

GM Reference Manual
2.0.6

Generated by Doxygen 1.2.15

Tue Sep 9 08:26:21 2003

Contents

1	GM: A message-passing system for Myrinet networks	1
1.1	Table of Contents:	1
2	GM Module Index	3
2.1	GM Modules	3
3	GM Data Structure Index	5
3.1	GM Data Structures	5
4	GM File Index	7
4.1	GM File List	7
5	GM Page Index	11
5.1	GM Related Pages	11
6	GM Module Documentation	13
6.1	Deprecated GM API functions	13
7	GM Data Structure Documentation	15
7.1	gm_free_mdebug Union Reference	15
7.2	gm_free_page Struct Reference	16
7.3	gm_hash Struct Reference	17
7.4	gm_hash_entry Struct Reference	18
7.5	gm_hash_segment Struct Reference	19
7.6	gm_lookaside Struct Reference	20

7.7	gm_lookaside::gm_lookaside_segment_list Struct Reference	21
7.8	gm_lookaside_segment Struct Reference	22
7.9	gm_mark_reference Union Reference	23
7.10	gm_mark_set Struct Reference	24
7.11	gm_mdebug Struct Reference	25
7.12	gm_on_exit_record Struct Reference	26
7.13	gm_page_allocation_record Struct Reference	27
7.14	gm_remote_ptr_n_t Struct Reference	28
7.15	gm_s16_n_t Struct Reference	29
7.16	gm_s32_n_t Struct Reference	30
7.17	gm_s64_n_t Struct Reference	31
7.18	gm_s8_n_t Struct Reference	32
7.19	gm_s_e_context_t Struct Reference	33
7.20	gm_u16_n_t Struct Reference	34
7.21	gm_u32_n_t Struct Reference	35
7.22	gm_u64_n_t Struct Reference	36
7.23	gm_u8_n_t Struct Reference	37
7.24	gm_up_n_t Struct Reference	38
7.25	gm_zone Struct Reference	39
7.26	gm_zone_area Struct Reference	40
7.27	hash_entry Struct Reference	41
7.28	preallocated_record_chunk Struct Reference	42
8	GM File Documentation	43
8.1	gm.h File Reference	43
8.2	gm_abort.c File Reference	101
8.3	gm_alloc_pages.c File Reference	102
8.4	gm_alloc_send_token.c File Reference	104
8.5	gm_allow_remote_memory_access.c File Reference	105
8.6	gm_bcopy.c File Reference	106
8.7	gm_blocking_receive.c File Reference	107

8.8	gm_blocking_receive_no_spin.c File Reference	109
8.9	gm_bzero.c File Reference	111
8.10	gm_calloc.c File Reference	112
8.11	gm_close.c File Reference	113
8.12	gm_crc.c File Reference	115
8.13	gm_datagram_send.c File Reference	117
8.14	gm_datagram_send_4.c File Reference	119
8.15	gm_debug_buffers.c File Reference	121
8.16	gm_deregister.c File Reference	124
8.17	gm_directcopy.c File Reference	126
8.18	gm_directed_send.c File Reference	128
8.19	gm_dma_calloc.c File Reference	129
8.20	gm_dma_malloc.c File Reference	130
8.21	gm_drop_sends.c File Reference	132
8.22	gm_eprintf.c File Reference	134
8.23	gm_exit.c File Reference	135
8.24	gm_flush_alarm.c File Reference	136
8.25	gm_free.c File Reference	137
8.26	gm_free_send_token.c File Reference	138
8.27	gm_free_send_tokens.c File Reference	139
8.28	gm_get.c File Reference	140
8.29	gm_get_host_name.c File Reference	142
8.30	gm_get_mapper_unique_id.c File Reference	143
8.31	gm_get_node_id.c File Reference	144
8.32	gm_get_node_type.c File Reference	145
8.33	gm_get_port_id.c File Reference	146
8.34	gm_get_unique_board_id.c File Reference	147
8.35	gm_getpid.c File Reference	148
8.36	gm_handle_sent_tokens.c File Reference	149
8.37	gm_hash.c File Reference	150
8.38	gm_hex_dump.c File Reference	160

8.39	gm_host_name_to_node_id.c File Reference	161
8.40	gm_init.c File Reference	163
8.41	gm_isprint.c File Reference	165
8.42	gm_log2.c File Reference	166
8.43	gm_lookaside.c File Reference	168
8.44	gm_malloc.c File Reference	172
8.45	gm_mark.c File Reference	173
8.46	gm_max_length_for_size.c File Reference	178
8.47	gm_max_node_id.c File Reference	179
8.48	gm_max_node_id_in_use.c File Reference	180
8.49	gm_memcmp.c File Reference	182
8.50	gm_memorize_message.c File Reference	183
8.51	gm_memset.c File Reference	185
8.52	gm_min_message_size.c File Reference	186
8.53	gm_min_size_for_length.c File Reference	187
8.54	gm_mtu.c File Reference	188
8.55	gm_mutex.c File Reference	189
8.56	gm_next_event_peek.c File Reference	191
8.57	gm_node_id_to_host_name.c File Reference	192
8.58	gm_node_id_to_unique_id.c File Reference	194
8.59	gm_num_ports.c File Reference	195
8.60	gm_num_receive_tokens.c File Reference	196
8.61	gm_num_send_tokens.c File Reference	197
8.62	gm_on_exit.c File Reference	198
8.63	gm_open.c File Reference	200
8.64	gm_page_alloc.c File Reference	202
8.65	gm_perror.c File Reference	204
8.66	gm_printf.c File Reference	205
8.67	gm_provide_receive_buffer.c File Reference	206
8.68	gm_put.c File Reference	208
8.69	gm_rand.c File Reference	210

8.70	gm_rand_mod.c File Reference	212
8.71	gm_receive.c File Reference	213
8.72	gm_receive_pending.c File Reference	215
8.73	gm_register.c File Reference	216
8.74	gm_resume_sending.c File Reference	219
8.75	gm_send.c File Reference	221
8.76	gm_send_to_peer.c File Reference	223
8.77	gm_send_token_available.c File Reference	225
8.78	gm_set_acceptable_sizes.c File Reference	226
8.79	gm_set_alarm.c File Reference	228
8.80	gm_set_enable_nack_down.c File Reference	231
8.81	gm_simple_example.h File Reference	232
8.82	gm_sleep.c File Reference	233
8.83	gm_strcmp.c File Reference	234
8.84	gm_strdup.c File Reference	235
8.85	gm_strerror.c File Reference	236
8.86	gm_strlen.c File Reference	237
8.87	gm_strncasecmp.c File Reference	238
8.88	gm_strncmp.c File Reference	239
8.89	gm_strncpy.c File Reference	240
8.90	gm_ticks.c File Reference	241
8.91	gm_unique_id.c File Reference	242
8.92	gm_unique_id_to_node_id.c File Reference	243
8.93	gm_unknown.c File Reference	245
8.94	gm_zone.c File Reference	246
9	GM Page Documentation	251
9.1	XI. Alarms	251
9.2	VII. Page Allocation	252
9.3	XV. GM Constants, Macros, and Enumerated Types	253
9.4	I. Copyright Notice	254

9.5	II. About This Document	256
9.6	X. Endian Conversion	257
9.7	XIV. Example Programs	259
9.8	XII. High Availability Extensions	260
9.9	V. Initialization	263
9.10	IV. Programming Model	264
9.11	III. Overview	270
9.12	IX. Receiving Messages	275
9.13	VIII. Sending Messages	282
9.14	VI. Memory Setup	285
9.15	XVI. Function Summary	286
9.16	XIII. Utility Modules	294

Chapter 1

GM: A message-passing system for Myrinet networks

1.1 Table of Contents:

- [I. Copyright Notice](#)
 - [II. About This Document](#)
 - [III. Overview](#)
 - [1. Definitions](#)
 - [2. Notation](#)
 - [IV. Programming Model](#)
 - [1. GM Endpoints \(Ports\)](#)
 - [2. User Token Flow \(Sending\)](#)
 - [3. User Token Flow \(Receiving\)](#)
 - [V. Initialization](#)
 - [VI. Memory Setup](#)
 - [VII. Page Allocation](#)
-

- [VIII. Sending Messages](#)
- [IX. Receiving Messages](#)
- [X. Endian Conversion](#)
- [XI. Alarms](#)
- [XII. High Availability Extensions](#)
- [XIII. Utility Modules](#)
 - [1. CRC Functions](#)
 - [2. Hash Table](#)
 - [3. Lookaside List](#)
 - [4. Marks](#)
 - [5. Zones](#)
 - [6. Mutexes](#)
 - [7. Buffer Debugging](#)
- [XIV. Example Programs](#)
- [XV. GM Constants, Macros, and Enumerated Types](#)
- [XVI. Function Summary](#)

If difficulties are encountered, please consult the [FAQ](#) and all technical support questions should be directed to help@myri.com.

