

XtratuM Hypervisor for LEON3

Volume 4: Reference Manual

Miguel Masmano, Javier O. Coronel, Alfons Crespo,
Patricia Balbastre

March, 2011
Reference: XM-3-reference-023d

This page is intentionally left blank.

DOCUMENT CONTROL PAGE

TITLE: XtratuM Hypervisor for LEON3 : Volume 4: Reference Manual

AUTHOR/S: Miguel Masmano, Javier O. Coronel, Alfons Crespo, Patricia Balbastre

LAST PAGE NUMBER: 91

VERSION OF SOURCE CODE: XtratuM 3 for LEON3 ()

REFERENCE ID: XM-3-reference-023d

SUMMARY: This document contains the set of manual pages of XtratuM.

DISCLAIMER: This documentation is currently under active development. Therefore, no explicit or implied warranties in respect of any properties, including, but not limited to, correctness and fitness for purpose. Contributions of material, suggestions and corrections are welcome.

REFERENCING THIS DOCUMENT:

```
@techreport {XM-3-reference-023d,
  title = {XtratuM Hypervisor for LEON3 : Volume 4: Reference Manual},
  authors = { Miguel Masmano & Javier O. Coronel & Alfons Crespo & Patricia
    Balbastre },
  institution = {Universidad polit cnica de Valencia},
  number = {XM-3-reference-023d},
  date={March, 2011},
}
```

Copyright   March, 2011 Miguel Masmano and Alfons Crespo

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

Changes:

Version	Date	Comments
0.1	February, 2010	[xm-3-reference-023] Code version 3.1.0.

Version	Date	Comments
0.2	May 2010	<p>[xm-3-reference-023b] Code version 3.1.2</p> <p>Renamed:</p> <ul style="list-style-type: none"> • <code>XM_system_get_status()</code> → <code>XM_get_system_status()</code>. • <code>XM_partition_get_status()</code> → <code>XM_get_partition_status()</code>. <p>Changed:</p> <ul style="list-style-type: none"> • <code>XM_unmask_irq()</code> → <code>XM_clear_irqmask()</code> • <code>XM_mask_irq()</code> → <code>XM_set_irqmask()</code> • <code>XM_enable_irqs()</code> → <code>XM_sparcv8_set_pil()</code> • <code>XM_disable_irqs()</code> → <code>XM_sparcv8_clear_pil()</code> <p>Added:</p> <ul style="list-style-type: none"> • <code>XM_clear_irqpend()</code> • <code>XM_set_irqpend()</code> • <code>XM_get_partition_mmap()</code> • <code>XM_get_sampling_port_info()</code> • <code>XM_get_queuing_port_info()</code> • <code>XM_sparcv8_set_psr()</code> • <code>XM_sparcv8_get_psr()</code> • <code>XM_sparcv8_flush_tlb()</code> • <code>XM_sparcv8_flush_cache()</code> • <code>XM_get_partition_mmap()</code> • <code>XM_set_plan()</code> • <code>XM_get_plan_status()</code> <p>Removed:</p> <ul style="list-style-type: none"> • <code>XM_get_physmem_map()</code>
0.3	October, 2010	<p>[xm-3-reference-023c] Code version 3.2.0.</p> <p>Added:</p> <ul style="list-style-type: none"> • <code>XM_update_page32()</code> • <code>XM_set_page_type()</code>
0.4	March, 2011	<p>[xm-3-reference-023d] Code version 3.2.0. Updates the definition of the hypercalls.</p> <p>Renamed:</p> <ul style="list-style-type: none"> • <code>XM_sparcv8*()</code> → <code>XM_sparc*()</code>.

Contents

Preface	ix
1 Introduction	1
1.1 XtratuM overview	2
1.2 Health Monitoring overview	6
2 Hypercalls	11
2.1 XM_clear_irqmask	12
2.2 XM_clear_irqpend	13
2.3 XM_create_queuing_port	14
2.4 XM_create_sampling_port	15
2.5 XM_disable_irqs	16
2.6 XM_enable_irqs	17
2.7 XM_get_partition_mmap	18
2.8 XM_get_partition_status	19
2.9 XM_get_plan_status	21
2.10 XM_get_queuing_port_info	22
2.11 XM_get_queuing_port_status	23
2.12 XM_get_sampling_port_info	24
2.13 XM_get_sampling_port_status	25
2.14 XM_get_system_status	26
2.15 XM_get_time	27
2.16 XM_halt_partition	28
2.17 XM_halt_system	29
2.18 XM_hm_open	30
2.19 XM_hm_read	31
2.20 XM_hm_seek	32
2.21 XM_hm_status	34
2.22 XM_idle_self	35

2.23 XM_mask_irq	37
2.24 XM_memory_copy	38
2.25 XM_multicall	40
2.26 XM_params.get_PCT	41
2.27 XM_read_console	42
2.28 XM_read_sampling_message	43
2.29 XM_receive_queuing_message	45
2.30 XM_request_irq	47
2.31 XM_reset_partition	48
2.32 XM_reset_system	50
2.33 XM_resume_partition	51
2.34 XM_route_irq	52
2.35 XM_send_queuing_message	53
2.36 XM_set_irqmask	55
2.37 XM_set_irqpend	56
2.38 XM_set_page_type	57
2.39 XM_set_plan	59
2.40 XM_set_timer	60
2.41 XM_shutdown_partition	62
2.42 XM_sparc_atomic_add	63
2.43 XM_sparc_atomic_and	64
2.44 XM_sparc_atomic_or	65
2.45 XM_sparc_clear_pil	66
2.46 XM_sparc_flush_cache	67
2.47 XM_sparc_flush_regwin	68
2.48 XM_sparc_flush_tlb	69
2.49 XM_sparc_get_flags	70
2.50 XM_sparc_get_psr	71
2.51 XM_sparc_inport	72
2.52 XM_sparc_iret	73
2.53 XM_sparc_outport	74
2.54 XM_sparc_set_flags	76
2.55 XM_sparc_set_pil	77
2.56 XM_sparc_set_psr	78
2.57 XM_suspend_partition	79
2.58 XM_trace_event	80
2.59 XM_trace_open	82

2.60 XM_trace_read	83
2.61 XM_trace_seek	84
2.62 XM_trace_status	85
2.63 XM_unmask_irq	86
2.64 XM_update_page32	87
2.65 XM_write_console	89
2.66 XM_write_register32	90
2.67 XM_write_sampling_message	91

This page is intentionally left blank.

Preface

The audience for this document is software developers that have to use directly the services of XtratuM. The reader is expected to have strong knowledge of the LEON3 (SPARC v8) architecture and experience in programming device drivers. It is also advisable to have some knowledge of the ARINC-653 and related standards.

Typographical conventions

The following font conventions are used in this document:

- **typewriter**: used in assembly and C code examples, and to show the output of commands.
- *italic*: used to introduce new terms.
- **bold face**: used to emphasize or highlight a word or paragraph.

Caution sign

The caution sign stresses information that is critical to the integrity or continuity of the system.



Support

Alfons Crespo
Universidad Politcnica de Valencia
Instituto de Automtica e Informtica Industrial
Camino de vera s/n
CP: 46022
Valencia, Spain

The official XtratuM web site is: <http://www.xtratum.org>

Acknowledgements

Porting to LEON3 with MMU has been done in the frame of the “ESA Project AO5829 Securely Partitioning Spacecraft Computing Resources” led by Astrium EADS.

This page is intentionally left blank.

Chapter 1

Introduction

1.1 XtratuM overview

Synopsis: Global data structures and types.

Global data types:

```
// Basic types
typedef unsigned char xm_u8_t;
typedef char xm_s8_t;
typedef unsigned short xm_u16_t;
typedef short xm_s16_t;
typedef unsigned int xm_u32_t;
typedef int xm_s32_t;
typedef unsigned long long xm_u64_t;
typedef long long xm_s64_t;
```

Listing 1.1: /core/include/sparcv8/arch.types.h

```
// Extended types
typedef xm_s64_t xmTime_t;
typedef xm_u32_t xmAddress_t;
typedef xmAddress_t xmIoAddress_t;
typedef xm_u32_t xmSize_t;
typedef xm_s32_t xmSSize_t;
typedef xm_u32_t xmId_t;
```

Listing 1.2: /core/include/sparcv8/arch.types.h

API and ABI version numbers:

```
#define XM_ABI_VERSION 3
#define XM_ABI_SUBVERSION 1
#define XM_ABI_REVISION 0

#define XM_API_VERSION 3
#define XM_API_SUBVERSION 1
#define XM_API_REVISION 2
```

Listing 1.3: /core/include/hypercalls.h

Error codes:

```
#define XM_OK (0)
#define XM_UNKNOWN_HYPERCALL (-1)
#define XM_INVALID_PARAM (-2)
#define XM_PERM_ERROR (-3)
#define XM_INVALID_CONFIG (-4)
#define XM_INVALID_MODE (-5)
#define XM_NOT_AVAILABLE (-6)
#define XM_OP_NOT_ALLOWED (-7)
#define XM_MULTICALL_ERROR (-8)
```

Listing 1.4: /core/include/hypercalls.h

Interrupts:

Native hardware interrupts (interrupts generated by the real hardware) are received by the partition in the same way than in the host machine. The first native interrupt causes the trap number 0x11. The last interrupt (15) causes the trap 0x1F. Although the SPARC v8 architecture defines only 15 interrupt numbers (levels), XtratuM reserves 32 for compatibility with other processors.

Extended interrupts are raised in response to XtratuM events. The new interrupts set are in the range from 32 to 63, and causes traps from 256 to 287 respectively.

Below the list of hardware and extended interrupts (terrupts available in the AT697 board and tsim simulator):

```
#define XM_VT_HW_FIRST          (0)
#define XM_VT_HW_LAST          (31)
#define XM_VT_HW_MAX           (32)

#define XM_VT_HW_INTERNAL_BUS_TRAP_NR (1+XM_VT_HW_FIRST)
#define XM_VT_HW_UART2_TRAP_NR      (2+XM_VT_HW_FIRST)
#define XM_VT_HW_UART1_TRAP_NR      (3+XM_VT_HW_FIRST)
#define XM_VT_HW_IO_IRQ0_TRAP_NR    (4+XM_VT_HW_FIRST)
#define XM_VT_HW_IO_IRQ1_TRAP_NR    (5+XM_VT_HW_FIRST)
#define XM_VT_HW_IO_IRQ2_TRAP_NR    (6+XM_VT_HW_FIRST)
#define XM_VT_HW_IO_IRQ3_TRAP_NR    (7+XM_VT_HW_FIRST)
#define XM_VT_HW_TIMER1_TRAP_NR     (8+XM_VT_HW_FIRST)
#define XM_VT_HW_TIMER2_TRAP_NR     (9+XM_VT_HW_FIRST)
#define XM_VT_HW_DSU_TRAP_NR        (11+XM_VT_HW_FIRST)
#define XM_VT_HW_PCI_TRAP_NR        (14+XM_VT_HW_FIRST)

#define XM_VT_EXT_FIRST         (0)
#define XM_VT_EXT_LAST         (31)
#define XM_VT_EXT_MAX          (32)

#define XM_VT_EXT_HW_TIMER      (0+XM_VT_EXT_FIRST)
#define XM_VT_EXT_EXEC_TIMER    (1+XM_VT_EXT_FIRST)
#define XM_VT_EXT_WATCHDOG_TIMER (2+XM_VT_EXT_FIRST)
#define XM_VT_EXT_SHUTDOWN      (3+XM_VT_EXT_FIRST)
#define XM_VT_EXT_OBJDESC       (4+XM_VT_EXT_FIRST)

#define XM_VT_EXT_CYCLIC_SLOT_START (8+XM_VT_EXT_FIRST)

#define XM_VT_EXT_MEM_PROTECT    (16+XM_VT_EXT_FIRST)

/* Inter-Partition Virtual Interrupts */
#define XM_MAX_IPVI 8
#define XM_VT_EXT_IPVIO      (24+XM_VT_EXT_FIRST)
#define XM_VT_EXT_IPVI1     (25+XM_VT_EXT_FIRST)
#define XM_VT_EXT_IPVI2     (26+XM_VT_EXT_FIRST)
#define XM_VT_EXT_IPVI3     (27+XM_VT_EXT_FIRST)
#define XM_VT_EXT_IPVI4     (28+XM_VT_EXT_FIRST)
#define XM_VT_EXT_IPVI5     (29+XM_VT_EXT_FIRST)
#define XM_VT_EXT_IPVI6     (30+XM_VT_EXT_FIRST)
#define XM_VT_EXT_IPVI7     (31+XM_VT_EXT_FIRST)
```

Listing 1.5: /core/include/guest.h

The ”_VT_” prefix stands for **Virtual Trap**.

Note that the trap table of a partition shall contain 32 more entries (to accommodate the extended interrupt set) than the original SPARC v8 trap table.

Exceptions:

List of exceptions (traps triggered by the processor):

```
#define DATA_STORE_ERROR 0x2b // 0
#define INSTRUCTION_ACCESS_MMU_MISS 0x3c // 1
#define INSTRUCTION_ACCESS_ERROR 0x21 // 2
#define R_REGISTER_ACCESS_ERROR 0x20 // 3
#define INSTRUCTION_ACCESS_EXCEPTION 0x1 // 4
#define PRIVILEGED_INSTRUCTION 0x03 // 5
#define ILLEGAL_INSTRUCTION 0x2 // 6
#define FP_DISABLED 0x4 // 7
#define CP_DISABLED 0x24 // 8
#define UNIMPLEMENTED_FLUSH 0x25 // 9
#define WATCHPOINT_DETECTED 0xb // 10
// #define WINDOW_OVERFLOW 0x5
// #define WINDOW_UNDERFLOW 0x6
#define MEM_ADDRESS_NOT_ALIGNED 0x7 // 11
#define FP_EXCEPTION 0x8 // 12
#define CP_EXCEPTION 0x28 // 13
#define DATA_ACCESS_ERROR 0x29 // 14
#define DATA_ACCESS_MMU_MISS 0x2c // 15
#define DATA_ACCESS_EXCEPTION 0x9 // 16
#define TAG_OVERFLOW 0xa // 17
#define DIVISION_BY_ZERO 0x2a // 18
```

Listing 1.6: /core/include/sparcv8/irqs.h

Processor exceptions are managed through the health monitoring system.

