, 326
- Transition Kit, 567
  - COMPAT\_rt\_alarm\_create, 567
  - COMPAT\_rt\_event\_clear, 568
  - COMPAT\_rt\_event\_create, 569
  - COMPAT\_rt\_event\_signal, 570
  - COMPAT\_rt\_pipe\_create, 571
  - COMPAT\_rt\_task\_create, 572
  - COMPAT\_rt\_task\_set\_periodic, 573
  - pthread\_make\_periodic\_np, 575
  - pthread\_wait\_np, 576
  - rt\_alarm\_wait, 577
- type
  - udd\_memregion, 624
- UDD\_IRQ\_CUSTOM
  - User-space driver core, 35
- UDD\_IRQ\_NONE
  - User-space driver core, 36
- UDD\_MEM\_LOGICAL
  - User-space driver core, 36
- UDD\_MEM\_NONE
  - User-space driver core, 36
- UDD\_MEM\_PHYS
  - User-space driver core, 36
- UDD\_MEM\_VIRTUAL
  - User-space driver core, 36
- UDD\_RTIOC\_IRQDIS
  - User-space driver core, 37
- UDD\_RTIOC\_IRQEN
  - User-space driver core, 37
- UDD\_RTIOC\_IRQSIG
  - User-space driver core, 37
- uapi/analog.h
  - A4L\_RNG\_FACTOR, 697
- udd\_device, 619
  - close, 620
  - device\_flags, 620
  - device\_subclass, 620
  - interrupt, 621
  - ioctl, 621
  - irq, 621
  - mem\_regions, 622
  - mmap, 622
  - open, 622
- udd\_device::udd\_reserved, 624
- udd\_disable\_irq
  - User-space driver core, 38
- udd\_enable\_irq
  - User-space driver core, 38
- udd\_get\_device
  - User-space driver core, 39
- udd\_memregion, 623
  - addr, 624

- len, [624](#)
- type, [624](#)
- udd\_notify\_event
  - User-space driver core, [39](#)
- udd\_register\_device
  - User-space driver core, [40](#)
- udd\_signotify, [625](#)
  - pid, [625](#)
  - sig, [625](#)
- udd\_unregister\_device
  - User-space driver core, [41](#)
- usablemem
  - RT\_HEAP\_INFO, [598](#)
- usedmem
  - RT\_HEAP\_INFO, [599](#)
- User-space driver core, [34](#)
  - UDD\_IRQ\_CUSTOM, [35](#)
  - UDD\_IRQ\_NONE, [36](#)
  - UDD\_MEM\_LOGICAL, [36](#)
  - UDD\_MEM\_NONE, [36](#)
  - UDD\_MEM\_PHYS, [36](#)
  - UDD\_MEM\_VIRTUAL, [36](#)
  - UDD\_RTIOC\_IRQDIS, [37](#)
  - UDD\_RTIOC\_IRQEN, [37](#)
  - UDD\_RTIOC\_IRQSIG, [37](#)
  - udd\_disable\_irq, [38](#)
  - udd\_enable\_irq, [38](#)
  - udd\_get\_device, [39](#)
  - udd\_notify\_event, [39](#)
  - udd\_register\_device, [40](#)
  - udd\_unregister\_device, [41](#)
- Utility Services, [231](#)
  - rtdm\_available, [232](#)
  - rtdm\_copy\_from\_user, [232](#)
  - rtdm\_copy\_to\_user, [233](#)
  - rtdm\_free, [234](#)
  - rtdm\_in\_rt\_context, [234](#)
  - rtdm\_iomap\_to\_user, [234](#)
  - rtdm\_malloc, [235](#)
  - rtdm\_mmap\_iomem, [236](#)
  - rtdm\_mmap\_kmem, [237](#)
  - rtdm\_mmap\_to\_user, [237](#)
  - rtdm\_mmap\_vmem, [238](#)
  - rtdm\_munmap, [239](#)
  - rtdm\_printk, [239](#)
  - rtdm\_printk\_ratelimited, [240](#)
  - rtdm\_read\_user\_ok, [240](#)
  - rtdm\_rt\_capable, [241](#)
  - rtdm\_rw\_user\_ok, [242](#)
  - rtdm\_safe\_copy\_from\_user, [242](#)
  - rtdm\_safe\_copy\_to\_user, [243](#)
  - rtdm\_strncpy\_from\_user, [243](#)
- vfile
  - xnvfile\_regular\_iterator, [630](#)
  - xnvfile\_snapshot\_iterator, [637](#)
- Virtual file services, [287](#)
  - cobalt\_vfroot, [294](#)
  - xnvfile\_destroy, [289](#)
  - xnvfile\_get\_blob, [289](#)
  - xnvfile\_get\_integer, [290](#)
  - xnvfile\_get\_string, [290](#)
  - xnvfile\_init\_dir, [291](#)
  - xnvfile\_init\_link, [291](#)
  - xnvfile\_init\_regular, [292](#)
  - xnvfile\_init\_snapshot, [293](#)
- VxWorks® emulator, [565](#)
- write\_nrt
  - rtdm\_fd\_ops, [612](#)
- write\_rt
  - rtdm\_fd\_ops, [612](#)
- XDDP\_BUFSZ
  - Real-time IPC, [85](#)
- XDDP\_EVTDOWN
  - Real-time IPC, [86](#)
- XDDP\_EVTIN
  - Real-time IPC, [86](#)