Chapter 2

GM Module Index

2.1 GM Modules

Here is a list of all modules:

Deprecated GM API functions	13
---------------------------------------	----

Chapter 3

GM Data Structure Index

3.1 GM Data Structures

Here are the data structures with brief descriptions:

gm_free_mdebug	15
gm_free_page	16
gm_hash	17
gm_hash_entry	18
gm_hash_segment	19
gm_lookaside	20
gm_lookaside::gm_lookaside_segment_list	21
gm_lookaside_segment	22
gm_mark_reference	23
gm_mark_set	24
gm_mdebug	25
gm_on_exit_record	26
gm_page_allocation_record	27
gm_remote_ptr_n_t	28
gm_s16_n_t	29
gm_s32_n_t	30
gm_s64_n_t	31
gm_s8_n_t	32
gm_s_e_context_t	33
gm_u16_n_t	34
gm_u32_n_t	35
gm_u64_n_t	36
gm_u8_n_t	37
gm_up_n_t	38
gm_zone	39

gm_zone_area	40
hash_entry	41
preallocated_record_chunk	42

Chapter 4

GM File Index

4.1 GM File List

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

gm.h	43
gm_abort.c	101
gm_alloc_pages.c	102
gm_alloc_send_token.c	104
gm_allow_remote_memory_access.c	105
gm_bcopy.c	106
gm_blocking_receive.c	107
gm_blocking_receive_no_spin.c	109
gm_bzero.c	111
gm_calloc.c	112
gm_close.c	113
gm_crc.c	115
gm_datagram_send.c	117
gm_datagram_send_4.c	119
gm_debug_buffers.c	121
gm_deregister.c	124
gm_directcopy.c	126
gm_directed_send.c	128
gm_dma_calloc.c	129
gm_dma_malloc.c	130
gm_drop_sends.c	132
gm_eprintf.c	134
gm_exit.c	135
gm_flush_alarm.c	136
gm_free.c	137

gm_free_send_token.c	138
gm_free_send_tokens.c	139
gm_get.c	140
gm_get_host_name.c	142
gm_get_mapper_unique_id.c	143
gm_get_node_id.c	144
gm_get_node_type.c	145
gm_get_port_id.c	146
gm_get_unique_board_id.c	147
gm_getpid.c	148
gm_handle_sent_tokens.c	149
gm_hash.c	150
gm_hex_dump.c	160
gm_host_name_to_node_id.c	161
gm_init.c	163
gm_isprint.c	165
gm_log2.c	166
gm_lookaside.c	168
gm_malloc.c	172
gm_mark.c	173
gm_max_length_for_size.c	178
gm_max_node_id.c	179
gm_max_node_id_in_use.c	180
gm_memcmp.c	182
gm_memorize_message.c	183
gm_memset.c	185
gm_min_message_size.c	186
gm_min_size_for_length.c	187
gm_mtu.c	188
gm_mutex.c	189
gm_next_event_peek.c	191
gm_node_id_to_host_name.c	192
gm_node_id_to_unique_id.c	194
gm_num_ports.c	195
gm_num_receive_tokens.c	196
gm_num_send_tokens.c	197
gm_on_exit.c	198
gm_open.c	200
gm_page_alloc.c	202
gm_perror.c	204
gm_printf.c	205
gm_provide_receive_buffer.c	206
gm_put.c	208
gm_rand.c	210
gm_rand_mod.c	212
gm_receive.c	213

gm_receive_pending.c	215
gm_register.c	216
gm_resume_sending.c	219
gm_send.c	221
gm_send_to_peer.c	223
gm_send_token_available.c	225
gm_set_acceptable_sizes.c	226
gm_set_alarm.c	228
gm_set_enable_nack_down.c	231
gm_simple_example.h	232
gm_sleep.c	233
gm_strcmp.c	234
gm_strdup.c	235
gm_strerror.c	236
gm_strlen.c	237
gm_strncasecmp.c	238
gm_strncmp.c	239
gm_strncpy.c	240
gm_ticks.c	241
gm_unique_id.c	242
gm_unique_id_to_node_id.c	243
gm_unknown.c	245
gm_zone.c	246

Chapter 5

GM Page Index

5.1 GM Related Pages

Here is a list of all related documentation pages:

Alarms	??
Page_Allocation	??
Constants	??
Copyright	??
Document	??
Endian_Conversion	??
Examples	??
High_Availability	??
Initialization	??
Programming_Model	??
Overview	??
Receiving_Messages	??
Sending_Messages	??
Memory_Setup	??
Summary	??
Utilities	??

Chapter 6

GM Module Documentation

6.1 Deprecated GM API functions

[gm_directed_send\(\)](#) is deprecated and included only for backward compatibility. Use [gm_directed_send_with_callback\(\)](#) instead.

See also:

[gm_directed_send_with_callback](#)

[gm_handle_sent_tokens\(\)](#) is deprecated and included only for backward compatibility. Use [gm_unknown\(\)](#) instead.

See also:

[gm_unknown](#)

[gm_provide_receive_buffer\(\)](#) is deprecated and included only for backward compatibility. Customers should instead use [gm_provide_receive_buffer_with_tag\(...,0\)](#).

See also:

[gm_provide_receive_buffer_with_tag](#)

[gm_send\(\)](#) is deprecated. Use [gm_send_with_callback\(\)](#) instead.

See also:

[gm_send_with_callback](#)

[gm_send_to_peer\(\)](#) is deprecated and included only for backward compatibility. Use [gm_send_to_peer_with_callback\(\)](#) instead.

See also:

[gm_send_to_peer_with_callback](#)

Chapter 7

GM Data Structure Documentation

7.1 gm_free_mdebug Union Reference

7.1.1 Detailed Description

Union to allow chaining of free mdebug records, without requiring an extra field in the mdebug record.

The documentation for this union was generated from the following file:

- gm_malloc_debug.c

7.2 gm_free_page Struct Reference

7.2.1 Detailed Description

A structure representing a free page. This structure is placed at the start of each free page, and used to link these free pages in a list.

The documentation for this struct was generated from the following file:

- [gm_page_alloc.c](#)

7.3 gm_hash Struct Reference

7.3.1 Detailed Description

The state of a GM hash table, referenced by the client only using opaque pointers.

The documentation for this struct was generated from the following file:

- [gm_hash.c](#)

7.4 gm_hash_entry Struct Reference

7.4.1 Detailed Description

A hash table entry.

The documentation for this struct was generated from the following file:

- [gm_hash.c](#)

7.5 gm_hash_segment Struct Reference

7.5.1 Detailed Description

Structure representing a segment of allocated hash table bins. To double the size of the hash table, we allocate a new segment with just enough bins to double the number of bins in the hash table and prepend it to the list of hash segments. This way, we don't have to double-buffer the hash table while growing it, and we can grow the table closer to the limits of available memory. While we sometimes have to walk the $O(\log(N))$ -segment list to find a bin, the average lookup only looks at $O(2)$ segments, so operations are still constant-average-time as expected for hash tables.

The documentation for this struct was generated from the following file:

- [gm_hash.c](#)

7.6 gm_lookaside Struct Reference

Data Fields

- [gm_lookaside::gm_lookaside_segment_list](#) `segment_list`

7.6.1 Detailed Description

State of a lookaside list, which is an very fast and space efficient memory allocator for fixed-sized buffers.

7.6.2 Field Documentation

7.6.2.1 struct [gm_lookaside::gm_lookaside_segment_list](#) [gm_lookaside::segment_list](#)

List of segments. Segments with free entries are at the front.

The documentation for this struct was generated from the following file:

- [gm_lookaside.c](#)

7.7 gm_lookaside::gm_lookaside_segment_list Struct Reference

7.7.1 Detailed Description

List of segments. Segments with free entries are at the front.