Note that WINDOW_OVERFLOW and WINDOW_UNDERFLOW exceptions are automatically managed by XtratuM and hidden to the partition.

Partition Control Table (PCT):

```
typedef struct {
    xm_u32_t magic;
    xm_u32_t xmVersion; // XM version
    xm_u32_t xmAbiVersion; // XM's abi version
    xm_u32_t xmApiVersion; // XM's api version
    xm_u32_t resetCounter;
    xm_u32_t resetStatus;
    xm_u32_t cpuKhz;
    xmId_t id;
    xm_u32_t flags;
    xm_u32_t hwIrqs; // Hw interrupts belonging to the partition
    xm_s32_t noPhysicalMemAreas;
    xm_u8_t name[CONFIG_ID_STRING_LENGTH];
    xm_u32_t iFlags;
    xm_u32_t hwIrqsPend; // pending hw irqs
    xm_u32_t hwIrqsMask; // masked hw irqs

    xm_u32_t extIrqsPend; // pending extended irqs
    xm_u32_t extIrqsMask; // masked extended irqs
}
```

```
struct pctArch arch;
struct {
    xm_u32_t noSlot:16, reserved:16;
    xm_u32_t id;
    xm_u32_t slotDuration;
} schedInfo;
xm_u16_t trap2Vector[NO_TRAPS];
xm_u16_t hwIrq2Vector[CONFIG_NO_HWIRQS];
xm_u16_t extIrq2Vector[XM_VT_EXT_MAX];
} partitionControlTable_t;
```

Listing 1.7: /core/include/guest.h

```
struct pctArch {
    xmAddress_t tbr;
#ifdef CONFIG_MMU
    xmAddress_t ptdL1;
    xm_u32_t faultStatusReg;
    xm_u32_t faultAddressReg;
#endif
};
```

Listing 1.8: /core/include/sparcv8/guest.h

1.2 Health Monitoring overview

Synopsis: Health Monitoring overview.

Health monitoring events:

There are tree groups of hm events:

Processor triggered

The processor traps caused by an incorrect behavior of the processor or the board are managed as health monitoring events.

```
#define XM_HM_EV_WRITE_ERROR (XM_HM_MAX_GENERIC_EVENTS+0)
#define XM_HM_EV_INSTR_ACCESS_MMU_MISS (XM_HM_MAX_GENERIC_EVENTS+1)
#define XM_HM_EV_INSTR_ACCESS_ERROR (XM_HM_MAX_GENERIC_EVENTS+2)
#define XM_HM_EV_REGISTER_HARDWARE_ERROR (XM_HM_MAX_GENERIC_EVENTS+3)
#define XM_HM_EV_INSTR_ACCESS_EXCEPTION (XM_HM_MAX_GENERIC_EVENTS+4)
#define XM_HM_EV_PRIVILEGED_INSTR (XM_HM_MAX_GENERIC_EVENTS+5)
#define XM_HM_EV_ILLEGAL_INSTR (XM_HM_MAX_GENERIC_EVENTS+6)
#define XM_HM_EV_FP_DISABLED (XM_HM_MAX_GENERIC_EVENTS+7)
#define XM_HM_EV_CP_DISABLED (XM_HM_MAX_GENERIC_EVENTS+8)
#define XM_HM_EV_UNIMPLEMENTED_FLUSH (XM_HM_MAX_GENERIC_EVENTS+9)
#define XM_HM_EV_WATCHPOINT_DETECTED (XM_HM_MAX_GENERIC_EVENTS+10)
#define XM_HM_EV_MEM_ADDR_NOT_ALIGNED (XM_HM_MAX_GENERIC_EVENTS+11)
#define XM_HM_EV_FP_EXCEPTION (XM_HM_MAX_GENERIC_EVENTS+12)
#define XM_HM_EV_CP_EXCEPTION (XM_HM_MAX_GENERIC_EVENTS+13)
#define XM_HM_EV_DATA_ACCESS_ERROR (XM_HM_MAX_GENERIC_EVENTS+14)
#define XM_HM_EV_DATA_ACCESS_MMU_MISS (XM_HM_MAX_GENERIC_EVENTS+15)
#define XM_HM_EV_DATA_ACCESS_EXCEPTION (XM_HM_MAX_GENERIC_EVENTS+16)
#define XM_HM_EV_TAG_OVERFLOW (XM_HM_MAX_GENERIC_EVENTS+17)
#define XM_HM_EV_DIVIDE_EXCEPTION (XM_HM_MAX_GENERIC_EVENTS+18)
```

Listing 1.9: /core/include/sparcv8/xmconf.h

User triggered

A trace message with the criticality level equal to XM_TRACE_UNRECOVERABLE raises the XM_HM_EV_PARTITION_ERROR event.

XtratuM triggered

```
#define XM_HM_EV_INTERNAL_ERROR 0
#define XM_HM_EV_UNEXPECTED_TRAP 1
#define XM_HM_EV_PARTITION_UNRECOVERABLE 2
#define XM_HM_EV_PARTITION_ERROR 3
#define XM_HM_EV_PARTITION_INTEGRITY 4
#define XM_HM_EV_MEM_PROTECTION 5

#define XM_HM_EV_OVERRUN 6

#define XM_HM_EV_SCHED_ERROR 7
#define XM_HM_EV_WATCHDOG_TIMER 8

#define XM_HM_EV_INCOMPATIBLE_INTERFACE 9
```

Listing 1.10: /core/include/xmconf.h

Default actions:

The following tables summarises the list of health monitoring events. The default action (marked with a * symbol). The allowed actions (those that can be selected in the XM_CF file) are marked with the • symbol. The last column indicates if the event is logger or not (if it is nit explicitly configured in the XM_CF file).

If this HM event is propagated to the partition, the trap number raised is the same than the one of the native hardware, and not the HM event number. Note that **XM_HM_EV_*** constants do not match the trap numbers (see Listing 1.6).

	XM_HM_AC_COLD_RESET	XM_HM_AC_HALT	XM_HM_AC_IGNORE	XM_HM_AC_PROPAGATE	XM_HM_AC_SHUTDOWN	XM_HM_AC_SUSPEND	XM_HM_AC_SWITCH_TO_MAINTENANCE	XM_HM_AC_WARM_RESET	Log event
XM_HM_EV_CP_DISABLED	•	*	•	•	•	•	•	•	No
XM_HM_EV_CP_EXCEPTION	•	*	•	•	•	•	•	•	No
XM_HM_EV_DATA_ACCESS_ERROR	•	*	•	•	•	•	•	•	Yes
XM_HM_EV_DATA_ACCESS_EXCEPTION	•	*	•	•	•	•	•	•	Yes
XM_HM_EV_DATA_ACCESS_MMU_MISS	•	*	•	•	•	•	•	•	No
XM_HM_EV_DIVIDE_EXCEPTION	•	*	•	•	•	•	•	•	No
XM_HM_EV_FP_DISABLED	•	*	•	•	•	•	•	•	No
XM_HM_EV_FP_EXCEPTION	•	*	•	•	•	•	•	•	No
XM_HM_EV_ILLEGAL_INSTR	•	*	•	•	•	•	•	•	No
XM_HM_EV_INCOMPATIBLE_INTERFACE	•	•			*		•	•	Yes
XM_HM_EV_INSTR_ACCESS_ERROR	•	*	•	•	•	•	•	•	Yes
XM_HM_EV_INSTR_ACCESS_EXCEPTION	•	*	•	•	•	•	•	•	Yes
XM_HM_EV_INSTR_ACCESS_MMU_MISS	•	*	•	•	•	•	•	•	No
XM_HM_EV_INTERNAL_ERROR	•	*	•		•	•	•	•	Yes
XM_HM_EV_MEM_ADDR_NOT_ALIGNED	•	*	•	•	•	•	•	•	No
XM_HM_EV_MEM_PROTECTION	•	*	•		•	•	•	•	Yes
XM_HM_EV_OVERRUN	•	*	•		•	•	•	•	Yes
XM_HM_EV_PARTITION_ERROR	•	*	•		•	•	•	•	Yes
XM_HM_EV_PARTITION_INTEGRITY	•	*	•		•	•	•	•	Yes
XM_HM_EV_PARTITION_UNRECOVERABLE	•	*	•		•	•	•	•	Yes
XM_HM_EV_PRIVILEGED_INSTR	•	*	•	•	•	•	•	•	Yes
XM_HM_EV_REGISTER_HARDWARE_ERROR	•	*	•	•	•	•	•	•	Yes
XM_HM_EV_SCHED_ERROR	•	*	•		•	•	•	•	Yes
XM_HM_EV_TAG_OVERFLOW	•	*	•	•	•	•	•	•	No
XM_HM_EV_UNEXPECTED_TRAP	•	*	•		•	•	•	•	Yes
XM_HM_EV_UNIMPLEMENTED_FLUSH	•	*	•	•	•	•	•	•	No
XM_HM_EV_WATCHDOG_TIMER	•	*	•		•	•	•	•	Yes
XM_HM_EV_WATCHPOINT_DETECTED	•	*	•	•	•	•	•	•	No
XM_HM_EV_WRITE_ERROR	•	*	•	•	•	•	•	•	No

Table 1.1: Partition scope HM events

	XM_HM_AC_COLD_RESET	XM_HM_AC_HALT	XM_HM_AC_IGNORE	XM_HM_AC_PROPAGATE	XM_HM_AC_SHUTDOWN	XM_HM_AC_SUSPEND	XM_HM_AC_SWITCH_TO_MAINTENANCE	XM_HM_AC_WARM_RESET	Log event
XM_HM_EV_CP_DISABLED	•	*	•				•	•	Yes
XM_HM_EV_CP_EXCEPTION	•	*	•				•	•	Yes
XM_HM_EV_DATA_ACCESS_ERROR	•	*	•				•	•	Yes
XM_HM_EV_DATA_ACCESS_EXCEPTION	•	*	•				•	•	Yes
XM_HM_EV_DATA_ACCESS_MMU_MISS	•	*	•				•	•	Yes
XM_HM_EV_DIVIDE_EXCEPTION	•	*	•				•	•	Yes
XM_HM_EV_FP_DISABLED	•	*	•				•	•	Yes
XM_HM_EV_FP_EXCEPTION	•	*	•				•	•	Yes
XM_HM_EV_ILLEGAL_INSTR	•	*	•				•	•	Yes
XM_HM_EV_INCOMPATIBLE_INTERFACE	•	*			•		•	•	Yes
XM_HM_EV_INSTR_ACCESS_ERROR	•	*	•				•	•	Yes
XM_HM_EV_INSTR_ACCESS_EXCEPTION	•	*	•				•	•	Yes
XM_HM_EV_INSTR_ACCESS_MMU_MISS	•	*	•				•	•	Yes
XM_HM_EV_INTERNAL_ERROR	•	*	•				•	•	Yes
XM_HM_EV_MEM_ADDR_NOT_ALIGNED	•	*	•				•	•	Yes
XM_HM_EV_MEM_PROTECTION	•	*	•				•	•	Yes
XM_HM_EV_OVERRUN	•	*	•				•	•	Yes
XM_HM_EV_PARTITION_ERROR	•	*	•				•	•	Yes
XM_HM_EV_PARTITION_INTEGRITY	•	*	•				•	•	Yes
XM_HM_EV_PARTITION_UNRECOVERABLE	•	*	•				•	•	Yes
XM_HM_EV_PRIVILEGED_INSTR	•	*	•				•	•	Yes
XM_HM_EV_REGISTER_HARDWARE_ERROR	•	*	•				•	•	Yes
XM_HM_EV_SCHED_ERROR	•	*	•				•	•	Yes
XM_HM_EV_TAG_OVERFLOW	•	*	•				•	•	Yes
XM_HM_EV_UNEXPECTED_TRAP	•	*	•				•	•	Yes
XM_HM_EV_UNIMPLEMENTED_FLUSH	•	*	•				•	•	Yes
XM_HM_EV_WATCHDOG_TIMER	•	*	•		•	•	•	•	Yes
XM_HM_EV_WATCHPOINT_DETECTED	•	*	•				•	•	Yes
XM_HM_EV_WRITE_ERROR	•	*	•				•	•	Yes

Table 1.2: System scope HM events

Health monitoring actions:

```
#define XM_HM_AC_IGNORE 0
#define XM_HM_AC_SHUTDOWN 1
#define XM_HM_AC_COLD_RESET 2
#define XM_HM_AC_WARM_RESET 3
#define XM_HM_AC_SUSPEND 4
#define XM_HM_AC_HALT 5
#define XM_HM_AC_PROPAGATE 6
```

```
#define XM_HM_AC_SWITCH_TO_MAINTENANCE 7
```

Listing 1.11: /core/include/xmconf.h

[XM_HM_AC_IGNORE] No action is performed.

[XM_HM_AC_SHUTDOWN] The shutdown event is sent to the offending partition.

[XM_HM_AC_COLD_RESET] The offending partition or the hypervisor is cold reset.

[XM_HM_AC_WARM_RESET] The offending partition or the hypervisor is warm reset.

[XM_HM_AC_SUSPEND] The offending partition is suspended.

[XM_HM_AC_HALT] The offending partition or the hypervisor is halted.

[XM_HM_AC_PROPAGATE] The original trap is delivered to the faulting partition. This action can only be used in the partition's health monitoring tables.