- XDDP\_EVTNOBUF
  - Real-time IPC, [86](#)
- XDDP\_EVTOUT
  - Real-time IPC, [87](#)
- XDDP\_LABEL
  - Real-time IPC, [87](#)
- XDDP\_MONITOR
  - Real-time IPC, [88](#)
- XDDP\_POOLSZ
  - Real-time IPC, [88](#)
- XNHELD
  - Thread state flags, [44](#)
- XNMIGRATE
  - Thread state flags, [44](#)
- XNPEND
  - Thread state flags, [44](#)
- XNREADY
  - Thread state flags, [44](#)
- XNSUSP
  - Thread state flags, [45](#)
- XNTRAPLB
  - Thread state flags, [45](#)
- xnapc\_alloc
  - Asynchronous Procedure Calls, [98](#)
- xnapc\_free
  - Asynchronous Procedure Calls, [99](#)
- xnapc\_schedule
  - Asynchronous Procedure Calls, [100](#)
- xnarch\_generic\_full\_divmod64
  - In-kernel arithmetics, [101](#)
- xnbufd\_copy\_from\_kmem
  - Buffer descriptor, [104](#)
- xnbufd\_copy\_to\_kmem
  - Buffer descriptor, [105](#)
- xnbufd\_invalidate
  - Buffer descriptor, [106](#)
- xnbufd\_map\_kread
  - Buffer descriptor, [107](#)
- xnbufd\_map\_kwrite

- Buffer descriptor, [107](#)
- xnbufd\_map\_uread
  - Buffer descriptor, [108](#)
- xnbufd\_map\_uwrite
  - Buffer descriptor, [108](#)
- xnbufd\_reset
  - Buffer descriptor, [109](#)
- xnbufd\_unmap\_kread
  - Buffer descriptor, [109](#)
- xnbufd\_unmap\_kwrite
  - Buffer descriptor, [109](#)
- xnbufd\_unmap\_uread
  - Buffer descriptor, [110](#)
- xnbufd\_unmap\_uwrite
  - Buffer descriptor, [110](#)
- xnclock\_adjust
  - Clock services, [112](#)
- xnclock\_deregister
  - Clock services, [113](#)
- xnclock\_register
  - Clock services, [113](#)
- xnclock\_tick
  - Clock services, [114](#)
- xnheap\_alloc
  - Dynamic memory allocation services, [116](#)
- xnheap\_destroy
  - Dynamic memory allocation services, [117](#)
- xnheap\_free
  - Dynamic memory allocation services, [117](#)
- xnheap\_set\_name
  - Dynamic memory allocation services, [118](#)
- xnintr\_affinity
  - Interrupt management, [123](#)
- xnintr\_attach
  - Interrupt management, [125](#)
- xnintr\_destory
  - Interrupt management, [125](#)
- xnintr\_detach
  - Interrupt management, [126](#)
- xnintr\_disable
  - Interrupt management, [126](#)
- xnintr\_enable
  - Interrupt management, [127](#)
- xnintr\_init
  - Interrupt management, [127](#)
- xnmap\_create
  - Lightweight key-to-object mapping service, [133](#)
- xnmap\_delete
  - Lightweight key-to-object mapping service, [133](#)
- xnmap\_enter
  - Lightweight key-to-object mapping service, [134](#)
- xnmap\_fetch
  - Lightweight key-to-object mapping service, [134](#)
- xnmap\_fetch\_nocheck
  - Lightweight key-to-object mapping service, [135](#)
- xnmap\_remove
  - Lightweight key-to-object mapping service, [136](#)
- xnregistry\_bind
  - Registry services, [137](#)
- xnregistry\_enter
  - Registry services, [138](#)
- xnregistry\_lookup
  - Registry services, [139](#)
- xnregistry\_remove
  - Registry services, [140](#)
- xnregistry\_unlink
  - Registry services, [141](#)
- xnsched, [626](#)
  - cpu, [626](#)
  - curr, [626](#)
  - htimer, [627](#)
  - inesting, [627](#)
  - lflags, [627](#)
  - resched, [627](#)
  - rrbtimer, [627](#)
  - rt, [627](#)
  - status, [627](#)
- xnsched\_rotate
  - Thread scheduling control, [246](#)
- xnsched\_run
  - Thread scheduling control, [247](#)
- xnselect
  - Synchronous I/O multiplexing, [249](#)
- xnselect\_bind
  - Synchronous I/O multiplexing, [250](#)
- xnselect\_destroy
  - Synchronous I/O multiplexing, [250](#)
- xnselect\_init
  - Synchronous I/O multiplexing, [251](#)
- xnselect\_signal
  - Synchronous I/O multiplexing, [251](#)