The documentation for this struct was generated from the following file:

- [gm_lookaside.c](#)

7.8 gm_lookaside_segment Struct Reference

7.8.1 Detailed Description

Structure holding a group of allocated lookaside entries. It is used primarily when entries are smaller than a page to allocate a page worth of entries at a time. This is important since some kernel memory allocators (Windows) round all allocations up to a page length, but this implementation allows us to be efficient for small allocations despite this.

The documentation for this struct was generated from the following file:

- [gm_lookaside.c](#)

7.9 gm_mark_reference Union Reference

7.9.1 Detailed Description

a reference to a mark, or a free entry. The reference is used to validate the mark. That mark is considered valid if and only if (set->reference[mark->tag] == mark).

The documentation for this union was generated from the following file:

- [gm_mark.c](#)

7.10 gm_mark_set Struct Reference

7.10.1 Detailed Description

The state of a mark set, the database that tracks the location and validity of marks.

The documentation for this struct was generated from the following file:

- [gm_mark.c](#)

7.11 gm_mdebug Struct Reference

7.11.1 Detailed Description

A record holding a pointer free function that an allocation should be used to free the allocation, and the top of the call stack when the memory was allocated. This allows us to report when the wrong function is used to free the memory, and to report the call stack of the allocation in this case or if a memory leak is detected.

The documentation for this struct was generated from the following file:

- gm_malloc_debug.c

7.12 gm_on_exit_record Struct Reference

7.12.1 Detailed Description

List element storing the details of a callback that should be called upon exit.

The documentation for this struct was generated from the following file:

- [gm_on_exit.c](#)

7.13 gm_page_allocation_record Struct Reference

7.13.1 Detailed Description

A record of a large buffer allocation used to allocate a bunch of pages. (We allocate pages in blocks of many pages, in case the OS's memory allocator does not give us a page-aligned allocation, and use only the aligned pages within this large buffer.

The documentation for this struct was generated from the following file:

- [gm_page_alloc.c](#)

7.14 gm_remote_ptr_n_t Struct Reference

```
#include <gm.h>
```

7.14.1 Detailed Description

A pointer to memory on a (potentially) remote machine. Such pointers are always 64-bits, since we don't know if the remote pointer is 32- or 64-bits.

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.15 gm_s16_n_t Struct Reference

```
#include <gm.h>
```

7.15.1 Detailed Description

A 16-bit signed value in network byte order.

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.16 gm_s32_n_t Struct Reference

```
#include <gm.h>
```

7.16.1 Detailed Description

A 32-bit signed value in network byte order.

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.17 gm_s64_n_t Struct Reference

```
#include <gm.h>
```

7.17.1 Detailed Description

A 64-bit signed value in network byte order.

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.18 gm_s8_n_t Struct Reference

```
#include <gm.h>
```

7.18.1 Detailed Description

A 8-bit signed value in network byte order. (Silly, I know.)

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.19 gm_s_e_context_t Struct Reference

7.19.1 Detailed Description

A structure used to store the state of this simple example program.

The documentation for this struct was generated from the following file:

- gm_simple_example_rcv.c

7.20 gm_u16_n_t Struct Reference

```
#include <gm.h>
```

7.20.1 Detailed Description

A 16-bit unsigned value in network byte order.

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.21 gm_u32_n_t Struct Reference

```
#include <gm.h>
```

7.21.1 Detailed Description

A 32-bit unsigned value in network byte order.

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.22 gm_u64_n_t Struct Reference

```
#include <gm.h>
```

7.22.1 Detailed Description

A 64-bit unsigned value in network byte order.

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.23 gm_u8_n_t Struct Reference

```
#include <gm.h>
```

7.23.1 Detailed Description

An byte value in network byte order. (Silly, I know.)

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.24 gm_up_n_t Struct Reference

```
#include <gm.h>
```

7.24.1 Detailed Description

A user-space pointer in network byte order. (On systems supporting multiple pointer sizes, this is large enough to store the largest pointer size, even if the process is using smaller pointers.)

The documentation for this struct was generated from the following file:

- [gm.h](#)

7.25 gm_zone Struct Reference

7.25.1 Detailed Description

State of a zone, which is a region of memory from which one can allocate memory using the zone allocation functions.

The documentation for this struct was generated from the following file:

- [gm_zone.c](#)

7.26 gm_zone_area Struct Reference

7.26.1 Detailed Description

Zones are divided into managed buffers called "areas", which may either represent free buffers or buffers holding user data.

The documentation for this struct was generated from the following file:

- [gm_zone.c](#)

7.27 hash_entry Struct Reference

7.27.1 Detailed Description

An entry in a custom hash table used to track memory allocations.

The documentation for this struct was generated from the following file:

- gm_malloc_debug.c

7.28 preallocated_record_chunk Struct Reference

7.28.1 Detailed Description

Holds a groups of preallocated malloc debugging records. We allocated these records in chunks to reduce the average allocation overhead, and for storage density. We allocate 2^N-1 records at a time because many `malloc()` implementations use memory most efficiently for allocations just smaller than 2^M bytes long.

The documentation for this struct was generated from the following file:

- `gm_malloc_debug.c`

Chapter 8

GM File Documentation

8.1 gm.h File Reference

Data Structures

- struct [gm_remote_ptr_n_t](#)
- struct [gm_s16_n_t](#)
- struct [gm_s32_n_t](#)
- struct [gm_s64_n_t](#)
- struct [gm_s8_n_t](#)
- struct [gm_u16_n_t](#)
- struct [gm_u32_n_t](#)
- struct [gm_u64_n_t](#)
- struct [gm_u8_n_t](#)
- struct [gm_up_n_t](#)

Defines

- #define [GM_API_VERSION_1_0](#) 0x100
 - #define [GM_API_VERSION_1_1](#) 0x101
 - #define [GM_API_VERSION_1_2](#) 0x102
 - #define [GM_API_VERSION_1_3](#) 0x103
 - #define [GM_API_VERSION_1_4](#) 0x104
 - #define [GM_API_VERSION_1_5](#) 0x105
 - #define [GM_API_VERSION_1_6](#) 0x106
-

- #define `GM_API_VERSION_2_0` 0x200
- #define `GM_API_VERSION_2_0_6` 0x20006
- #define `GM_API_VERSION` `GM_API_VERSION_2_0_6`
- #define `GM_MAX_HOST_NAME_LEN` 128
- #define `GM_MAX_PORT_NAME_LEN` 32
- #define `GM_NO_SUCH_NODE_ID` 0
- #define `GM_CPU_alpha` 0
- #define `GM_RDMA_GRANULARITY` 64
- #define `GM_MAX_DMA_GRANULARITY` 8
- #define `GM_STRUCT_CONTAINING`(type, field, field_instance) ((type *)((char *) (field_instance) - GM_OFFSETOF (type, field)))
- #define `GM_NUM_ELEM`(ar) (sizeof (ar) / sizeof (*ar))
- #define `GM_POWER_OF_TWO`(n) (!(n)&((n)-1))
- #define `GM_STRUCT_CONTAINING`(type, field, field_instance) ((type *)((char *) (field_instance) - GM_OFFSETOF (type, field)))
- #define `GM_NUM_ELEM`(ar) (sizeof (ar) / sizeof (*ar))
- #define `GM_POWER_OF_TWO`(n) (!(n)&((n)-1))

Typedefs

- typedef `gm_u64_t` `gm_remote_ptr_t`
- typedef enum `gm_status` `gm_status_t`
- typedef void(* `gm_send_completion_callback_t`)(struct `gm_port` *p, void *context, `gm_status_t` status)
- typedef `gm_u32_t` `gm_pid_t`