[XM_HM_AC_SWITCH_TO_MAINTENANCE] A synchronous plan switch to maintenance plan (Plan 2) is performed. That is, the current plan is terminated and the plan 2 is immediately started.

Health monitoring logs:

The HM log entries are a "C" structure with the following fields:

```
typedef struct {
    xm_u32_t eventId:13, system:1, reserved:1, validCpuCtxt:1,
        moduleId:8, partitionId:8;
    struct cpuCtxt cpuCtxt;
#define HM_PAYLOAD_WORDS 5
    xm_u32_t word[HM_PAYLOAD_WORDS]; /* Payload */
    xmTime_t timeStamp;
} xmHmLog_t;
```

Listing 1.12: /core/include/objects/hm.h

See also:

[XM_hm_open](#) [Page: 30], [XM_hm_read](#) [Page: 31], [XM_hm_seek](#) [Page: 32], [XM_hm_status](#) [Page: 34]

This page is intentionally left blank.

Chapter 2

Hypercalls

2.1 XM_clear_irqmask

Synopsis: Unmask interrupts.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_clear_irqmask(xm_u32_t hwIrqsMask,  
                          xm_u32_t extIrqMask);
```

Description:

This function unmask (allows) hardware and extended interrupts. The interrupt lines set in the `hwIrqsMask` and `extIrqMask` parameters are unmasked, hardware and extended interrupts respectively.

The interrupt handlers of the unmasked interrupts will be invoked as soon as the interrupt is delivered to the partition.

Return value: This hypercall always succeeds: `XM_OK`.

Rationale:

This function replaces the deprecated `XM_unmask_irq()`.

History:

Added in version 3.1.2

Usage examples:

```
XM_clear_irqmask(0, (1<<XM_VT_EXT_HW_TIMER));  
//This call doesn't modify the hardware interrupts mask and  
  only modifies  
  //the extended mask in one bit, that is to say, it only  
  unmask the  
  //HW_TIMER interrupt.  
  
XM_clear_irqmask(0xFFFFFFFF, 0);  
//This call unmask all hardware interrupts and leaves the  
  extended mask as it is.
```

See also:

`XM_set_irqmask` [Page: 55], `XM_set_irqpend` [Page: 56], `XM_clear_irqpend` [Page: 13].

2.2 XM_clear_irqpend

Synopsis: Clear pending interrupts.

Category:

Standard service.

Declaration:

```
xm_u32_t XM_clear_irqpend(xm_u32_t hwIrqsPend,  
                          xm_u32_t extIrqPend);
```

Description:

This function clears pending interrupts. The interrupt lines which are set in the `hwIrqsPend` and `extIrqPend` parameters are cleared, clearing hardware and extended pending interrupts respectively.

Return value: This hypercall always succeeds: `XM_OK`.

History:

Added in version 3.1.2

Usage examples:

```
XM_clear_irqpend(0, (1<<XM_VT_EXT_HW_TIMER));
```

This call doesn't modify the hardware pending interrupts and only modified the extended pending interrupts in one bit, that is to say, it only clears the `HW_TIMER` interrupt if it was pending.

```
XM_clear_irqpend(0xFFFFFFFF, 0); //This call clears all hardware interrupts and leaves the extended pending interrupts //as they are.
```

See also:

[XM_set_irqpend \[Page: 56\]](#), [XM_set_irqmask \[Page: 55\]](#), [XM_clear_irqmask \[Page: 12\]](#).

2.3 XM_create_queuing_port

Synopsis: Create a queuing port.

Category:

Standard service.

Declaration:

```

xm_s32_t XM_create_queuing_port (char      *portName,
                                xm_u32_t  maxNoMsgs,
                                xm_u32_t  maxMsgSize,
                                xm_u32_t  direction);

```

Description:

The `XM_create_queuing_port()` service is used to create a queuing port. Only those queuing ports specified in the configuration file `XM_CF` can be created by the partition. Note that all the parameters must match the information contained in the configuration file. New ports cannot be created dynamically.

It is not an error to invoke this hypercall with an already created port. The same port descriptor is returned.

Return value:

Upon successful completion, `XM_create_queuing_port()` returns a non-negative integer representing the port descriptor. Otherwise, the function returns a negative value:

[`XM_INVALID_CONFIG`]

- The maximum number of queuing ports has been created.
- No queuing port of the partition is named `portName` in the configuration file.
- `maxMsgSize` is out of range or not compatible with the configuration.
- `direction` is invalid or not compatible with the configuration.

Usage examples:

```

xm_s32_t portDesc;

portDesc = XM_create_queuing_port ("example",
                                   30,
                                   XM_DESTINATION_PORT);

if (portDesc < 0) {
    XM_write_console("Port example could not be created\n",
                    ,34);
    XM_halt_partition(0,-1);
}

```

See also:

`XM_send_queuing_message` [Page: 53], `XM_receive_queuing_message` [Page: 45],
`XM_get_queuing_port_status` [Page: 23], `XM_get_queuing_port_info` [Page: 22].

2.4 XM_create_sampling_port

Synopsis: Create a sampling port.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_create_sampling_port (char      *portName,
                                   xm_u32_t  maxMsgSize,
                                   xm_u32_t  direction);
```

Description:

The XM_create_sampling_port() service is used to create a sampling port. Only those sampling ports specified in the configuration file XM_CF can be created by the partition. Note that all the parameters must match the information contained in the configuration file. New ports cannot be created dynamically.

It is not an error to invoke this hypercall with an already created port. The same port descriptor is returned.

Return value:

Upon successful completion, XM_create_sampling_port() returns a non-negative integer representing the port descriptor. Otherwise, the function returns a negative value:

[XM_INVALID_CONFIG]

- The maximum number of sampling ports has been created.
- No sampling port of the partition is named portName in the configuration file.
- maxMsgSize is out of range or not compatible with the configuration.
- direction is invalid or not compatible with the configuration.
- If the port is a system port, the max_msg_size parameter does not match the size of the messages defined for those ports.
- The maximum number of open ports exceeded. This error condition should be also raised when compiling the configuration file (XM_CF).

Usage examples:

```
xm_s32_t portDesc;

portDesc = XM_create_sampling_port ("sample",
                                   30,
                                   XM_DESTINATION_PORT);

if (portDesc < 0) {
    XM_write_console("Port sample could not be created\n",34);
    XM_halt_partition(XM_PARTITION_SELF);
}
```

See also:

XM_write_sampling_message [Page: 91], XM_read_sampling_message [Page: 43],
 XM_get_sampling_port_status [Page: 25], XM_get_sampling_port_info [Page: 24].

2.5 XM_disable_irqs

Synopsis: Replaced by `XM_sparc_clear_pil()` on Sparc processors.

See also:

`XM_clear_irqmask` [Page: 12], `XM_set_irqmask` [Page: 55],
`XM_sparc_clear_pil` [Page: 66], `XM_sparc_set_pil` [Page: 77]

2.6 XM_enable_irqs

Synopsis: Replaced by XM_sparc_set_pil() on Sparc processors.

Declaration:

```
void XM_enable_irqs(void);
```

Description:

All native external interrupts and virtual device events are “enabled”. Note that this function operates on the virtual state on the partition and not on the native hardware.

If an non-masked interrupt occurs while interrupts are disabled, then the highest priority interrupt will then be received by the partition when interrupts are enabled.

Return value:

This hypercall always succeed.

See also:

XM_clear_irqmask [Page: 12], XM_set_irqmask [Page: 55],
XM_sparc_clear_pil [Page: 66], XM_sparc_set_pil [Page: 77]

2.7 XM_get_partition_mmap

Synopsis: This function returns a pointer to the memory map table (MMT).

Category:

Library service

Declaration:

```
struct xmPhysicalMemMap *XM_get_partition_mmap (void);
```

Description:

XM_get_partition_mmap returns the pointer to the memory map table (MMT) data structure. The MMT is a read-only data structure with the following fields:

```
struct xmPhysicalMemMap {
    xmAddress_t startAddr;
    xmAddress_t mappedAt;
    xmSize_t size;
#define XM_MEM_AREA_SHARED      (1<<0)
#define XM_MEM_AREA_UNMAPPED  (1<<1)
#define XM_MEM_AREA_READONLY  (1<<2)
#define XM_MEM_AREA_UNCACHEABLE (1<<3)
#define XM_MEM_AREA_ROM       (1<<4)
#define XM_MEM_AREA_FLAG0     (1<<5)
#define XM_MEM_AREA_FLAG1     (1<<6)
#define XM_MEM_AREA_FLAG2     (1<<7)
#define XM_MEM_AREA_FLAG3     (1<<8)
#define XM_MEM_AREA_SHLIB     (1<<9)
    xm_u32_t flags;
};
```

Listing 2.1: /core/include/guest.h

Return value: This function always succeeds.

Rationale:

With this table, the partition knows the amount of memory it is using, and where it is initially mapped. It can also be used to translate virtual to physical addresses.

History:

Added in version 3.1.2

Usage examples:

```
XM_get_partition_mmap();
```

2.8 XM_get_partition_status

Synopsis: Get the current status of a partition.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_get_partition_status (xmId_t id,
                                 xmPartitionStatus_t *state);
```

Description:

Returns in state a structure containing the current state of the partition referenced by id.

The xmPartitionStatus_t is a data type containing the following fields:

```
typedef struct {
    /* Current state of the partition: ready, suspended ... */
    xm_u32_t state;
#define XM_STATUS_IDLE 0x0
#define XM_STATUS_READY 0x1
#define XM_STATUS_SUSPENDED 0x2
#define XM_STATUS_HALTED 0x3
    /* Number of virtual interrupts received. */
    xm_u64_t noVirqs; /* [[OPTIONAL]] */
    /* Reset information */
    xm_u32_t resetCounter;
    xm_u32_t resetStatus;
    xmTime_t execClock;
    /* Total number of partition messages: */
    xm_u64_t noSamplingPortMsgsRead; /* [[OPTIONAL]] */
    xm_u64_t noSamplingPortMsgsWritten; /* [[OPTIONAL]] */
    xm_u64_t noQueuingPortMsgsSent; /* [[OPTIONAL]] */
    xm_u64_t noQueuingPortMsgsReceived; /* [[OPTIONAL]] */
} xmPartitionStatus_t;
```

Listing 2.2: /core/include/objects/status.h

The fields labeled as OPTIONAL will not be updated if the “Enable system/partition status accounting” source configuration parameter is not set. It is disabled by default .

Return value:

[XM_OK]

The operation succeeded.

[XM_PERM_ERROR]

The calling partition is not supervisor.

[XM_INVALID_PARAM]

id is not a valid identifier, or state is not a valid address.

Rationale:

This function replaces the deprecated XM_unmask_irq().

History:

Introduced in XtratuM 2.2.0.

Usage examples:

```
xmPartitionStatus_t state;  
...  
XM_get_partition_status (XM_PARTITION_SELF, &state);
```

See also:

`XM_get_system_status` [Page: 26].

2.9 XM_get_plan_status

Synopsis: Return information about the scheduling plans.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_get_plan_status (xmSystemStatus_t *state);
```

Description:

Returns in state the information about the previous, current and next plans.

The xmPlanStatus_t is a data type which contains the following fields:

```
typedef struct {
    xmTime_t switchTime;
    xm_s32_t next;
    xm_s32_t current;
    xm_s32_t prev;
} xmPlanStatus_t;
```

Listing 2.3: /core/include/objects/status.h

[switchTime]

The absolute time of the last plan switch request. After a reset (both warm and cold), the value is set to zero.

[current]

Identifier of the current plan.

[next]

The plan identifier that will be active on the next major frame. If no plan switch is going to occur, then the value of next is equal to the value of current.

[prev]

The identifier of the plan executed before the current one. After a reset (both warm and cold) the value is set to (-1).

Return value:

[XM_OK]

The operation succeeded.

[XM_PERM_ERROR]

The calling partition is not a system partition.

[XM_INVALID_PARAM]

state is not a valid partition address.

History:

Introduced in XtratuM 3.1.2.

See also:

[XM_set_plan](#) [Page: 59].

2.10 XM_get_queuing_port_info

Synopsis: Get the info of a queuing port.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_get_queuing_port_info (xm_u32_t portDesc,
                                   xmQueuingPortInfo_t *info);
```

Description:

This function fills the structure `info`, with the following information:

```
typedef struct {
    char *_gParam portName;
    xm_u32_t maxMsgSize; // Max message size.
    xm_u32_t maxNoMsgs; // Max number of messages.
    xm_u32_t direction;
} xmQueuingPortInfo_t;
```

Listing 2.4: /core/include/objects/commports.h

Return value:

The return values are:

[XM_OK]

The operation succeeds.

[XM_INVALID_PARAM]

`portDesc` does not identify an existing destination queuing port, or `info` is not a valid partition address.

Usage examples:

```
xmQueuingPortInfo_t info;
xm_u32_t portDesc = XM_create_queuing_port(portName,
                                           maxNoMsgs,

                                           ...
                                           XM_get_queuing_port_info (portDesc, &info);
```

See also:

[XM_send_queuing_message](#) [Page: 53], [XM_receive_queuing_message](#) [Page: 45],

[XM_create_queuing_port](#) [Page: 14], [XM_get_queuing_port_status](#) [Page: 23].

2.11 XM_get_queuing_port_status

Synopsis: Get the status of a queuing port.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_get_queuing_port_status (xm_u32_t portDesc,  
                                     xmQueuingPortStatus_t *status  
                                     );
```

Description:

This function fills the structure status, with the following information:

```
typedef struct {  
    xm_u32_t maxMsgSize; // Max message size.  
    xm_u32_t maxNoMsgs;  // Max number of messages.  
    xm_u32_t noMsgs;     // Current number of messages.  
    xm_u32_t flags;  
} xmQueuingPortStatus_t;
```

Listing 2.5: /core/include/objects/commports.h

Return value:

The return values are:

[XM_OK]

The operation succeeds.