- xnselector\_destroy
  - Synchronous I/O multiplexing, [252](#)
- xnselector\_init
  - Synchronous I/O multiplexing, [252](#)
- xnsynch\_acquire
  - Thread synchronization services, [254](#)
- xnsynch\_destroy
  - Thread synchronization services, [254](#)
- xnsynch\_flush
  - Thread synchronization services, [255](#)
- xnsynch\_init
  - Thread synchronization services, [256](#)
- xnsynch\_init\_protect
  - Thread synchronization services, [257](#)
- xnsynch\_peek\_pendq
  - Thread synchronization services, [257](#)
- xnsynch\_release
  - Thread synchronization services, [258](#)
- xnsynch\_sleep\_on

- Thread synchronization services, [259](#)
- xnsynch\_try\_acquire
  - Thread synchronization services, [259](#)
- xnsynch\_wakeup\_one\_sleeper
  - Thread synchronization services, [260](#)
- xnsynch\_wakeup\_this\_sleeper
  - Thread synchronization services, [261](#)
- xnthread\_cancel
  - Thread services, [263](#)
- xnthread\_current
  - Thread services, [264](#)
- xnthread\_from\_task
  - Thread services, [264](#)
- xnthread\_harden
  - Thread services, [265](#)
- xnthread\_init
  - Thread services, [265](#)
- xnthread\_join
  - Thread services, [267](#)
- xnthread\_map
  - Thread services, [267](#)
- xnthread\_relax
  - Thread services, [268](#)
- xnthread\_resume
  - Thread services, [269](#)
- xnthread\_set\_mode
  - Thread services, [269](#)
- xnthread\_set\_periodic
  - Thread services, [270](#)
- xnthread\_set\_schedparam
  - Thread services, [271](#)
- xnthread\_set\_slice
  - Thread services, [272](#)
- xnthread\_start
  - Thread services, [273](#)
- xnthread\_suspend
  - Thread services, [274](#)
- xnthread\_test\_cancel
  - Thread services, [275](#)
- xnthread\_unblock
  - Thread services, [275](#)
- xnthread\_wait\_period
  - Thread services, [276](#)
- xntimer\_destroy
  - Timer services, [281](#)
- xntimer\_get\_date
  - Timer services, [281](#)
- xntimer\_get\_overruns
  - Timer services, [282](#)
- xntimer\_get\_timeout
  - Timer services, [282](#)
- xntimer\_grab\_hardware
  - Timer services, [283](#)
- xntimer\_init
  - Timer services, [283](#)
- xntimer\_interval
  - Timer services, [284](#)
- xntimer\_release\_hardware
  - Timer services, [285](#)
- xntimer\_start
  - Timer services, [285](#)
- xntimer\_stop
  - Timer services, [286](#)
- xnvfile\_destroy
  - Virtual file services, [289](#)
- xnvfile\_get\_blob
  - Virtual file services, [289](#)
- xnvfile\_get\_integer
  - Virtual file services, [290](#)
- xnvfile\_get\_string
  - Virtual file services, [290](#)
- xnvfile\_init\_dir
  - Virtual file services, [291](#)
- xnvfile\_init\_link
  - Virtual file services, [291](#)
- xnvfile\_init\_regular
  - Virtual file services, [292](#)
- xnvfile\_init\_snapshot
  - Virtual file services, [293](#)
- xnvfile\_lock\_ops, [628](#)
  - get, [628](#)
  - put, [629](#)
- xnvfile\_regular\_iterator, [629](#)
  - pos, [629](#)
  - private, [630](#)
  - seq, [630](#)
  - vfile, [630](#)
- xnvfile\_regular\_ops, [630](#)
  - begin, [631](#)
  - end, [631](#)
  - next, [632](#)
  - rewind, [632](#)
  - show, [633](#)
  - store, [633](#)
- xnvfile\_rev\_tag, [634](#)
  - rev, [634](#)
- xnvfile\_snapshot, [635](#)
- xnvfile\_snapshot\_iterator, [635](#)
  - databuf, [636](#)
  - endfn, [637](#)
  - nrddata, [637](#)
  - private, [637](#)
  - seq, [637](#)
  - vfile, [637](#)
- xnvfile\_snapshot\_ops, [638](#)
  - begin, [638](#)
  - end, [639](#)
  - next, [639](#)
  - rewind, [640](#)
  - show, [640](#)
  - store, [641](#)