Enumerations

- enum `gm_status` { `GM_SUCCESS` = 0, `GM_FAILURE` = 1, `GM_INPUT_BUFFER_TOO_SMALL` = 2, `GM_OUTPUT_BUFFER_TOO_SMALL` = 3, `GM_TRY_AGAIN` = 4, `GM_BUSY` = 5, `GM_MEMORY_FAULT` = 6, `GM_INTERRUPTED` = 7, `GM_INVALID_PARAMETER` = 8, `GM_OUT_OF_MEMORY` = 9, `GM_INVALID_COMMAND` = 10, `GM_PERMISSION_DENIED` = 11, `GM_INTERNAL_ERROR` = 12, `GM_UNATTACHED` = 13, `GM_UNSUPPORTED_DEVICE` = 14, `GM_SEND_TIMED_OUT` = 15, `GM_SEND_REJECTED` = 16, `GM_SEND_TARGET_PORT_CLOSED` = 17, `GM_SEND_TARGET_NODE_UNREACHABLE` = 18, `GM_SEND_DROPPED` = 19, `GM_SEND_PORT_CLOSED` = 20, `GM_NODE_ID_NOT_YET_SET` = 21, `GM_STILL_SHUTTING_DOWN` = 22, `GM_CLONE_BUSY` = 23, `GM_NO_SUCH_DEVICE` = 24, `GM_ABORTED` = 25, `GM_INCOMPATIBLE_LIB_AND_DRIVER` = 26, `GM_UNTRANSLATED_SYSTEM_ERROR` = 27, `GM_ACCESS_DENIED` =

- ```

28, GM_NO_DRIVER_SUPPORT = 29, GM_PTE_REF_CNT_OVERFLOW
= 30, GM_NOT_SUPPORTED_IN_KERNEL = 31,
GM_NOT_SUPPORTED_ON_ARCH = 32, GM_NO_MATCH = 33,
GM_USER_ERROR = 34, GM_TIMED_OUT = 35, GM_DATA_CORRUPTED
= 36, GM_HARDWARE_FAULT = 37, GM_SEND_ORPHANED =
38, GM_MINOR_OVERFLOW = 39, GM_PAGE_TABLE_FULL =
40, GM_UC_ERROR = 41, GM_INVALID_PORT_NUMBER = 42,
GM_DEV_NOT_FOUND = 43, GM_FIRMWARE_NOT_RUNNING = 44,
GM_YP_NO_MATCH = 45 }

```
- enum `gm_priority` { `GM_LOW_PRIORITY` = 0, `GM_HIGH_PRIORITY` = 1, `GM_NUM_PRIORITIES` = 2 }
  - enum `gm_recv_event_type` { `GM_NO_RECV_EVENT` = 0, `GM_SENDS_FAILED_EVENT` = 1, `GM_ALARM_EVENT` = 2, `GM_RECV_EVENT` = 11, `GM_HIGH_RECV_EVENT` = 12, `GM_PEER_RECV_EVENT` = 13, `GM_HIGH_PEER_RECV_EVENT` = 14, `GM_FAST_RECV_EVENT` = 15, `GM_FAST_HIGH_RECV_EVENT` = 16, `GM_FAST_PEER_RECV_EVENT` = 17, `GM_FAST_HIGH_PEER_RECV_EVENT` = 18, `GM_NEW_SENDS_FAILED_EVENT` = 129, `GM_NEW_RECV_EVENT` = 139, `GM_NEW_FAST_RECV_EVENT` = 143, `_GM_NEW_PUT_NOTIFICATION_EVENT` = 149, `GM_NUM_RECV_EVENT_TYPES` }

## Functions

- GM\_ENTRY\_POINT void `gm_abort` (void)
- GM\_ENTRY\_POINT int `gm_alloc_send_token` (struct `gm_port` \*p, unsigned int priority)
- GM\_ENTRY\_POINT `gm_status_t` `gm_allow_remote_memory_access` (struct `gm_port` \*p)
- GM\_ENTRY\_POINT void `gm_bcopy` (const void \*from, void \*to, `gm_size_t` len)
- GM\_ENTRY\_POINT union `gm_recv_event` \* `gm_blocking_receive` (struct `gm_port` \*p)
- GM\_ENTRY\_POINT union `gm_recv_event` \* `gm_blocking_receive_no_spin` (struct `gm_port` \*p)
- GM\_ENTRY\_POINT void `gm_bzero` (void \*ptr, `gm_size_t` len)
- GM\_ENTRY\_POINT void \* `gm_calloc` (`gm_size_t` len, `gm_size_t` cnt)
- GM\_ENTRY\_POINT void `gm_cancel_alarm` (struct `gm_alarm` \*gm\_alarm)
- GM\_ENTRY\_POINT void `gm_close` (struct `gm_port` \*p)
- GM\_ENTRY\_POINT void `gm_datagram_send` (struct `gm_port` \*p, void \*message, unsigned int size, `gm_size_t` len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, `gm_send_completion_callback_t` callback, void \*context)

- GM\_ENTRY\_POINT void [gm\\_datagram\\_send\\_4](#) (struct gm\_port \*p, gm\_u32\_t message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_deregister\\_memory](#) (struct gm\_port \*p, void \*ptr, gm\_size\_t length)
- GM\_ENTRY\_POINT void [gm\\_directed\\_send\\_with\\_callback](#) (struct gm\_port \*p, void \*source\_buffer, [gm\\_remote\\_ptr\\_t](#) target\_buffer, gm\_size\_t len, enum [gm\\_priority](#) priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)
- GM\_ENTRY\_POINT void \* [gm\\_dma\\_alloc](#) (struct gm\_port \*p, gm\_size\_t count, gm\_size\_t length)
- GM\_ENTRY\_POINT void [gm\\_dma\\_free](#) (struct gm\_port \*p, void \*addr)
- GM\_ENTRY\_POINT void \* [gm\\_dma\\_malloc](#) (struct gm\_port \*p, gm\_size\_t length)
- GM\_ENTRY\_POINT void [gm\\_flush\\_alarm](#) (struct gm\_port \*p)
- GM\_ENTRY\_POINT void [gm\\_free](#) (void \*ptr)
- GM\_ENTRY\_POINT void [gm\\_free\\_send\\_token](#) (struct gm\_port \*p, unsigned int priority)
- GM\_ENTRY\_POINT void [gm\\_free\\_send\\_tokens](#) (struct gm\_port \*p, unsigned int priority, unsigned int count)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_host\\_name](#) (struct gm\_port \*port, char name[GM\_MAX\_HOST\_NAME\_LEN])
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_node\\_type](#) (struct gm\_port \*port, int \*node\_type)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_node\\_id](#) (struct gm\_port \*port, unsigned int \*n)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_unique\\_board\\_id](#) (struct gm\_port \*port, char unique[6])
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_mapper\\_unique\\_id](#) (struct gm\_port \*port, char unique[6])
- GM\_ENTRY\_POINT void [gm\\_hex\\_dump](#) (const void \*ptr, gm\_size\_t len)
- GM\_ENTRY\_POINT unsigned int [gm\\_host\\_name\\_to\\_node\\_id](#) (struct gm\_port \*port, char \*\_host\_name)
- GM\_ENTRY\_POINT void [gm\\_initialize\\_alarm](#) (struct gm\_alarm \*my\_alarm)
- GM\_ENTRY\_POINT int [gm\\_isprint](#) (int c)
- GM\_ENTRY\_POINT void \* [gm\\_malloc](#) (gm\_size\_t len)
- GM\_ENTRY\_POINT void \* [gm\\_page\\_alloc](#) (void)
- GM\_ENTRY\_POINT void [gm\\_page\\_free](#) (void \*addr)
- GM\_ENTRY\_POINT void \* [gm\\_alloc\\_pages](#) (gm\_size\_t len)
- GM\_ENTRY\_POINT void [gm\\_free\\_pages](#) (void \*addr, gm\_size\_t len)
- GM\_ENTRY\_POINT gm\_size\_t [gm\\_max\\_length\\_for\\_size](#) (unsigned int size)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_max\\_node\\_id](#) (struct gm\_port \*port, unsigned int \*n)