[XM_INVALID_PARAM]

portDesc does not identify an existing destination queuing port, or status is not a valid partition address.

Usage examples:

```
xmQueuingPortStatus_t info;  
xm_u32_t portDesc = XM_create_queuing_port(portName,  
      maxNoMsgs,maxMsgSize, direction);  
...  
XM_queuing_port_status (portDesc, &info);
```

See also:

XM_send_queuing_message [Page: 53], XM_receive_queuing_message [Page: 45],

XM_create_queuing_port [Page: 14].

2.12 XM_get_sampling_port_info

Synopsis: Get the info of a sampling port.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_get_sampling_port_info (xm_u32_t portDesc,
                                   xmSamplingPortInfo_t *info);
```

Description:

This function fills the structure `info`, with the configuration information:

```
typedef struct {
    char *_gParam portName;
    xmTime_t refreshPeriod; // Refresh period.
    xm_u32_t maxMsgSize; // Max message size.
    xm_u32_t direction;
} xmSamplingPortInfo_t;
```

Listing 2.6: /core/include/objects/commports.h

Return value:

[XM_OK]

The operation succeeds.

[XM_INVALID_PARAM]

`portDesc` does not identify an existing destination queuing port, or `info` is not a valid partition address.

Usage examples:

```
xmSamplingPortInfo_t info;
xm_u32_t portDesc = XM_create_sampling_port(portName,

...
XM_get_sampling_port_info (portDesc, &info);
```

See also:

`XM_write_sampling_message` [Page: 91], `XM_read_sampling_message` [Page: 43],
`XM_create_sampling_port` [Page: 15], `XM_get_sampling_port_status` [Page: 25].

2.13 XM_get_sampling_port_status

Synopsis: Get the status of a sampling port.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_get_sampling_port_status (xm_u32_t portDesc,
                                     xmSamplingPortStatus_t *status)
                                     ;
```

Description:

This function fills the structure status, with the following information:

```
typedef struct {
    xmTime_t refreshPeriod; // Refresh period.
    xm_u32_t maxMsgSize; // Max message size.
    xm_u32_t flags;
} xmSamplingPortStatus_t;
```

Listing 2.7: /core/include/objects/commports.h

Return value:

[XM_OK]

The operation succeeds.

[XM_INVALID_PARAM]

portDesc does not identify an existing destination queuing port, or status is not a valid partition address.

Usage examples:

```
xmSamplingPortStatus_t status;
xm_u32_t portDesc = XM_create_sampling_port(portName,
      maxMsgSize, direction);
...
XM_get_sampling_port_status (portDesc, &status);
```

See also:

XM_write_sampling_message [Page: 91], XM_read_sampling_message [Page: 43],

XM_create_sampling_port [Page: 15], XM_get_sampling_port_info [Page: 24].

2.14 XM_get_system_status

Synopsis: Get the current status of the system.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_get_system_status (xmSystemStatus_t *state);
```

Description:

Returns the current state of the system into the state structure.

The `xmSystemStatus_t` is a data type which contains the following fields:

```
typedef struct {
    xm_u32_t resetCounter;
    /* Number of HM events emitted. */
    xm_u64_t noHmEvents;          /* [[OPTIONAL]] */
    /* Number of HW interrupts received. */
    xm_u64_t noIrqs;             /* [[OPTIONAL]] */
    /* Current major cycle interation. */
    xm_u64_t currentMaf;         /* [[OPTIONAL]] */
    /* Total number of system messages: */
    xm_u64_t noSamplingPortMsgsRead; /* [[OPTIONAL]] */
    xm_u64_t noSamplingPortMsgsWritten; /* [[OPTIONAL]] */
    xm_u64_t noQueuingPortMsgsSent; /* [[OPTIONAL]] */
    xm_u64_t noQueuingPortMsgsReceived; /* [[OPTIONAL]] */
} xmSystemStatus_t;
```

Listing 2.8: `/core/include/objects/status.h`

The fields labeled as `OPTIONAL` will not be updated if the “Enable system/partition status accounting” source configuration parameter is not set. By default it is disabled.

Return value:

[XM_OK]

The operation succeeded.

[XM_PERM_ERROR]

The calling partition is not supervisor.

[XM_INVALID_PARAM]

state is not a valid partition address.

History:

Introduced in XtratuM 2.2.0.

Usage examples:

```
xmSystemStatus_t status;
XM_get_system_status (&status);
```

See also:

`XM_get_partition_status` [Page: 19].

2.15 XM_get_time

Synopsis: Retrieve the time of the clock specified in the parameter.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_get_time (xm_u32_t clockId, *xmTime_t time);
```

Description:

This function retrieves the number of **microseconds** elapsed since the last hardware reset as seen by the clock specified in the parameter. The time is stored in the memory pointed to by `time`.

XtratuM provides the next clocks:

```
#define XM_HW_CLOCK (0x0)
#define XM_EXEC_CLOCK (0x1)
```

Listing 2.9: /core/include/hypcalls.h

These are monotonic non-decreasing clocks. The resolution is one microsecond, and also the time is represented in microseconds.

Time is represented with an unsigned 64bit integer, which can hold sufficient microseconds to represent more than 290 thousand years.

Return value:

[XM_OK]

The operation succeeds.

[XM_INVALID_PARAM]

If the `clockId` is a non valid virtual clock.

Rationale:

Since the resolution of the clock is 1 microsecond, the same time may be returned if the function is called twice rapidly.

A data type of 64 bits is large enough even to measure the time in nanoseconds, and produce an overflow after a reasonable amount of time.

Usage examples:

```
xmTime_t t1,t2,t3;
char msg[100];

XM_get_time(XM_HW_CLOCK, &t1);
XM_get_time(XM_HW_CLOCK, &t2);
do_some_thing();
XM_get_time(XM_HW_CLOCK, &t3);
snprintf(msg, 100, "Measured duration: %lld      ",(t3-t2)-(t2-t1
));
XM_write_console(msg,29);
```

See also: [XM_set_timer](#) [Page: 60]

2.16 XM_halt_partition

Synopsis: Terminates a partition.

Category:

Standard service/Supervisor reserved.

Declaration:

```
xm_s32_t XM_halt_partition (xmId_t partitionId);
```

Description:

The `XM_halt_partition()` hypercall causes the termination of the `partitionId` partition. The partition is set in halt state.

The hypervisor will not schedule the target partition until the partition is reset (`XM_reset_partition()`). If the partition is scheduled by a cyclic scheduler, the time slot allocated to the halted partition is left idle. All the resources allocated to the partition are released: Interrupt lines, I/O ports and communication ports. The RAM memory used by the partition is not deleted.

Only supervisor partitions can halt other partitions different from itself. Any partition can halt itself. The constant `XM_PARTITION_SELF` represents the calling partition.

The target partition is in charge of copying in a non-volatile medium the information to be saved, if any. Although the RAM memory is not erased when this hypercall is called, depending on how the partition will be reset (system hardware reset, cold partition reset or warm partition reset) the memory may be deleted.

If the target partition was in halt state, then this hypercall has no effect.

Return value:

If the target partition (`partitionId`) is the calling partition then this function always succeeds, and the hypercall does not return. Otherwise, the return value is:

[XM_OK]

The target partition has been successfully halted, or the target partition was already in halt state.

[XM_INVALID_PARAM]

`partitionId` is not a valid partition identifier.

[XM_PERM_ERROR]

The calling partition is not supervisor and the target partition is not itself.

Usage examples:

```
...
XM_halt_partition(XM_PARTITION_SELF);
/* This code is never executed unless the partition is restored
   by a system partition*/
```

See also:

`XM_reset_partition` [Page: 48], `XM_suspend_partition` [Page: 79],

`XM_resume_partition` [Page: 51].

2.17 XM_halt_system

Synopsis: Stop the system.

Category:

Supervisor reserved.

Declaration:

```
xm_s32_t XM_halt_system();
```

Description:

The board is halted immediately. The whole system is stopped: interrupts are disabled and XtratuM executes an endless loop.

This function shall be used with extreme caution. Only a hardware reset can reboot the system. If the watchdog is armed, when the watchdog timer will expire the board will restart.



This is a supervisor reserved function.

Return value:

The function does not return if the operation succeeds. In the case of error, the return code is:

[XM_PERM_ERROR]

The calling partition is not a supervisor partition.

Rationale:

This service is provided to allow to stop the system in case of non-recoverable malfunctioning of the hardware.

Usage examples:

```
XM_halt_system();  
//This code will never be executed
```

See also:

XM_reset_system [Page: 50]

2.18 XM_hm_open

Synopsis: DEPRECATED

Category:

Standard service.

Declaration:

```
xm_s32_t XM_hm_open ();
```

Description:

This function has no real effect and will be removed of XtratuM in future versions.

Return value:

[XM_OK]

The operation succeeded.

[XM_PERM_ERROR]

The calling partition is not supervisor.

History:

Introduced in XtratuM 2.2.0, DEPRECATED.

Usage examples:

```
XM_hm_open ();
```

See also:

XM_hm_read [Page: 31], XM_hm_seek [Page: 32], XM_hm_status [Page: 34]

2.19 XM_hm_read

Synopsis: Read a health monitoring log entry.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_hm_read (xmHmLog_t *hmLogPtr);
```

Description:

This hypercall attempts to read one health monitoring log entry from the current read position into the structure `hmLogPtr`. If no log entry is available, the call returns immediately.

If the read pointer is not at the end of the stream, then it moves forward to the next log entry.

Note that this operation is not destructive. That is, the log entries are not removed from the internal buffer once read. It is possible to retrieve already read entries changing the read position with the hypercall `XM_hm_seek()`.

Every health monitoring log entry is a data `xmTraceEvent_t` structure (see [Health Monitoring overview \[Page: 6\]](#)).

Return value:

[0] There are no new log entries in the HM log stream. No data is returned.

[1]

The next log entry is returned and the current read position is increased by one.

[XM_PERM_ERROR]

The calling partition is not supervisor.

[XM_INVALID_PARAM]

The `hmLogPtr` address is not a valid address.

History:

Introduced in XtratuM 2.2.0.

Usage examples:

```
xmHmLog_t hmLogEntry;
...
if (XM_hm_open() != XM_OK) {
    XM_halt_partition(XM_PARTITION_SELF);
}
while (1) {
    XM_idle_self();
    XM_hm_read(&hmLogEntry);
    ProcessHmEntry(&hmLogEntry);
}
```

See also:

[XM_hm_open \[Page: 30\]](#), [XM_hm_seek \[Page: 32\]](#), [XM_hm_status \[Page: 34\]](#)

2.20 XM_hm_seek

Synopsis: Sets the read position in the health monitoring stream.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_hm_seek (xm_s32_t offset, xm_u32_t whence);
```

Description:

This hypercall repositions the read position of the health monitoring stream to the argument `offset` according to the directive `whence` as follows:

[XM_OBJ_SEEK_SET]

The new read position is `offset` entries from the oldest one stored health monitoring log entry. Only positive values (or zero) of `offset` are valid.

[XM_OBJ_SEEK_CUR]

The new position to read is set to the current position plus the `offset` value. Negative values of `offset` refer to older (previous) log entries. Positive values refer to newer (posterior) log entries. If `offset` is zero then the current position is not changed.

[XM_OBJ_SEEK_END]

The new read position is set to the end of the stream. Negative values of `offset` refer to older (previous) entries. If `offset` is zero then the next read operation will not return a log entry (unless a new hm event is recorded after the completion of this call).

Note that the `offset` value represents the number of log entries, and not the size in bytes.

`offset` shall not be greater than the size of the health monitoring log buffer.

If the computed new read position is beyond the beginning or the end of the stream, then it will be truncated to the start or the end of the stream respectively.



Since the stream is a circular buffer, it may happen that right after the `XM_hm_seek` operation, one or more new hm log entries are stored in the stream. In this case, the pointer may not be positioned at the start or the end of the stream, but a few entries after or before it. To avoid this concurrency issues, it is advisable not to move the pointer too close to the start of the buffer when the buffer is full.

Return value:

[XM_OK]

The position has been set.

[XM_PERM_ERROR]

The calling partition is not supervisor.

[XM_INVALID_PARAM]

The `offset` or `whence` are not valid.

History:

Introduced in XtratuM 2.2.0.

Usage examples:

```
xmHmLog_t hmBuffer[10];
XM_hm_seek (0, XM_OBJ_SEEK_CUR);
for (x=0; x<10 ;x++){
    /* Read the oldest 10 entries */
    XM_hm_read (&hmBuffer[x]);
}
```

See also:

XM_hm_open [Page: 30], XM_hm_read [Page: 31], XM_hm_status [Page: 34]

2.21 XM_hm_status

Synopsis: Get the status of the health monitoring log stream.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_hm_status (xmHmStatus_t *hmStatusPtr);
```

Description:

This hypercall returns information about the XtratuM health monitoring log stream in the structure `hmStatusPtr`.

This service return a stat structure, which contains the following fields:

```
typedef struct {
    xm_s32_t noEvents;
    xm_s32_t maxEvents;
    xm_s32_t currentEvent;
} xmHmStatus_t;
```

Listing 2.10: `/core/include/objects/hm.h`

Return value:

[XM_OK]

The operation succeeded.

[XM_PERM_ERROR]

The calling partition is not supervisor.

[XM_INVALID_PARAM]

The `hmStatusPtr` address is not a valid partition address.

History:

Introduced in XtratuM 2.2.0.

Usage examples:

```
xmHmLogStatus_t hmLogStatus;
...
XM_hm_status (&hmLogStatus);
```

See also:

`XM_hm_open` [Page: 30], `XM_hm_read` [Page: 31], `XM_hm_seek` [Page: 32]

2.22 XM_idle_self

Synopsis: Idles the execution of the calling partition.

Category:

Standard service.

Declaration:

```
xm_u32_t XM_idle_self (void);
```

Description:

Suspends the execution of the calling partition until a non masked interrupt is received by the partition or at the start of the next scheduling slot, whatever happens first.

Return value: This function always succeeds: XM_OK.

Rationale:

The use of this function improves the overall system performance. Rather than waiting on a busy loop for a trap, the partition can yield the processor to do other activities or to reduce power consumption by lowering the processor frequency (if supported).

The partition developer can use this service to synchronize the execution of the partition with the scheduling plan.

Usage examples:

The next example shows how the information reported in the `partitionControlTable_t` data structure can be used to synchronise the execution of a partition with the system cyclic plan.

```
#include <guest.h>

#define HW_IRQS 16

partitionControlTable_t pct;

...

/* Mask all partition hw interrupts */
for (i=XM_VT_HW_FIRST; i <= HW_VT_HW_LAST; i++) {
    XM_mask_irq(i);
}
/* Mask all partition extended interrupts */
for (i=XM_VT_EXT_FIRST; i <= HW_VT_EXT_LAST; i++) {
    XM_mask_irq(i);
}
/* Unmask the start slot interrupt */
XM_unmask_irq(XM_VT_EXT_CYCLIC_SLOT_START);
XM_enable_irqs();
/* Wait the start of the next MAF */
if (pct.tag == 0) { /* If we are in the first slot of
                    the current MAF, then skip it. */
    XM_idle_self();
}
while (pct.tag != 0) {
    XM_idle_self(); /* Skip all the non-starting slots */
}
```

```
}
/* Implementation of a cyclic plan to run partition tasks: */
while (1) {
    TaskA(); TaskB();
    XM_idle_self(); /* Wait for the first slot. */
    TaskC(); TaskB();
    XM_idle_self(); /* Wait for the second slot. */
    TaskD(); TaskE(); TaskB();
    XM_idle_self(); /* Wait for the third slot. */
}
XM_write_console("PANIC: unreachable code\n",25)
XM_halt_partition(XM_PARTITION_SELF);

...
```

Note that this code is not robust, and should be improved if used in a product.

2.23 XM_mask_irq

Synopsis: Obsoleted by XM_set_irqmask().

Description:

None

History:

Removed in version 3.1.2

See also:

[XM_set_irqmask](#) [Page: 55].

2.24 XM_memory_copy

Synopsis: Copy copies data from/to address spaces.

Category:

Standard service/Supervisor reserved.

Declaration:

```

xm_s32_t XM_memory_copy (xmId_t destId, xm_u32_t destAddr,
                        xmId_t srcId, xm_u32_t srcAddr,
                        xm_u32_t size);

```

Description:

This function copies data from/to address spaces. The addresses (source and/or destination) can point to both memory or mapped I/O registers.

This function copies `size` bytes from the area pointed by `srcAddr` located in the address space of `srcId` partition to the address `destAddr` in the memory space of `destId` partition.

The following considerations shall be taken into account:

- The source and destination areas shall not overlap to avoid data corruption.
- When copying from memory to memory, for efficiency reasons, both areas shall be 8 bytes aligned.
- Since this function allows to copy large blocks of memory, care must be taken to avoid breaking temporal isolation.
- This function has been implemented using one object descriptor (associated with the `/dev/mem` object). The maximum number of object descriptors per partition is a source code configuration parameter, and shall be large enough.
- If the source or destination addresses do not belong to the space of a partition (for example, ROM areas) then the `XM_SYSTEM_ID` shall be used.
Only supervisor partitions are allowed to perform a copy for/to address spaces other than its own (i.e. other partitions or system space).
- The source or/and destination are mapped I/O registers, then the I/O mapped addresses shall be word (4 bytes) aligned.
- Note that the `size` parameter indicates the number of **bytes** to be copied, and an I/O register is 4 bytes.
- Only memory mapped I/O registers fully allocated to the partition (that is, using the `Range` element in the `XM_CF` configuration file) can be addressed by the `XM_memory_copy` function.

Return value:

[`XM_OK`]

The operation succeeds.

[`XM_INVALID_PARAMS`]

- `destId` or `srcId` are not valid partition Id's,
- `destAddr` does not belong to the `destId` partition,
- `srcAddr` does not belong to the `srcId` partition.

[`XM_INVALID_CONFIG`]

The maximum number of open object descriptors has been exceeded. If this happens, then XtratuM shall be reconfigured to increase this limit and recompiled accordingly.



[XM_PERM_ERROR]

The calling partition is not supervisor.

Rationale:

Since the LEON2 does not have MMU, both address (`destAddr` and `srcAddr`) are physical memory addresses. Partition identifiers are not needed, but has been included for future compatibility.

Partition identifiers are not checked in the current implementation.

Usage examples:

```
xm_u32_t i=10;
xm_u32_t j;
XM_memcpy(XM_PARTITION_SELF,&i,XM_PARTITION_SELF,&j,sizeof(
    xm_u32_t));
```

2.25 XM_multicall

Synopsis: Execute a sequence of hypercalls.

Category:

Standard service. Optimization. [Not ported to LEON2]

Declaration:

```
xm_s32_t XM_multicall(void *buffer, xm_s32_t nr);
```

Description:

In some cases, the partition has to perform a large sequence of hypercalls to perform an operation. For example, several `XM_sparc_outport()` calls may be needed to program a single peripheral operation; another common case is when managing memory maps (not implemented in this version).

In order to reduce the overhead caused by the hypercall mechanism, XtratuM provides this hypercall to **pack** several hypercalls in a buffer, and then execute them using a single hypercall.

This hypercall does not provide new functionality, it is only a mechanism that can be used to improve the performance.

Return value: TBD.

2.26 XM_params_get_PCT

Synopsis: Return the address of the PCT.

Category:

Library service.

Declaration:

```
partitionControlTable_t * XM_params_get_PCT (void);
```

Description:

XM_params_get_PCT returns the pointer to the PCT (Partition Control Table) data structure. The PCT is a read-only data structure.

```
typedef struct {
    xm_u32_t magic;
    xm_u32_t xmVersion; // XM version
    xm_u32_t xmAbiVersion; // XM's abi version
    xm_u32_t xmApiVersion; // XM's api version
    xm_u32_t resetCounter;
    xm_u32_t resetStatus;
    xm_u32_t cpuKhz;
    xmId_t id;
    xm_u32_t flags;
    xm_u32_t hwIrqs; // Hw interrupts belonging to the partition
    xm_s32_t noPhysicalMemAreas;
    xm_u8_t name[CONFIG_ID_STRING_LENGTH];
    xm_u32_t iFlags;
    xm_u32_t hwIrqsPend; // pending hw irqs
    xm_u32_t hwIrqsMask; // masked hw irqs

    xm_u32_t extIrqsPend; // pending extended irqs
    xm_u32_t extIrqsMask; // masked extended irqs

    struct pctArch arch;
    struct {
        xm_u32_t noSlot:16, reserved:16;
        xm_u32_t id;
        xm_u32_t slotDuration;
    } schedInfo;
    xm_u16_t trap2Vector[NO_TRAPS];
    xm_u16_t hwIrq2Vector[CONFIG_NO_HWIRQS];
    xm_u16_t extIrq2Vector[XM_VT_EXT_MAX];
} partitionControlTable_t;
```

Listing 2.11: /core/include/guest.h

Return value: This function always succeeds.

History:

Added in version 3.1.2

2.27 XM_read_console

Synopsis: Print a string in the hypervisor console.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_read_console (char *buf, xm_s32_t counter);
```

Description:

`XM_read_console` attempts to read up to `counter` bytes from the device configured as the main console into the buffer `buf`.

Return value: On success, the number of bytes read is returned.

[`XM_PERM_ERROR`]

The partition is not a systems partition.

[`XM_INVALID_PARAM`]

Invalid buffer pointer. Or invalid counter.

Usage examples:

```
char message[30];  
message[29]=0x0;  
XM_read_console(message,29);
```

See also:

`XM_write_console` [Page: 89]

2.28 XM_read_sampling_message

Synopsis: Reads a message from the specified sampling port.

Category:

Standard service.

Declaration:

```

xm_s32_t XM_read_sampling_message (xm_s32_t portDesc,
                                   void *msgPtr,
                                   xm_u32_t size,
                                   xm_u32_t *flags);

```

Description:

The `XM_read_sampling_message()` service is used to read a message from the specified sampling port. If succeed, at most `size` bytes are copied into the buffer pointed by `msgPtr`.

If `flags` is not a null pointer, then the validity bit (`XM_MSG.VALID`) is set accordingly: the bit is set if the age of the read message is consistent with the `@validPeriod` optional attribute of the channel. Otherwise the bit is reset. Note that the message is considered to be correct, even if the `XM_MSG.VALID` bit is reset.

Return value:

When successful, the minimum between the length of the received message and the `size` parameter is returned. On error, one of the following negative values is returned:

[`XM_INVALID_PARAM`]

- `portDesc` does not identify a valid port.
- The specified sampling port is not configured as `XM_PORT_DESTINATION`.
- The value of `size` is zero.

[`XM_INVALID_CONFIG`]

The `size` is not compatible with the configuration of the specified port.

Optimization:

If the buffer is aligned to 8 bytes, then the copy operation is performed faster. Also, the copy operation is more efficient if the length of the message is multiple of 8 bytes.

Usage examples:

```

xm_s32_t portDesc, ret_code;
char message[30];

xm_u32_t flags;

portDesc = XM_create_sampling_port ("sample",
                                   30,
                                   XM_PORT_DESTINATION);

if (portDesc < 0) {
    XM_halt_partition(XM_PARTITION_SELF);
}

```

```
ret_code = XM_read_sampling_message (portDesc, message, 30, &
    flags);

if (ret_code < 0) {
    XM_write_console("Error reading sampling port\n",28);
    return;
}

if (!(flags & XM_MSG_VALID)) {
    XM_write_console("The message is not valid!\n",27);
    return;
}

message[29]=0x0; /* For safety */
XM_write_console(message,ret_code);
```

See also:

[XM_create_sampling_port \[Page: 15\]](#), [XM_write_sampling_message \[Page: 91\]](#),
[XM_get_sampling_port_status \[Page: 25\]](#), [XM_get_sampling_port_info \[Page: 24\]](#).

2.29 XM_receive_queuing_message

Synopsis: Receive a message from the specified queuing port.

Category:

Standard service.

Declaration:

```

xm_s32_t XM_receive_queuing_message (xm_s32_t portDesc,
                                     void      *msgPtr,
                                     xm_u32_t  size);

```

Description:

If the port channel is not empty, then the oldest message is retrieved from the channel. At most size bytes are copied into the buffer pointed by msgPtr. The actual number of bytes received is returned by the call (if the hypercall succeeds).

If the hypercall succeeds, then the message is removed from the XtratuM channel. Note that the message is consumed (considered as correctly read) even if the receiving partition only reads the message partially (the parameter size is smaller than the actual size of the message in the channel).

Return value:

On success, the minimum between the length of the received message and the size parameter is returned. On error, one of the following negative values is returned:

[XM_NOT_AVAILABLE]

The channel is empty or the port is not connected to a channel in the XM_CF configuration file.

[XM_INVALID_PARAM]

portDesc does not identify an existing queuing port or it was not configured as a destination port; the value of size is zero.

[XM_INVALID_CONFIG]

The size is not compatible with the configuration of the specified port.

Optimization:

If the buffer is aligned to 8 bytes, then the copy operation is performed faster. Also, the copy operation is more efficient if the length of the message is multiple of 8 bytes.

Usage examples:

```

xm_s32_t portDesc, ret_code;
char message[30];

xm_u32_t flags;

portDesc = XM_create_queuing_port ("queue", 4, 30,
                                   XM_PORT_DESTINATION);
if (portDesc < 0) {
    XM_halt_partition(XM_PARTITION_SELF);
}
ret_code = XM_receive_queuing_message (portDesc, message, 30,
                                       &flags);

```

```
if (ret_code < 0) {
    XM_write_console("Error reading sampling port\n",28);
    return;
}
if (!(flags & XM_MSG_VALID)) {
    XM_write_console("The message is not valid!\n",27);
    return;
}
message[29]=0x0; /* For safety */
XM_write_console(message,ret_code);
```

See also:

[XM_create_queuing_port](#) [Page: 14], [XM_send_queuing_message](#) [Page: 53],

[XM_get_queuing_port_status](#) [Page: 23], [XM_get_queuing_port_status](#) [Page: 23].

2.30 XM_request_irq

Synopsis: Request to receive an interrupt.

Category:

Removed.

Declaration:

```
xm_s32_t XM_request_irq (xm_u32_t irq)
```

Description:

This hypercall is no longer available.

Rationale:

This service was initially provided for compatibility with the ARINC-653 standard. It checks that the requested interrupt line was effectively allocated to the calling partition.

Since all hypercalls that deal with interrupts also perform that check, this service is redundant.

2.31 XM_reset_partition

Synopsis: Reset a partition.

Category:

Standard service/Supervisor reserved.

Declaration:

```
xm_s32_t XM_reset_partition (xm_s32_t partitionId,
                             xm_u32_t resetMode,
                             xm_u32_t status);
```

Description:

The partition `partitionId` is reset. There are two ways to reset a partition depending on the `resetMode` value:

XM_WARM_RESET:

1. The `resetCounter` field of the partition information table (PIT) is incremented.
2. The parameter `status` is copied in the field `resetStatus` of the partition information table
3. The program counter is set to the partition entry point.
4. The partition is set in normal/ready state.

XM_COLD_RESET:

1. All communication ports are closed.
2. The PCT and PIT data structures are initialised in the partition memory space.
3. The `resetCounter` field of the partition information table (PIT) is set to zero.
4. The parameter `status` is copied in the field `resetStatus` of the partition information table
5. The program counter is set to the partition entry point.
6. The partition is set in normal/ready state.