- GM\_ENTRY\_POINT int [gm\\_memcmp](#) (const void \*a, const void \*b, gm\_size\_t len)
- GM\_ENTRY\_POINT void \* [gm\\_memorize\\_message](#) (void \*message, void \*buffer, unsigned int len)
- GM\_ENTRY\_POINT unsigned int [gm\\_min\\_message\\_size](#) (struct gm\_port \*port)
- GM\_ENTRY\_POINT unsigned int [gm\\_min\\_size\\_for\\_length](#) (gm\_size\_t length)
- GM\_ENTRY\_POINT unsigned int [gm\\_mtu](#) (struct gm\_port \*port)
- GM\_ENTRY\_POINT char \* [gm\\_node\\_id\\_to\\_host\\_name](#) (struct gm\_port \*port, unsigned int node\_id)
- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_node\\_id\\_to\\_unique\\_id](#) (struct gm\_port \*port, unsigned int n, char unique[6])
- GM\_ENTRY\_POINT unsigned int [gm\\_num\\_ports](#) (struct gm\_port \*p)
- GM\_ENTRY\_POINT unsigned int [gm\\_num\\_send\\_tokens](#) (struct gm\_port \*p)
- GM\_ENTRY\_POINT unsigned int [gm\\_num\\_receive\\_tokens](#) (struct gm\_port \*p)
- GM\_ENTRY\_POINT unsigned int [gm\\_get\\_port\\_id](#) (struct gm\_port \*p)
- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_open](#) (struct gm\_port \*\*p, unsigned int unit, unsigned int port, const char \*port\_name, enum gm\_api\_version version)
- GM\_ENTRY\_POINT void [gm\\_provide\\_receive\\_buffer\\_with\\_tag](#) (struct gm\_port \*p, void \*ptr, unsigned int size, unsigned int priority, unsigned int tag)
- GM\_ENTRY\_POINT int [gm\\_receive\\_pending](#) (struct gm\_port \*p)
- GM\_ENTRY\_POINT int [gm\\_next\\_event\\_peek](#) (struct gm\_port \*p, gm\_u16\_t \*sender)
- GM\_ENTRY\_POINT union gm\_recv\_event \* [gm\\_receive](#) (struct gm\_port \*p)
- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_register\\_memory](#) (struct gm\_port \*p, void \*ptr, gm\_size\_t length)
- GM\_ENTRY\_POINT int [gm\\_send\\_token\\_available](#) (struct gm\_port \*p, unsigned int priority)
- GM\_ENTRY\_POINT void [gm\\_send\\_with\\_callback](#) (struct gm\_port \*p, void \*message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)
- GM\_ENTRY\_POINT void [gm\\_send\\_to\\_peer\\_with\\_callback](#) (struct gm\_port \*p, void \*message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)
- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_set\\_acceptable\\_sizes](#) (struct gm\_port \*p, enum [gm\\_priority](#) priority, gm\_size\_t mask)
- GM\_ENTRY\_POINT void [gm\\_set\\_alarm](#) (struct gm\_port \*p, struct gm\_alarm \*my\_alarm, gm\_u64\_t usecs, void(\*callback)(void \*), void \*context)
- GM\_ENTRY\_POINT gm\_size\_t [gm\\_strlen](#) (const char \*cptr)
- GM\_ENTRY\_POINT char \* [gm\\_strncpy](#) (char \*to, const char \*from, int len)
- GM\_ENTRY\_POINT int [gm\\_strcmp](#) (const char \*a, const char \*b)
- GM\_ENTRY\_POINT int [gm\\_strncmp](#) (const char \*a, const char \*b, int len)
- GM\_ENTRY\_POINT int [gm\\_strcasecmp](#) (const char \*a, const char \*b, int len)
- GM\_ENTRY\_POINT gm\_u64\_t [gm\\_ticks](#) (struct gm\_port \*port)

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_unique\\_id](#) (struct gm\_port \*port, char unique[6])
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_unique\\_id\\_to\\_node\\_id](#) (struct gm\_port \*port, char unique[6], unsigned int \*node\_id)
- GM\_ENTRY\_POINT void [gm\\_unknown](#) (struct gm\_port \*p, union gm\_recv\_event \*e)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_route](#) (struct gm\_port \*p, unsigned int node\_id, char \*route, unsigned int \*len)
- GM\_ENTRY\_POINT void [gm\\_dump\\_buffers](#) (void)
- GM\_ENTRY\_POINT void [gm\\_register\\_buffer](#) (void \*addr, int size)
- GM\_ENTRY\_POINT int [gm\\_unregister\\_buffer](#) (void \*addr, int size)
- GM\_ENTRY\_POINT struct [gm\\_lookaside](#) \* [gm\\_create\\_lookaside](#) (gm\_size\_t entry\_len, gm\_size\_t min\_entry\_cnt)
- GM\_ENTRY\_POINT void [gm\\_destroy\\_lookaside](#) (struct [gm\\_lookaside](#) \*l)
- GM\_ENTRY\_POINT void \* [gm\\_lookaside\\_alloc](#) (struct [gm\\_lookaside](#) \*l)
- GM\_ENTRY\_POINT void \* [gm\\_lookaside\\_zalloc](#) (struct [gm\\_lookaside](#) \*l)
- GM\_ENTRY\_POINT void [gm\\_lookaside\\_free](#) (void \*ptr)
- GM\_ENTRY\_POINT struct [gm\\_hash](#) \* [gm\\_create\\_hash](#) (long(\*gm\_user\_compare)(void \*key1, void \*key2), unsigned long(\*gm\_user\_hash)(void \*key1), gm\_size\_t key\_len, gm\_size\_t data\_len, gm\_size\_t gm\_min\_entries, int flags)
- GM\_ENTRY\_POINT void [gm\\_destroy\\_hash](#) (struct [gm\\_hash](#) \*h)
- GM\_ENTRY\_POINT void \* [gm\\_hash\\_remove](#) (struct [gm\\_hash](#) \*hash, void \*key)
- GM\_ENTRY\_POINT void \* [gm\\_hash\\_find](#) (struct [gm\\_hash](#) \*hash, void \*key)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_hash\\_insert](#) (struct [gm\\_hash](#) \*hash, void \*key, void \*datum)
- GM\_ENTRY\_POINT void [gm\\_hash\\_rekey](#) (struct [gm\\_hash](#) \*hash, void \*old\_key, void \*new\_key)
- GM\_ENTRY\_POINT long [gm\\_hash\\_compare\\_strings](#) (void \*key1, void \*key2)
- GM\_ENTRY\_POINT unsigned long [gm\\_hash\\_hash\\_string](#) (void \*key)
- GM\_ENTRY\_POINT long [gm\\_hash\\_compare\\_longs](#) (void \*key1, void \*key2)
- GM\_ENTRY\_POINT unsigned long [gm\\_hash\\_hash\\_long](#) (void \*key)
- GM\_ENTRY\_POINT long [gm\\_hash\\_compare\\_ints](#) (void \*key1, void \*key2)
- GM\_ENTRY\_POINT unsigned long [gm\\_hash\\_hash\\_int](#) (void \*key)
- GM\_ENTRY\_POINT long [gm\\_hash\\_compare\\_ptrs](#) (void \*key1, void \*key2)
- GM\_ENTRY\_POINT unsigned long [gm\\_hash\\_hash\\_ptr](#) (void \*key)
- GM\_ENTRY\_POINT unsigned long [gm\\_crc](#) (void \*ptr, gm\_size\_t len)
- GM\_ENTRY\_POINT unsigned long [gm\\_crc\\_str](#) (const char \*ptr)
- GM\_ENTRY\_POINT int [gm\\_rand](#) (void)
- GM\_ENTRY\_POINT void [gm\\_srand](#) (int seed)
- GM\_ENTRY\_POINT unsigned int [gm\\_rand\\_mod](#) (unsigned int modulus)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_init](#) (void)
- GM\_ENTRY\_POINT void [gm\\_finalize](#) (void)
- GM\_ENTRY\_POINT unsigned long [gm\\_log2\\_roundup](#) (unsigned long n)



- GM\_ENTRY\_POINT struct gm\_mutex \* gm\_create\_mutex (void)
- GM\_ENTRY\_POINT void gm\_destroy\_mutex (struct gm\_mutex \*mu)
- GM\_ENTRY\_POINT void gm\_mutex\_enter (struct gm\_mutex \*mu)
- GM\_ENTRY\_POINT void gm\_mutex\_exit (struct gm\_mutex \*mu)
- GM\_ENTRY\_POINT struct gm\_zone \* gm\_zone\_create\_zone (void \*base, gm\_size\_t len)
- GM\_ENTRY\_POINT void gm\_zone\_destroy\_zone (struct gm\_zone \*zone)
- GM\_ENTRY\_POINT void \* gm\_zone\_free (struct gm\_zone \*zone, void \*a)
- GM\_ENTRY\_POINT void \* gm\_zone\_malloc (struct gm\_zone \*zone, gm\_size\_t length)
- GM\_ENTRY\_POINT void \* gm\_zone\_calloc (struct gm\_zone \*zone, gm\_size\_t count, gm\_size\_t length)
- GM\_ENTRY\_POINT int gm\_zone\_addr\_in\_zone (struct gm\_zone \*zone, void \*p)
- GM\_ENTRY\_POINT void gm\_resume\_sending (struct gm\_port \*p, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, gm\_send\_completion\_callback\_t callback, void \*context)
- GM\_ENTRY\_POINT void gm\_drop\_sends (struct gm\_port \*port, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, gm\_send\_completion\_callback\_t callback, void \*context)
- GM\_ENTRY\_POINT gm\_pid\_t gm\_getpid (void)
- GM\_ENTRY\_POINT gm\_status\_t gm\_directcopy\_get (struct gm\_port \*p, void \*source\_addr, void \*target\_addr, gm\_size\_t length, unsigned int source\_instance\_id, unsigned int source\_port\_id)
- GM\_ENTRY\_POINT void gm\_perror (const char \*message, gm\_status\_t error)
- GM\_ENTRY\_POINT int gm\_sleep (unsigned seconds)
- GM\_ENTRY\_POINT void gm\_exit (gm\_status\_t status)
- GM\_ENTRY\_POINT int gm\_printf (const char \*format,...)
- GM\_ENTRY\_POINT char \* gm\_strerror (gm\_status\_t error)
- GM\_ENTRY\_POINT gm\_status\_t gm\_set\_enable\_nack\_down (struct gm\_port \*port, int flag)
- GM\_ENTRY\_POINT gm\_status\_t gm\_max\_node\_id\_in\_use (struct gm\_port \*port, unsigned int \*n)
- GM\_ENTRY\_POINT int gm\_eprintf (const char \*format,...)
- GM\_ENTRY\_POINT void \* gm\_memset (void \*s, int c, gm\_size\_t n)
- GM\_ENTRY\_POINT char \* gm\_strdup (const char \*)
- GM\_ENTRY\_POINT gm\_status\_t gm\_mark (struct gm\_mark\_set \*set, gm\_mark\_t \*m)
- GM\_ENTRY\_POINT int gm\_mark\_is\_valid (struct gm\_mark\_set \*set, gm\_mark\_t \*m)
- GM\_ENTRY\_POINT gm\_status\_t gm\_create\_mark\_set (struct gm\_mark\_set \*\*set, unsigned long init\_count)
- GM\_ENTRY\_POINT void gm\_destroy\_mark\_set (struct gm\_mark\_set \*set)
- GM\_ENTRY\_POINT void gm\_unmark (struct gm\_mark\_set \*set, gm\_mark\_t \*m)

- GM\_ENTRY\_POINT void [gm\\_unmark\\_all](#) (struct [gm\\_mark\\_set](#) \*set)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_on\\_exit](#) ([gm\\_on\\_exit\\_callback\\_t](#), void \*arg)
- GM\_ENTRY\_POINT void [gm\\_put](#) (struct [gm\\_port](#) \*p, void \*local\_buffer, [gm\\_remote\\_ptr\\_t](#) remote\_buffer, [gm\\_size\\_t](#) len, enum [gm\\_priority](#) priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_global\\_id\\_to\\_node\\_id](#) (struct [gm\\_port](#) \*port, unsigned int global\_id, unsigned int \*node\_id)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_node\\_id\\_to\\_global\\_id](#) (struct [gm\\_port](#) \*port, unsigned int node\_id, unsigned int \*global\_id)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_node\\_id\\_to\\_host\\_name\\_ex](#) (struct [gm\\_port](#) \*port, unsigned int timeout\_usecs, unsigned int node\_id, char(\*name)[GM\_MAX\_HOST\_NAME\_LEN+1])
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_host\\_name\\_to\\_node\\_id\\_ex](#) (struct [gm\\_port](#) \*port, unsigned int timeout\_usecs, const char \*host\_name, unsigned int \*node\_id)

## Variables

- GM\_ENTRY\_POINT const unsigned char [gm\\_log2\\_roundup\\_table](#) [257]

## 8.1.1 Detailed Description

The official GM API include file.

author: [glenn@myri.com](mailto:glenn@myri.com)

## 8.1.2 Define Documentation

### 8.1.2.1 #define GM\_API\_VERSION\_1\_0 0x100

Hex equivalent of GM\_API\_VERSION\_1\_0

### 8.1.2.2 #define GM\_API\_VERSION\_1\_1 0x101

Hex equivalent of GM\_API\_VERSION\_1\_1

### 8.1.2.3 #define GM\_API\_VERSION\_1\_2 0x102

Hex equivalent of GM\_API\_VERSION\_1\_2

**8.1.2.4 #define GM\_API\_VERSION\_1\_3 0x103**

Hex equivalent of GM\_API\_VERSION\_1\_3

**8.1.2.5 #define GM\_API\_VERSION\_1\_4 0x104**

Hex equivalent of GM\_API\_VERSION\_1\_4

**8.1.2.6 #define GM\_API\_VERSION\_1\_5 0x105**

Hex equivalent of GM\_API\_VERSION\_1\_5

**8.1.2.7 #define GM\_API\_VERSION\_1\_6 0x106**

Hex equivalent of GM\_API\_VERSION\_1\_6

**8.1.2.8 #define GM\_API\_VERSION\_2\_0 0x200**

Hex equivalent of GM\_API\_VERSION\_2\_0

**8.1.2.9 #define GM\_API\_VERSION\_2\_0\_6 0x20006**

Hex equivalent of GM\_API\_VERSION\_2\_0\_6

**8.1.2.10 #define GM\_API\_VERSION GM\_API\_VERSION\_2\_0\_6**

Set the default API version used in this file.

**8.1.2.11 #define GM\_MAX\_HOST\_NAME\_LEN 128**

Maximum length of GM host name

**8.1.2.12 #define GM\_MAX\_PORT\_NAME\_LEN 32**

Maximum length of GM port name

**8.1.2.13 #define GM\_NO\_SUCH\_NODE\_ID 0**

No such GM node id

**8.1.2.14 #define GM\_CPU\_alpha 0**

Define all undefined GM\_CPU switches to 0 to prevent problems with "gcc -Wundef"

**8.1.2.15 #define GM\_RDMA\_GRANULARITY 64**

GM RDMA GRANULARITY

**8.1.2.16 #define GM\_MAX\_DMA\_GRANULARITY 8**

GM MAX DMA GRANULARITY

**8.1.2.17 #define GM\_STRUCT\_CONTAINING(type, field, field\_instance) ((type \*)((char \*)field\_instance) - GM\_OFFSETOF (type, field))**

Given a pointer to an instance of a field in a structure of a certain type, return a pointer to the containing structure.

**8.1.2.18 #define GM\_NUM\_ELEM(ar) (sizeof (ar) / sizeof (\*ar))**

Given an array, return the number of elements in the array.

**8.1.2.19 #define GM\_POWER\_OF\_TWO(n) (!(n)&((n)-1))**

Return nonzero if the input is neither a power of two nor zero. Otherwise, return zero.

**8.1.2.20 #define GM\_STRUCT\_CONTAINING(type, field, field\_instance) ((type \*)((char \*)field\_instance) - GM\_OFFSETOF (type, field))**

Given a pointer to an instance of a field in a structure of a certain type, return a pointer to the containing structure.

**8.1.2.21 #define GM\_NUM\_ELEM(ar) (sizeof (ar) / sizeof (\*ar))**

Given an array, return the number of elements in the array.

**8.1.2.22 #define GM\_POWER\_OF\_TWO(n) (!(n)&((n)-1))**

Return nonzero if the input is neither a power of two nor zero. Otherwise, return zero.

### 8.1.3 Typedef Documentation

#### 8.1.3.1 typedef gm\_u64\_t gm\_remote\_ptr\_t

The "gm\_up\_t" is a LANai-compatible representation of a user virtual memory address. If the host has 32-bit pointers, then a gm\_up\_t has 32 bits. If the host has 64-bit pointers, but only 32-bits of these pointers are used, then a gm\_up\_t still has 32 bits (for performance reasons). Finally, if the host has 64-bit pointers and more than 32 bits of the pointer is used, then a gm\_up\_t has 64-bits.

#### 8.1.3.2 typedef enum gm\_status gm\_status\_t

GM Send Completion Status codes

#### 8.1.3.3 typedef void(\* gm\_send\_completion\_callback\_t)(struct gm\_port \* p, void \*context, gm\_status\_t status)

gm\_send\_completion\_callback\_t typedef function.

#### 8.1.3.4 typedef gm\_u32\_t gm\_pid\_t

typedef for gm\_pid\_t.