The partition can use the values of `resetCounter` and `resetStatus` to perform different actions after the reset.



Note that the memory of the partition is not modified, i.e, the content of the memory will be the same (except the PIT and PCT) than before the reset.

Only supervisor partitions can reset other partitions than itself. Any partition can reset itself. The constant `XM_PARTITION_SELF` represents the calling partition.

Return value:

If the target partition (`partitionId`) is the calling partition then his function always succeeds, and the hypercall does not return. Otherwise, the return value is:

[XM_OK]

The target partition has been successfully reset.

[XM_INVALID_PARAM]

- `partitionId` is not a valid partition identifier.
- `resetMode` is not a valid reset mode (`XM_COLD_RESET` or `XM_WARM_RESET`).

[XM_PERM_ERROR]

A non supervisor partition attempted to reset other partition.

Usage examples:

```
...
XM_reset_partition(XM_PARTITION_SELF, XM_COLD_RESET, -9);
/* This code is never reached */
```

See also:

XM_halt_partition [Page: 28], XM_suspend_partition [Page: 79],
XM_resume_partition [Page: 51], XM_memory_copy [Page: 38]

2.32 XM_reset_system

Synopsis: Reset the system.

Category:

Supervisor reserved.

Declaration:

```
xm_s32_t XM_reset_system(xm_u32_t mode);
```

Description:

The system is reset immediately. There are two ways to reset the system depending on the mode value:

XM_WARM_RESET:

XtratuM unconditionally jumps to the XtratuM initialization. All the XtratuM data structures are initialised. Partitions are cold reset.

XM_COLD_RESET:

XtratuM unconditionally jumps back to the resident software entry point. This value shall be specified attribute `/SystemDescription/ResidentSw/@entryPoint` of the XML configuration file. If the `./ResidentSw/@entryPoint` attribute is not specified then XtratuM will jump to address `0x000000`.



This function shall be used with extreme caution. The state of the partitions will be lost unless it was saved in permanent memory before the reset.

Return value:

The function does not return if the operation succeeds. In case of error, the return code is:

[XM_PERM_ERROR]

The calling partition is not a supervisor partition.

[XM_INVALID_PARAM]

mode is not a valid mode.

Rationale:

This service can be used to try and recover from hardware errors.

Usage examples:

```
XM_reset_system(XM_COLD_RESET);
//This code is never executed
```

See also:

XM_halt_system [Page: 29]

2.33 XM_resume_partition

Synopsis: Resume the execution of a partition.

Category:

Supervisor reserved.

Declaration:

```
xm_s32_t XM_resume_partition (xmId_t partitionId);
```

Description:

Resumes the execution of the target partition, partitionId. If the target partition is not in suspended state then this function has no effect.

All the pending interrupts will be delivered when the partition is resumed.

Only supervisor partitions can invoke this service.

Return value:

[XM_OK]

The target partition has been resumed.

[XM_INVALID_PARAM]

partitionId is not a valid partition identifier.

[XM_PERM_ERROR]

The calling partition does not have supervisor rights.

Usage examples:

```
XM_resume_partition(suspendedPartitionId);
```

See also:

XM_suspend_partition [Page: 79].

2.34 XM_route_irq

Synopsis: Link an interrupt with the vector which is generated when the interrupt is issued.

Category:

Standard service

Declaration:

```
xm_s32_t XM_route_irq (xm_u32_t type, xm_u32_t irq, xm_u16_t
vector);
```

Description:

The `XM_route_irq()` hypercall binds an exception/hw-interrupt/extended-interrupt (given by the combination `type`: `XM_TRAP_TYPE`, `XM_HWIRQ_TYPE`, `XM_EXTIRQ_TYPE` and the number of interrupt) with an interrupt vector.

An interrupt vector is the index into the interrupt vector table. When an interrupt is triggered, this index is used by the hypervisor jointly with the base address of the interrupt vector table to jump to interrupt handler.

Return value:

This hypercall returns:

[`XM_OK`]

The parameter (`type` is either `XM_TRAP_TYPE` or `XM_HWIRQ_TYPE` or `XM_EXTIRQ_TYPE`, `irq` is a valid exception/hw-interrupt/extended-interrupt and `vector` is a valid interrupt vector).

[`XM_INVALID_PARAM`]

Either `type` or `irq` or `vector` is not valid.

2.35 XM_send_queuing_message

Synopsis: Send a message in the specified queuing port.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_send_queuing_message (xm_s32_t portDesc,  
                                void *msgPtr,  
                                xm_u32_t size);
```

Description:

The message is inserted into the XtratuM internal channel of the port, if enough space. Otherwise, the operation fails.

Return value:

[XM_OK]

The message has been successfully written into the port.

[XM_NOT_AVAILABLE]

Insufficient space in the channel or the port is not connected to a channel in the XM_CF configuration file.

[XM_INVALID_PARAM]

portDesc does not identify an existing destination queuing port; or size is zero.

[XM_INVALID_CONFIG]

The size is not compatible with the configuration of the specified port.

Optimization:

If the buffer is aligned to 8 bytes, then the copy operation is performed faster. Also, the copy operation is more efficient if the length of the message is multiple of 8 bytes.

Usage examples:

```
xm_s32_t portDesc, ret_code;  
char message="This is a dummy message";  
  
xm_u32_t flags;  
  
portDesc = XM_create_queuing_port ("queue", 4, 30,  
                                XM_PORT_DESTINATION);  
if (portDesc < 0) {  
    XM_halt_partition(XM_PARTITION_SELF);  
}  
ret_code = XM_write_queuing_message (portDesc, message, 24, &  
                                flags);  
if (ret_code < 0) {  
    XM_write_console("Error writing sampling port\n",28);  
    return;  
}
```

See also:

`XM_create_queuing_port` [Page: 14], `XM_receive_queuing_message` [Page: 45],
`XM_get_queuing_port_status` [Page: 23], `XM_get_queuing_port_status` [Page: 23].

2.36 XM_set_irqmask

Synopsis: Mask interrupts.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_set_irqmask(xm_u32_t hwIrqsMask, xm_u32_t
                        extIrqsMask);
```

Description:

This function masks (blocks) hardware and extended interrupts. The interrupt lines which are set in the `hwIrqsMask` and `extIrqsMask` parameters are unmasked, hardware and extended interrupts respectively.

The interrupt handlers of the masked interrupts will not be invoked until the interrupt line will be unmasked again.

Return value: This hypercall always succeed: `XM_OK`.

Rationale:

This function replaces the deprecated `XM_mask_irq()`.

History:

Added in version 3.1.2

Usage examples:

```
XM_set_irqmask(0, 0xffffffff); //this call leaves the HW_IRQ mask as it is and masks all extended
interruptions
```

See also:

`XM_clear_irqmask` [Page: 12]

2.37 XM_set_irqpend

Synopsis: Mask interrupts.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_set_irqpend(xm_u32_t hwIrqsMask, xm_u32_t
                        extIrqsMask);
```

Description:

This function simulates interrupt arrivals (i.e, forces) hardware and/or extended interrupts. The interrupt lines which are set in the hwIrqsPend and extIrqPend parameters are set as pending, hardware and extended interrupts respectively.

Return value: This hypercall always succeeds.

History:

Added in version 3.1.2

Usage examples:

```
XM_set_irq_pend(0,0x1); //This call leaves the pending HW_IRQ's as they are and sets all //EXT_IRQ
as pending.
```

See also:

[XM_clear_irqpend \[Page: 13\]](#), [XM_set_irqmask \[Page: 55\]](#), [XM_clear_irqmask \[Page: 12\]](#).

2.38 XM_set_page_type

Synopsis: Changes the type of the physical page pAddr to type.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_set_page_type(xmAddress_t pAddr, xm_u32_t type);
```

Description:

Partitions are able to declare new page tables and use them. In order to perform this operation, the type of the page selected to act as part of a page table must be changed to the appropriate type:

PPAG_STD Standard page, there is no restrictions to the permissions given to this kind of page. They can be mapped as read/write or read only.

PPAG_PTD1 A page of this type will hold a first level page table, any of its entries must either map a PPAG_PTD2 page or not to map anything (not present page). This sort of pages can only be mapped as read-only. This restriction is enforced by XtratuM.

PPAG_PTD2

A page of this type will act as a second level page table, there is no restrictions about the pages mapped by this page, however, PPAG_PTD1 and PPAG_PTD2 must be set as read-only pages.

Before change the type of a page, this page must be unmapped from all existing page tables, otherwise, the hypercall returns a XM_OP_NOT_ALLOWED error.

Return value:

[XM_OK] The operation finished successfully.

[XM_INVALID_PARAM]The page pAddr does not belong to the partition or its content breaks some rule for its new type

[XM_OP_NOT_ALLOWED]The page pAddr is mapped in an existing page table.

Usage examples:

```
#include <xm_inc/arch/physmm.h/>
#include <xm.h>

unsigned long pgd[1024] __attribute__((aligned(4096)));
unsigned long pgt[1024] __attribute__((aligned(4096)));
...

void PartitionMain(void) {
    memset(ptdL1, 0, 4096);
    memset(ptdL2, 0, 4096);
    if (XM_set_page_type(ptdL1, PPAG_PTD1)!=XM_OK) {
        xprintf("ptdL1 couldn't be set as PTD1\n");
        XM_halt_partition(XM_PARTITION_SELF);
    }
    if (XM_set_page_type(ptdL2, PPAG_PTD2)!=XM_OK) {
        xprintf("ptdL2 couldn't be set as PTD2\n");
    }
}
```

```
    XM_halt_partition(XM_PARTITION_SELF);
}
if (XM_update_page32(&ptdL1[0], (unsigned long)ptdL2|0x2)!=
    XM_OK) { // PTDL2 PRESENT
    xprintf("ptdL1[0] couldn't be updated\n");
    XM_halt_partition(XM_PARTITION_SELF);
}
if (XM_update_page32(&ptdL2[0], 0x40100000|0xe)!=XM_OK) { //
    PAGE PRESENT / USER_RW
    xprintf("ptdL2[0] couldn't be updated\n");
    XM_halt_partition(XM_PARTITION_SELF);
}
...
}
```

See also:

[XM_update_page32 \[Page: 87\]](#)

2.39 XM_set_plan

Synopsis: Request a plan switch at the end of the current MAF.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_set_plan(xm_u32_t planId);
```

Description:

The plan `planId` is scheduled to be started at the end of the current MAF.

Note that this hypercall does not force an immediate plan switch, but prepares the system to perform it. If this hypercall is called multiple times, the last call determines the new plan.

The plan zero is the initial plan and cannot be called back. Plan zero can only be activated by causing a system reset.

Return value:

[XM_OK]

The operation succeeded.

[XM_PERM_ERROR]

The calling partition is not system.

[XM_INVALID_PARAM]

`planId` is not a valid plan identifier. Valid plan identifiers are those specified in the XM_CF file, except the plan zero.

History:

Introduced in XtratuM 3.1.2.

See also:

`XM_get_plan_status` [Page: 21].

2.40 XM_set_timer

Synopsis: Arm a timer.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_set_timer (xm_u32_t clockId,
                      xmTime_t absTime,
                      xmTime_t interval);
```

Description:

If `interval` is zero, then the timer is armed only once. The timer associated with the virtual clock `clockId` is armed to expire at the absolute instant `absTime`. That is, the timer will expire when the clock reaches the value specified by the `absTime` parameter.

If `interval` is not zero, then the timer will expire periodically at absolute times: `absTime + n * interval`; where "n" starts in zero and repeats until the timer is re-armed.

If the specified `absTime` time has already passed, the function succeeds and the expiration interrupt happens immediately.

Once the timer is armed, the partition will receive a virtual timer interrupt on every timer expiration. It is responsibility of the partition code to install the corresponding interrupt handler.



Clock	Associated extended interrupt
XM_HW_CLOCK	XM_VT_EXT_HW_TIMER
XM_EXEC_CLOCK	XM_VT_EXT_EXEC_TIMER



There is only one timer per clock. Therefore, if the timer was already armed when the `XM_set_timer()` function is called, the previous value is reset and the timer is reprogrammed with the new values.

If `absTime` is zero then the timer is disarmed. If at the time of disarming the timer, there was pending timer interrupts, then the interrupt will not be removed, and will be delivered when appropriate. This may happen if `XM_set_timer()` function is called while interrupts are disabled or masked.

If a periodic timer expires several times before the interrupt is received (interrupt line masked, interrupts disabled, or the partition is not ready), then only one interrupt is delivered to the partition.

Return value:

[XM_OK]

The timer has been successfully armed.

[XM_INVALID_PARAM]

- The specified `clockId` is a non valid virtual clock (not a valid clock or the clock can not be used to arm timers).

Rationale:

Internally, the timers are managed in **one shot mode**. That is, the hardware timer is not programmed to generate an interrupt in a periodic way, but it is re-programmed to cause the interrupt exactly when the timer closer timer expires.

On systems where the cost of reprogramming the timer hardware is low (the case of the LEON2 board), the one shot timer mode is the best technique because it provides an high resolution and small overhead.

Although an absolute time point should be a positive number, the time is represented with a signed integer to detect incorrect time values.

Usage examples:

```
void ExtIrqHandler(int irqnr) {
    if (irqnr == XM_VT_EXT_HW_TIMER) {
        XM_unmask_irq(XM_VT_EXT_HW_TIMER);
        XM_write_console("Periodic..\n",11);
    } else {
        XM_write_console("Unexpected Irq\n",15);
    }
}

...
xmTime_t Start;
xmTime_t Period = (xmTime_t)10000;

XM_get_time(XM_HW_CLOCK, &Start);
Start += (xmTime_t)1000000;
XM_set_timer(XM_HW_CLOCK, Start, Period);
XM_enable_irqs();
XM_unmask_irq(XM_VT_EXT_HW_TIMER);
```

See also: [XM_get_time](#) [Page: 27]

2.41 XM_shutdown_partition

Synopsis: Send a shutdown interrupt to a partition.

Category:

Standard service/Supervisor reserved.

Declaration:

```
xm_s32_t XM_shutdown_partition (xmId_t partitionId);
```

Description:

The `XM_shutdown_partition()` hypercall raises the extended interrupt `XM_VT_EXT_SHUTDOWN` on the target partition.

On receiving a shutdown interrupt, the target partition shall close and terminate the ongoing tasks and finally call the `XM_halt_partition` [Page: 28] hypercall. XtratuM does not control the state the target partition after a shutdown request.

Only supervisor partitions can invoke this service to shutdown other partitions than itself. Any partition can shutdown itself. The constant `XM_PARTITION_SELF` represents the calling partition.

If the target partition was in halt state, then this hypercall has no effect.

Return value:

If the target partition (`partitionId`) is the calling partition then this function always succeeds. Otherwise, the return value is:

[`XM_OK`]

The shutdown trap has been successfully delivered.

[`XM_INVALID_PARAM`]

`partitionId` is not a valid partition identifier.

[`XM_PERM_ERROR`]

The calling partition is not supervisor and the target partition is not itself.

Usage examples:

```
XM_shutdown_partition(partitionId);
```

See also:

`XM_reset_partition` [Page: 48], `XM_suspend_partition` [Page: 79],

`XM_resume_partition` [Page: 51].

2.42 XM_sparc_atomic_add

Synopsis: Atomic add.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_sparc_atomic_add(xm_s32_t *target, xm_s32_t delta);
```

Description:

Adds the value `delta` to the value stored in `target` in an atomic manner.

Arithmetic overflow or underflow of the target value is not checked.

Return value:

[XM_OK]

If the operation has been done.

[XM_INVALID_PARAM]

If an invalid target pointer is given a XM_INVALID_PARAM value is returned.

Rationale:

SparcV8 processor lacks support for complex atomic operations. Since a partition is always executed with interrupts enabled, atomic operations have to be provided by XtratuM.

See also:

`XM_sparc_atomic_and` [Page: 64], `XM_sparc_atomic_or` [Page: 65]

2.43 XM_sparc_atomic_and

Synopsis: Atomic bitwise AND.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_sparc_atomic_and(xm_u32_t *target, xm_u32_t bits);
```

Description:

Performs the bitwise AND of bits to the value stored in target in an atomic manner.

Return value:

[XM_OK]

If the operation has been done.

[XM_INVALID_PARAM]

If an invalid target pointer is given a XM_INVALID_PARAM value is returned.

Rationale:

This service can be used to implement the atomic clear bit operation:

```
static inline void XM_sparc_atomic_clearbit(xm_u32_t *target,
                                             xm_u32_t pos){
    XM_sparc_atomic_and(target, ~(1 << pos));
}
```

See also:

XM_sparc_atomic_add [Page: 63], XM_sparc_atomic_or [Page: 65]

2.44 XM_sparc_atomic_or

Synopsis: Atomic bitwise OR.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_sparc_atomic_or(xm_u32_t *target, xm_u32_t bits);
```

Description:

Performs the bitwise OR of bits to the value stored in target in an atomic manner.

Return value:

[XM_OK]

If the operation has been done.

[XM_INVALID_PARAM]

If an invalid target pointer is given a XM_INVALID_PARAM value is returned.

Rationale:

This service can be used to implement the atomic set bit operation:

```
static inline void XM_sparc_atomic_setbit(xm_u32_t *target,
                                          xm_u32_t pos){
    XM_sparc_atomic_or(target, (1 << pos));
}
```

See also:

XM_sparc_atomic_add [Page: 63], XM_sparc_atomic_and [Page: 64]

2.45 XM_sparc_clear_pil

Synopsis: Clear the PIL field of the PSR (allows interrupts).

Category:

Standard service. Assembly hypercall.

Declaration:

```
void XM_sparc_clear_pil(void);
```

Description:

Clear the PIL field of the PSR register (bit [11:8]) to the value 0x0.

As a result of this operation, interrupts will be attended by the processor.

Return value:

This hypercall always succeed.

History:

Added in version 3.1.2

See also:

[XM_clear_irqmask \[Page: 12\]](#), [XM_set_irqmask \[Page: 55\]](#),

[XM_sparc_set_pil \[Page: 77\]](#).

2.46 XM_sparc_flush_cache

Synopsis: Flush data cache. Assembly hypercall.

Category:

Standard service.

Declaration:

```
void XM_sparc_flush_cache(void);
```

Description:

Invalidates the data cache. All the cache lines are marked as empty.

Return value:

This hypercall always succeed.

History:

Added in version 3.1.2

See also:

[XM_sparc_flush_tlb](#) [Page: 69]

2.47 XM_sparc_flush_regwin

Synopsis: Save the contents of the register window.

Category:

Standard service. Assembly hypercall.

Declaration:

```
void XM_sparc_flush_regwin(void);
```

Description:

Saves the contents of all the register windows but the active one when the hypercall is invoked.

This hypercall assumes that the stack registers contain valid (and correct) addresses.

Return value:

No value is returned.

Rationale:

This function is needed if the partition code has to implement threads.

A counterpart function to reload the register window is not needed because XtratuM manages transparently the stack. That is, it will be reloaded automatically when the register window underflows.

2.48 XM_sparc_flush_tlb

Synopsis: The TLB cache is invalidated. Assembly hypercall.

Category:

Standard service.

Declaration:

```
void XM_sparc_flush_tlb(void);
```

Description:

The TLB cache is invalidated.

Although this service is available on all processors (even on those without MMU), it is only effective on systems with MMU.

Return value:

This hypercall always succeed.

History:

Added in version 3.1.2

See also:

[XM_sparc_flush_cache](#) [Page: 67]

2.49 XM_sparc_get_flags

Synopsis: Replaced by `XM_sparc_get_psr()` on Sparc processors.

Category:

Standard service. Assembly hypercall.

Declaration:

```
xm_u32_t XM_sparc_get_flags(void);
```

Description:

Returns the four ICC (Integer Condition Codes) fields of the PSR (Processor Status Register). The bits 20 to 23 contains the ICC flags.

2.50 XM_sparc_get_psr

Synopsis: Get the ICC and PIL flags from the virtual PSR processor register.

Category:

Standard service. Assembly hypercall.

Declaration:

```
xm_u32_t XM_sparc_get_psr(void);
```

Description:

Returns ICC (Integer Condition Codes) fields and PIL (Processor Interrupt Level) fields of the PSR (Processor Status Register). The bits 20 to 23 contains the ICC flags; bits 8 to 11 contains the PIL.

Rationale:

This function replaces the deprecated XM_sparc_get_flags().

History:

Added in version 3.1.2

See also:

XM_sparc_set_psr [Page: 78].

2.51 XM_sparc_inport

Synopsis: Read from a hardware I/O port.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_sparc_inport(xm_u32_t portAddr,  
                        xm_u32_t *ptrValue);
```

Description:

This service provides direct hardware access to the peripherals. The content of the io hardware port `portAddr` is returned in the address pointed by `ptrValue`.

The configuration file lists the hardware ports that can be used by each partition. In the case of reading from a restricted port, only the mask attribute is applied to filter out the bits not allocated to the partition.

Return value:

[XM_OK]

The operation succeeded.

[XM_INVALID_PARAM]

- `portAddr` is not allocated to this partition,
- `portAddr` it is not aligned to 4 bytes,
- `ptrValue` is not a valid partition address.

Rationale:

Note that the SPARC v8 architecture does not define separate address spaces for memory and peripherals. This hypercall is implemented as a standard load instruction, with the corresponding security checks.

See also: `XM_sparc_outport` [Page: 74], `XM_memory_copy` [Page: 38]

2.52 XM_sparc_iret

Synopsis: Return from an interrupt.

Category:


Standard service. Assembly hypercall.

Declaration:

```
XM_sparc_iret();
```

Description:

Emulates a `rett` (RETurn from Trap) assembly instruction.

This service should be called at the end of an exception handler; otherwise the result is undetermined. The program counter of the partition will be corrupted. 

The ICC fields of the PSR are restored with the content of the register 10. 

The hypercall is provided as an assembly macro.

Return value:

The partition will continue with the normal execution flow, previous to the last attended trap.

Rationale:

When a trap is triggered, the current register window is decremented by one; and the processor register contain the following values:

- The register 10 contains a copy of the ICC flags.
- The register 11 contains the PC (program counter).
- The register 12 contains the nPC (next program counter).

Usage examples:

```
exception_handler_asm:
    call exception_handler
    nop
    XM_sparc_iret()

/* Code expanded in the trap table. */
#define BUILD_TRAP(trapnr) \
    sethi %hi(exception_handler_asm), %l4 ;\
    jmpl %l4 + %lo(exception_handler_asm), %g0 ;\
    mov trapnr, %o0 ;\
    nop
```

2.53 XM_sparc_outport

Synopsis: Write in a hardware I/O port.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_sparc_outport(xm_u32_t portAddr,
                          xm_u32_t value);
```

Description:

This service provides direct hardware access to the peripherals. On response to this call, XtratuM performs the write of value into the port portAddr on the native hardware.

The configuration file lists the hardware ports that can be used by each partition. There are two methods to allocate ports to a partition in the configuration file:

Range of ports:

A range of ports, **with no restriction**, allocated to the partition. The Range element is used.

Example:

```
<Partition ..... >
  <HwResources>
    <IoPorts>
      <Range base="0xc0000008" noPorts="2"/>
    </IoPorts>
  </HwResources>
</Partition>
```

In this example, the ports 0xc0000008 and 0xc0000009 are allocated to the partition.

The attributes base and noPorts are mandatory.

Restricted ports:

A single port with restrictions on the bits the the partition is allowed to write in. Only those bits that are set in the mask, can be modified by the partition. In order to implement this feature, XtratuM executed the code below:

```
oldValue=LoadIoReg(port);
StoreIoReg(port, ((oldValue&~(mask))|(value&mask)));
```

Listing 2.12: /core/kernel/sparcv8/hypercalls.c



Note that the port is read before it is finally written. **The read operation shall not cause side effects on the associated peripheral.** Some devices may interpret as interrupt acknowledge to read from a control port. Another source of error happen then the restricted is implemented as an open collector.

If the bitmask restriction is used, then the bits of the port that are not set in the mask can be allocated to other partitions. This way, it is possible to perform a fine grain (bit level) port allocation to partitions.

```
<Partition ..... >
  <HwResources>
    <IoPorts>
      <Restricted address="0x3000" mask="0xff" />
```

```
        </IoPorts>
    </HwResources>
</Partition>
```

Return value:

[XM.OK]

The value has been successfully written.

[XM.INVALID_PARAM]

- portAddr is not allocated to this partition,
- portAddr it is not aligned to 4 bytes,
- the port has been configured as a restricted port and value is out of range.

Rationale:

Note that the SPARC v8 architecture does not define separate address spaces for memory and peripherals. This hypercall is implemented as a standard store instruction, with the corresponding security checks.

See also: XM_sparc_inport [Page: 72], XM_memory_copy [Page: 38]

2.54 XM_sparc_set_flags

Synopsis: Set the ICC flags on the PSR processor register.

Category:

Standard service. Assembly hypercall.

Declaration:

```
void XM_sparc_set_flags(xm_u32_t flags);
```

Description:

Sets four ICC (Integer Condition Codes) fields of the PSR (Processor Status Register). The bits 20 to 23 of the flags parameter are copied into the ICC flags of the PSR processor register.

2.55 XM_sparc_set_pil

Synopsis: Set the PIL field of the PSR (disallow interrupts).

Category:

Standard service. Assembly hypercall.

Declaration:

```
void XM_sparc_set_pil(void);
```

Description:

Set the PIL field of the PSR register (bit [11:8]) to the value 0xF.

As a result of this operation, interrupts will not be attended by the processor. Note that interrupt level 15 is always attended (non-disable, non-mask-arable).

Return value:

This hypercall always succeed.

History:

Added in version 3.1.2

See also:

[XM_clear_irqmask \[Page: 12\]](#), [XM_set_irqmask \[Page: 55\]](#),

[XM_sparc_clear_pil \[Page: 66\]](#)

2.56 XM_sparc_set_psr

Synopsis: Set the ICC and PIL flags on the virtual PSR processor register.

Category:

Standard service. Assembly hypercall.

Declaration:

```
void XM_sparc_set_psr(xm_u32_t psr);
```

Description:

Sets ICC (Integer Condition Codes) fields and PIL (Processor Interrupt Level) fields of the PSR (Processor Status Register). The bits 20 to 23 contains the ICC flags; bits 8 to 11 contains the PIL.

Rationale:

This function replaces the deprecated `XM_sparc_set_flags()`.

History:

Added in version 3.1.2

See also:

`XM_sparc_get_psr` [Page: 71].

2.57 XM_suspend_partition

Synopsis: Suspend the execution of a partition.

Category:

Standard service/Supervisor reserved.

Declaration:

```
xm_s32_t XM_suspend_partition (xmId_t partitionId);
```

Description:

Suspends the execution of the target partition, `partitionId` until it will be resumed. If the target partition is in suspended state then this hypercall has no effect. The target partition is set in suspend state.

In suspend state, a partition will not be scheduled and interrupts are not delivered. Interrupts raised while in suspended state are left pending. If the partition returns to normal state then pending interrupts will be delivered then.

If a partition is in suspended state for a long period of time, some hardware peripherals may get unattended for a unacceptable amount of time which may cause improper peripheral operation.

Only supervisor partitions can suspend other partitions than itself. Any partition can halt itself. The constant `XM_PARTITION_SELF` represents the calling partition.



Return value:

[XM_OK]

The target partition has been suspended or it was already in suspend state.

[XM_INVALID_PARAM]

`partitionId` is not a valid partition identifier.

[XM_PERM_ERROR]

The calling partition is not supervisor and the target partition is not itself.

Usage examples:

```
XM_suspend_partition(XM_PARTITION_SELF);
```

See also:

`XM_reset_partition` [Page: 48], `XM_halt_partition` [Page: 28],

`XM_resume_partition` [Page: 51].

2.58 XM_trace_event

Synopsis: Records a trace entry.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_trace_event (xm_u32_t  bitmask,
                        xmTraceEvent_t *event);
```

Description:

This service records trace event pointed by `event` into the partition's trace stream if the logical and operation between the `bitmak` parameter and the `/XMHypervisor/PartitionTable/Partition/-Trace/@bitmask` attribute of the configuration file.