### 8.1.4 Enumeration Type Documentation

#### 8.1.4.1 enum gm\_status

GM Send Completion Status codes

**Enumeration values:**

**GM\_SUCCESS** The send succeeded. This status code does not indicate an error.

**GM\_FAILURE** Operation Failed

**GM\_INPUT\_BUFFER\_TOO\_SMALL** Input buffer is too small

**GM\_OUTPUT\_BUFFER\_TOO\_SMALL** Output buffer is too small

**GM\_TRY\_AGAIN** Try Again

- GM\_BUSY** GM Port is Busy
- GM\_MEMORY\_FAULT** Memory Fault
- GM\_INTERRUPTED** Interrupted
- GM\_INVALID\_PARAMETER** Invalid input parameter
- GM\_OUT\_OF\_MEMORY** Out of Memory
- GM\_INVALID\_COMMAND** Invalid Command
- GM\_PERMISSION\_DENIED** Permission Denied
- GM\_INTERNAL\_ERROR** Internal Error
- GM\_UNATTACHED** Unattached
- GM\_UNSUPPORTED\_DEVICE** Unsupported Device
- GM\_SEND\_TIMED\_OUT** The target port is open and responsive and the message is of an acceptable size, but the receiver failed to provide a matching receive buffer within the timeout period. This error can be caused by the receive neglecting its responsibility to provide receive buffers in a timely fashion or crashing. It can also be caused by severe congestion at the receiving node where many senders are contending for the same receive buffers on the target port for an extended period. This error indicates a programming error in the client software.
- GM\_SEND\_REJECTED** The receiver indicated (in a call to [gm\\_set\\_acceptable\\_sizes\(\)](#)) the size of the message was unacceptable. This error indicates a programming error in the client software.
- GM\_SEND\_TARGET\_PORT\_CLOSED** The message cannot be delivered because the destination port has been closed.
- GM\_SEND\_TARGET\_NODE\_UNREACHABLE** The target node could not be reached over the Myrinet. This error can be caused by the network becoming disconnected for too long, the remote node being powered off, or by network links being rearranged when the Myrinet mapper is not running.
- GM\_SEND\_DROPPED** The send was dropped at the client's request. (The client called [gm\\_drop\\_sends\(\)](#).) This status code does not indicate an error.
- GM\_SEND\_PORT\_CLOSED** Clients should never see this internal error code.
- GM\_NODE\_ID\_NOT\_YET\_SET** Node ID is not yet set
- GM\_STILL\_SHUTTING\_DOWN** GM Port is still shutting down
- GM\_CLONE\_BUSY** GM Clone Busy
- GM\_NO\_SUCH\_DEVICE** No such device
- GM\_ABORTED** Aborted.
- GM\_INCOMPATIBLE\_LIB\_AND\_DRIVER** Incompatible GM library and driver
- GM\_UNTRANSLATED\_SYSTEM\_ERROR** Untranslated System Error

**GM\_ACCESS\_DENIED** Access Denied  
**GM\_NO\_DRIVER\_SUPPORT** No Driver Support  
**GM\_PTE\_REF\_CNT\_OVERFLOW** PTE Ref Cnt Overflow  
**GM\_NOT\_SUPPORTED\_IN\_KERNEL** Not supported in the kernel  
**GM\_NOT\_SUPPORTED\_ON\_ARCH** Not supported for this architecture  
**GM\_NO\_MATCH** No match  
**GM\_USER\_ERROR** User error  
**GM\_TIMED\_OUT** Timed out  
**GM\_DATA\_CORRUPTED** Data has been corrupted  
**GM\_HARDWARE\_FAULT** Hardware fault  
**GM\_SEND\_ORPHANED** Send orphaned  
**GM\_MINOR\_OVERFLOW** Minor overflow  
**GM\_PAGE\_TABLE\_FULL** Page Table is Full  
**GM\_UC\_ERROR** UC Error  
**GM\_INVALID\_PORT\_NUMBER** Invalid Port Number  
**GM\_DEV\_NOT\_FOUND** No device files found  
**GM\_FIRMWARE\_NOT\_RUNNING** Lanai not running  
**GM\_YP\_NO\_MATCH** No match for yellow pages query.

#### 8.1.4.2 enum gm\_priority

Priority Levels

##### Enumeration values:

**GM\_LOW\_PRIORITY** Low priority message  
**GM\_HIGH\_PRIORITY** High priority message  
**GM\_NUM\_PRIORITIES** Number of priority types

#### 8.1.4.3 enum gm\_recv\_event\_type

Receive Event Types

##### Enumeration values:

**GM\_NO\_RECV\_EVENT** No significant receive event is pending.

**GM\_SENDS\_FAILED\_EVENT** deprecated

**GM\_ALARM\_EVENT** This event should be treated as an unknown event (passed to [gm\\_unknown\(\)](#))

**GM\_RECV\_EVENT** This event indicates that a normal receive (GM\_LOW\_PRIORITY) has occurred.

**GM\_HIGH\_RECV\_EVENT** This event indicates that a normal receive (GM\_HIGH\_PRIORITY) has occurred.

**GM\_PEER\_RECV\_EVENT** This event indicates that a normal receive (GM\_LOW\_PRIORITY) has occurred, and the PEER indicates that the sender/receiver ports are the same.

**GM\_HIGH\_PEER\_RECV\_EVENT** This event indicates that a normal receive (GM\_HIGH\_PRIORITY) has occurred, and the PEER indicates that the sender/receiver ports are the same.

**GM\_FAST\_RECV\_EVENT** A small-message receive occurred (GM\_LOW\_PRIORITY) with the small message stored in the receive queue for improved small-message performance.

**GM\_FAST\_HIGH\_RECV\_EVENT** A small-message receive occurred (GM\_HIGH\_PRIORITY) with the small message stored in the receive queue for improved small-message performance.

**GM\_FAST\_PEER\_RECV\_EVENT** A small-message receive occurred (GM\_LOW\_PRIORITY) with the small message stored in the receive queue for improved small-message performance. The PEER indicates that the sender/receiver ports are the same.

**GM\_FAST\_HIGH\_PEER\_RECV\_EVENT** A small-message receive occurred (GM\_HIGH\_PRIORITY) with the small message stored in the receive queue for improved small-message performance. The PEER indicates that the sender/receiver ports are the same.

**GM\_NEW\_SENDS\_FAILED\_EVENT** deprecated

**GM\_NEW\_RECV\_EVENT** normal receives

**GM\_NEW\_FAST\_RECV\_EVENT** streamlined small message receives

**\_GM\_NEW\_PUT\_NOTIFICATION\_EVENT** Directed send notification

**GM\_NUM\_RECV\_EVENT\_TYPES** DO NOT add new types here.

## 8.1.5 Function Documentation

### 8.1.5.1 GM\_ENTRY\_POINT void gm\_abort (void)

[gm\\_abort\(\)](#) aborts the current process, and is a wrapper around the system function abort().



**See also:**

[gm\\_init](#) [gm\\_open](#) [gm\\_close](#) [gm\\_exit](#) [gm\\_finalize](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.2 GM\_ENTRY\_POINT int gm\_alloc\_send\_token (struct gm\_port \* p, unsigned int priority)**

Allocates a send token ([Details](#)).

**8.1.5.3 GM\_ENTRY\_POINT gm\_status\_t gm\_allow\_remote\_memory\_access (struct gm\_port \* port)**

[gm\\_allow\\_remote\\_memory\\_access\(\)](#) allows any remote GM port to modify the contents of any GM DMAable memory using the [gm\\_directed\\_send\(\)](#) function. This is a significant security hole, but is very useful on tightly coupled clusters on trusted networks.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

**Parameters:**

*port* (IN) Handle to the GM port.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.4 GM\_ENTRY\_POINT void gm\_bcopy (const void \* from, void \* to, gm\_size\_t len)**

[gm\\_bcopy\(\)](#) copies len bytes starting at **from** to location **to**. This function does not handle overlapping regions.

**Parameters:**

*from* (IN) The starting location for the region to be copied.

*to* (IN) The ending location for the region to be copied.

*len* (IN) The length in bytes of the region to be copied.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.5 GM\_ENTRY\_POINT union gm\_recv\_event\* gm\_blocking\_receive (struct gm\_port \* p)**

Blocks until there is a receive event and then returns a pointer to the event. ([Details](#)).

**8.1.5.6 GM\_ENTRY\_POINT union gm\_recv\_event\* gm\_blocking\_receive\_no\_spin (struct gm\_port \* p)**

Like [gm\\_blocking\\_receive\(\)](#), except it sleeps the current thread immediately if no receive is pending. ([gm\\_blocking\\_receive\\_no\\_spin](#) "Details").