Each partition has its own trace stream to store the trace events generated by the partition. This trace stream buffer has to be configured in the `XM_CF` configuration file, and is automatically opened after a reset.

A health monitoring log entry is a data `xmTraceEvent_t` structure.

Where the `xmTraceOpCode_t` is a 32bits word with the following bit fields:

```
typedef struct {
    xm_u32_t code:13, criticality:3, moduleId:8, partitionId:8;
#define XM_TRACE_UNRECOVERABLE 0x3 // This level triggers a health
                                  // monitoring fault
#define XM_TRACE_WARNING 0x2
#define XM_TRACE_DEBUG 0x1
#define XM_TRACE_NOTIFY 0x0
} xmTraceOpCode_t;
```

Listing 2.13: `/core/include/objects/trace.h`

`partitionId`

Constains the identifier of the calling partition. It is filled by `XtratuM`.

`moduleId`

User defined field.

`criticality`

Defines the importance of the trace event.

`code`

An number which identifies cause of the trace.

The event shall be aligned to 8 bytes.

If the trace stream is stored in non-volatile memory, then the internal trace stream may become full. In this case, the oldest traced event is overwritten.

Return value:

`[XM_OK]`

The trace event has been stored successfully.

[XM_INVALID_PARAM]

event address is not a valid.

History:

Introduced in XtratuM 2.2.0.

See also:

[XM_trace_open](#) [Page: 82], [XM_trace_read](#) [Page: 83], [XM_trace_seek](#) [Page: 84], [XM_trace_status](#) [Page: 85].

2.59 XM_trace_open

Synopsis: Open a trace stream.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_trace_open (xmId_t id);
```

Description:

Returns the trace stream descriptor of the partition whose id is: `id`. If the `id` parameter is `XM_HYPERVISOR_ID`, then the trace stream of XtratuM is returned.

The trace events generated by a partition are stored in a trace stream with the same name than the partition that created the traces. The trace events generated by XtratuM are stored in the trace stream named "xm".

The first time the trace stream is opened, the read position is set to the oldest trace event.

Opening multiple times the same trace stream returns the same stream descriptor, and not a duplicate of it. Therefore, the read position is not changed as a result of an open operation.

The trace stream has to be allocated (in the XML configuration file) to the partition that opens it. Otherwise an error is returned. At most one partition can open a trace stream.



Return value:

Return a positive number which represents the trace stream descriptor. A negative number is returned in case of an error. The returned error values are:

[`XM_PERM_ERROR`]

The calling partition is not supervisor.

[`XM_INVALID_PARAM`]

The `id` is not a valid partition address, or a valid trace stream name.

History:

Introduced in XtratuM 2.2.0.

See also:

`XM_trace_event` [Page: 80], `XM_trace_read` [Page: 83], `XM_trace_seek` [Page: 84], `XM_trace_status` [Page: 85].

2.60 XM_trace_read

Synopsis: Read a trace event.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_trace_read (xm_s32_t  traceStream,  
                        xmTraceEvent_t *traceEventPtr);
```

Description:

Attempts to read one trace event from the stream `traceStream`. If no log entry is available, the call returns immediately.

If there were unread events in the stream, a positive value is returned **and the read position is advanced by one**. If there was no new trace events to be read, then a zero value is returned.

Note that this operation is not destructive. That is, the trace events are not removed from the internal stream once read. It is possible to retrieve old events with the hypercall `XM_trace_seek()`.

A health monitoring log entry is a data `xmTraceEvent_t` structure.

Return value:

[0] There are no new log entries in the trace stream. No data is returned.

[1]

The next trace entry is returned and the current read position is increased by one.

[XM_PERM_ERROR]

The calling partition is not supervisor.

[XM_INVALID_PARAM]

The `traceEventPtr` address is not a valid partition address; or `traceStream` is not a valid trace stream descriptor.

History:

Introduced in XtratuM 2.2.0.

See also:

`XM_trace_event` [Page: 80], `XM_trace_open` [Page: 82], `XM_trace_seek` [Page: 84], `XM_trace_status` [Page: 85].

2.61 XM_trace_seek

Synopsis: Sets the read position in a trace stream.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_trace_seek (xm_s32_t traceStream,
                       xm_s32_t offset,
                       xm_u32_t whence);
```

Description:

This hypercall repositions the read position of the open file associated with the trace stream descriptor `traceStream` to the argument `offset` according to the directive `whence` as follows:

[XM_OBJ_SEEK_SET]

The new read position is `offset` events from the oldest one stored in the given stream. Only positive values (or zero) of `offset` are valid.

[XM_OBJ_SEEK_CUR]

The new read position is set to the current read position plus the `offset` value. Negative values of `offset` refer to older (previous) trace events. Positive values refer to newer (posterior) events. If `offset` is zero then the current position is not changed.

[XM_OBJ_SEEK_END]

The new read position is set to the end of the stream. Negative values of `offset` refer to older (previous) events. If `offset` is zero then the next read operation will not return a trace event (unless a new event is recorded after the completion of this call).

Note that the `offset` value represents the number of trace events, and not the size in bytes.

`offset` has to be no greater than the size of the internal buffer stream.

If the computed new read position is beyond the beginning or the end of the stream, then it will be trunked to the start or the end of the stream respectively.



Since the stream is a circular buffer, it may happen that right after the `XM_trace_seek()` operation, one or more new trace events are recorded in the stream. In this case, the pointer may not be positioned at the start or the end of the stream, but a few entries after or before it. To avoid this concurrency issues, it is advisable not to move the pointer too close to the start of the buffer when the buffer is full.

Return value:

[XM_OK]

The operation succeeds.

[XM_PERM_ERROR]

The calling partition is not supervisor.

[XM_INVALID_PARAM]

The `offset` or `whence` are not valid. Or `traceStream` is not a valid trace stream descriptor.

History:

Introduced in XtratuM 2.2.0.

See also:

`XM_trace_event` [Page: 80], `XM_trace_open` [Page: 82], `XM_trace_read` [Page: 83], `XM_trace_status` [Page: 85].

2.62 XM_trace_status

Synopsis: Get the status of a trace stream.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_trace_status (xm_s32_t      traceStream,
                          xmTraceStatus_t *traceStatusPtr);
```

Description:

This hypercall returns information of the trace stream `traceStream` in the structure pointed by `traceStatusPtr`.

This service returns a structure which contains the following fields:

```
typedef struct {
    xm_s32_t noEvents;
    xm_s32_t maxEvents;
    xm_s32_t currentEvent;
} xmTraceStatus_t;
```

Listing 2.14: /core/include/objects/trace.h

Return value:

[XM_OK]

The operation succeeded.

[XM_PERM_ERROR]

The calling partition is not supervisor.

[XM_INVALID_PARAM]

The `traceStatusPtr` address is not a valid partition address; or `traceStream` is not a valid trace stream descriptor.

History:

Introduced in XtratuM 2.2.0.

See also:

`XM_trace_event` [Page: 80], `XM_trace_open` [Page: 82], `XM_trace_read` [Page: 83],

`XM_trace_seek` [Page: 84].

2.63 XM_unmask_irq

Synopsis: Obsoleted by `XM_clear_irqmask()`.

Description:

None

History:

Removed in version 3.1.2

See also:

`XM_clear_irqmask` [Page: 12].

2.64 XM_update_page32

Synopsis: Writes `val` in `pAddr`.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_update_page32(xmAddress_t pAddr, xm_u32_t val);
```

Description:

Pages declared as page table are marked as read-only, preventing the partition to freely update its content. Updating an entry in these pages requires the use of the hypercall `XM_update_page32`.

`pAddr` is the physical address to be updated, and `val` is the new value.

Before performing any change to the page, XtratuM validates that `val` does not breaks any security rule as mapping pages which do not belong to the partition, etc. If any rule is broken, the update is done successfully and the hypercall returns with `XM_OK`. Otherwise, `XM_INVALID_PARAM` is returned.

A partition can update the content of any page, no matter the type of the page by using this hypercall. The rules described above are only applied to the pages marked as page table ones.

Return value:

[`XM_OK`] The operation finished successfully.

[`XM_INVALID_PARAM`]The page `pAddr` does not belong to the partition or its content breaks some rule for its new type

Rationale:

Any partition has read-access to its own page tables avoiding the overhead of hypercall anytime a query to them. However, the page tables are marked as read-only, so updating any entry requires the partition to invoke `XM_update_page32`. This enable XtratuM to validate any update performed to the page tables.

Usage examples:

```
#include <xm_inc/arch/physmm.h/>
#include <xm.h/>

unsigned long pgd[1024] __attribute__((aligned(4096)));
unsigned long pgt[1024] __attribute__((aligned(4096)));
...

void PartitionMain(void) {
    memset(ptdL1, 0, 4096);
    memset(ptdL2, 0, 4096);
    if (XM_set_page_type(ptdL1, PPAG_PTD1)!=XM_OK) {
        xprintf("ptdL1 couldn't be set as PTD1\n");
        XM_halt_partition(XM_PARTITION_SELF);
    }
    if (XM_set_page_type(ptdL2, PPAG_PTD2)!=XM_OK) {
        xprintf("ptdL2 couldn't be set as PTD2\n");
        XM_halt_partition(XM_PARTITION_SELF);
    }
}
```

```
if (XM_update_page32(&ptdL1[0], (unsigned long)ptdL2|0x2)!=
    XM_OK) { // PTDL2 PRESENT
    xprintf("ptdL1[0] couldn't be updated\n");
    XM_halt_partition(XM_PARTITION_SELF);
}
if (XM_update_page32(&ptdL2[0], 0x40100000|0xe)!=XM_OK) { //
    PAGE PRESENT / USER_RW
    xprintf("ptdL2[0] couldn't be updated\n");
    XM_halt_partition(XM_PARTITION_SELF);
}

...
}
```

See also:

[XM_set_page_type](#) [Page: 57]

2.65 XM_write_console

Synopsis: Print a string in the hypervisor console.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_write_console (char *str, xm_s32_t length);
```

Description:

Writes up to `length` bytes from the buffer `str` to the default output console of XtratuM.

The target device where the messages are printed depends on the configuration of XtratuM. During the debugging phase, the XtratuM console is attached to a serial port.

This function is intended only for development and testing purposes, and should not be used in real operation.

The message is completely written on the output device before the function returns.

Return value:

This function always succeeds, returning the number of characters written.

Usage examples:

```
/* Initialization code */  
XM_write_console("Partition 2: Initialization succeed.\n", 37);
```

See also:

`XM_read_console` [Page: 42]

2.66 XM_write_register32

Synopsis: Modify a processor control register.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_write_register32(xm_s32_t register, xm_u32_t value)
;
```

Description:

This para-virtualises the processor. The parameter `register` identifies the processor control register to be modified with the new value. The result of this service depends on the target register.



A wrong value may have fatal consequences on the calling partition. Since this service operates in the virtual processor, no effects can be produced on other partitions or the hypervisor.

Below is the list of valid processor registers:

```
#define TBR_REG32 0
```

Listing 2.15: `/core/include/sparcv8/processor.h`

Return value:

[XM_OK]

The requested operation has been successful.

[XM_INVALID_PARAM]

- The `register` is not a valid register of this architecture.
- The given `value` is not a valid value for the specified register.

Usage examples:

```
XM_write_register32(reg,value);
```

2.67 XM_write_sampling_message

Synopsis: Writes a message in the specified sampling port.

Category:

Standard service.

Declaration:

```
xm_s32_t XM_write_sampling_message (xm_s32_t portDesc,  
                                   void      *msgPtr,  
                                   xm_u32_t  size);
```

Description:

The message is automatically copied into the internal XtratuM buffer of the channel.

If the `validPeriod` parameter is specified in the configuration file, then a timestamp is attached to the message when it is copied in the channel.

Return value:

[XM_OK]

If the message has been successfully written into internal buffer.

[XM_INVALID_PARAM]

`portDesc` does not identify an existing sampling port or it is not a destination port; or `size` is zero.

[XM_INVALID_CONFIG]

The size is not compatible with the configuration of the specified port.

Optimization:

If the buffer is aligned to 8 bytes, then the copy operation is performed faster. Also, the copy operation is more efficient if the length of the message is multiple of 8 bytes.

Usage examples:

```
xm_s32_t portDesc, retCode;  
  
char message[] = "This is a sample message";  
  
portDesc = XM_create_sampling_port ("sample",  
                                   30,  
                                   XM_PORT_DESTINATION);  
if (portDesc < 0) XM_halt_partition(XM_PARTITION_SELF);  
retCode = XM_write_sampling_message (portDesc,  
                                     message,  
                                     strlen(message));  
  
if (retCode != XM_OK)  
    XM_halt_partition(XM_PARTITION_SELF);
```

See also:

[XM_create_sampling_port](#) [Page: 15], [XM_read_sampling_message](#) [Page: 43],

[XM_get_sampling_port_status](#) [Page: 25], [XM_get_sampling_port_info](#) [Page: 24].