**8.1.5.7 GM\_ENTRY\_POINT void gm\_bzero (void \* ptr, gm\_size\_t len)**

[gm\\_bzero\(\)](#) clears the **len** bytes of memory starting at **ptr**. This function does not use partword I/O unless it must (for speed, especially when doing PIO), and does not rely on the system `bzero()` functionality, which may not be safe for PIO mapped memory.

**Parameters:**

*ptr* (IN) The pointer to the memory location.

*len* (IN) The number of bytes of memory to be bzero'ed.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.8 GM\_ENTRY\_POINT void\* gm\_malloc (gm\_size\_t len, gm\_size\_t cnt)**

[gm\\_malloc\(\)](#) allocates and clears an array of **cnt** elements of length **len**.

**Parameters:**

*len* (IN) The number of bytes in each element.

*cnt* (IN) The number of elements in the array.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.9 GM\_ENTRY\_POINT void gm\_cancel\_alarm (struct gm\_alarm \* gm\_alarm)**

Cancels alarm ([Details](#)).

**8.1.5.10 GM\_ENTRY\_POINT void gm\_close (struct gm\_port \* p)**

Closes a GM port ([Details](#)).

**8.1.5.11 GM\_ENTRY\_POINT void gm\_datagram\_send (struct gm\_port \* p, void \* message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, gm\_send\_completion\_callback\_t callback, void \* context)**

Unreliable send ([Details](#)).

**8.1.5.12 GM\_ENTRY\_POINT void gm\_datagram\_send\_4 (struct gm\_port \* p, gm\_u32\_t message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, gm\_send\_completion\_callback\_t callback, void \* context)**

Unreliable send of gm\_u32\_t messages ([Details](#)).

**8.1.5.13 GM\_ENTRY\_POINT gm\_status\_t gm\_deregister\_memory (struct gm\_port \* p, void \* ptr, gm\_size\_t length)**

Deregisters memory ([Details](#)).

**8.1.5.14** **GM\_ENTRY\_POINT** void **gm\_directed\_send\_with\_callback** (struct **gm\_port** \* *p*, void \* *source\_buffer*, **gm\_remote\_ptr\_t** *target\_buffer*, **gm\_size\_t** *len*, enum **gm\_priority** *priority*, unsigned int *target\_node\_id*, unsigned int *target\_port\_id*, **gm\_send\_completion\_callback\_t** *callback*, void \* *context*)

Directed send (PUT) ([Details](#)).

**8.1.5.15** **GM\_ENTRY\_POINT** void\* **gm\_dma\_calloc** (struct **gm\_port** \* *p*, **gm\_size\_t** *count*, **gm\_size\_t** *length*)

**gm\_dma\_calloc**() allocates and clears **count** \* **length** bytes of DMAable memory aligned on a 4-byte boundary. Memory should be freed using **gm\_dma\_free**().

**Parameters:**

*p* (IN) Handle to the GM port.

*count* (IN) The number of elements to be calloc'ed.

*length* (IN) The size of each element to be calloc'ed.

**See also:**

[gm\\_dma\\_malloc](#) [gm\\_dma\\_free](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.16** **GM\_ENTRY\_POINT** void **gm\_dma\_free** (struct **gm\_port** \* *p*, void \* *addr*)

Frees a region of DMAable memory ([Details](#)).

**8.1.5.17** **GM\_ENTRY\_POINT** void\* **gm\_dma\_malloc** (struct **gm\_port** \* *p*, **gm\_size\_t** *length*)

**gm\_dma\_malloc**() allocates **length** bytes of DMAable memory aligned on a 4-byte boundary. Memory should be freed using **gm\_dma\_free**().

**Parameters:**

*p* (IN) Handle to the GM port.

*length* (IN) The number of bytes to be malloc'ed.

**See also:**

[gm\\_dma\\_malloc](#) [gm\\_dma\\_free](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.18 GM\_ENTRY\_POINT void gm\_flush\_alarm (struct gm\_port \* p)**

Flushes an alarm ([Details](#)).

**8.1.5.19 GM\_ENTRY\_POINT void gm\_free (void \* ptr)**

[gm\\_free\(\)](#) frees the memory buffer at **ptr**, which was previously allocated by [gm\\_malloc\(\)](#), or [gm\\_malloc\(\)](#).

**Parameters:**

*ptr* (IN) Address of the memory to be freed.

**See also:**

[gm\\_malloc](#) [gm\\_malloc](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.20 GM\_ENTRY\_POINT void gm\_free\_send\_token (struct gm\_port \* p,  
unsigned int priority)**

Frees a send token ([Details](#)).

**8.1.5.21 GM\_ENTRY\_POINT void gm\_free\_send\_tokens (struct gm\_port \* p,  
unsigned int priority, unsigned int count)**

Frees multiple send tokens ([Details](#)).

**8.1.5.22** GM\_ENTRY\_POINT [gm\\_status.t](#) `gm_get_host_name` (struct gm\_port \* *port*, char *name*[GM\_MAX\_HOST\_NAME\_LEN])

Copies the host name of the local node ([Details](#)).

**8.1.5.23** GM\_ENTRY\_POINT [gm\\_status.t](#) `gm_get_node_type` (struct gm\_port \* *port*, int \* *node\_type*)

Returns GM\_GET\_NODE\_TYPE ([Details](#)).

**8.1.5.24** GM\_ENTRY\_POINT [gm\\_status.t](#) `gm_get_node_id` (struct gm\_port \* *port*, unsigned int \* *n*)

Copies the GM ID of the interface ([Details](#)).

**8.1.5.25** GM\_ENTRY\_POINT [gm\\_status.t](#) `gm_get_unique_board_id` (struct gm\_port \* *port*, char *unique*[6])

Copies the board ID of the interface ([gm\\_get\\_unique\\_board\\_id](#) "Details").

**8.1.5.26** GM\_ENTRY\_POINT [gm\\_status.t](#) `gm_get_mapper_unique_id` (struct gm\_port \* *port*, char *unique*[6])

Copies copies the 6-byte ethernet address of the interface ([Details](#)).

**8.1.5.27** GM\_ENTRY\_POINT void `gm_hex_dump` (const void \* *ptr*, gm\_size\_t *len*)

[gm\\_hex\\_dump\(\)](#) prints the hex equivalent of data at `ptr`.

**Parameters:**

`ptr` (IN) Address of the data.

`len` (IN) The length (in bytes) of the data.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.28 GM\_ENTRY\_POINT unsigned int gm\_host\_name\_to\_node\_id (struct gm\_port \* *port*, char \* *host\_name*)**

This function is deprecated. Use [gm\\_host\\_name\\_to\\_node\\_id\\_ex\(\)](#) instead. Returns the GM ID associated with a *host\_name* ([Details](#)).

**8.1.5.29 GM\_ENTRY\_POINT void gm\_initialize\_alarm (struct gm\_alarm \* *my\_alarm*)**

Initializes user-allocated storage for an alarm. ([Details](#)).

**8.1.5.30 GM\_ENTRY\_POINT int gm\_isprint (int *c*)**

[gm\\_isprint\(\)](#) is just like ANSI `isprint()`, only it works in the kernel and MCP.

**Return values:**

*int* The return value is nonzero if the character is a printable character including a space, and a zero value if not.

**Parameters:**

*c* (OUT) Printable character.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.31 GM\_ENTRY\_POINT void\* gm\_malloc (gm\_size\_t *len*)**

[gm\\_malloc\(\)](#) returns a pointer to the specified amount of memory. In the kernel, the memory will be nonpageable.

**Return values:**

*ptr* Pointer to the specified amount of memory.

**Parameters:**

*len* (IN) The length in bytes to be malloc'ed.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.32 GM\_ENTRY\_POINT void\* gm\_page\_alloc (void)**

[gm\\_page\\_alloc\(\)](#) returns a ptr to a newly allocated page-aligned buffer of length GM\_PAGE\_LEN.

**Return values:**

*ptr* Page-aligned buffer of length GM\_PAGE\_LEN.

**See also:**

[gm\\_page\\_free](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.33 GM\_ENTRY\_POINT void gm\_page\_free (void \* ptr)**

[gm\\_page\\_free\(\)](#) frees the page of memory at **ptr** previously allocated by [gm\\_page\\_alloc\(\)](#). If all pages have been freed, free all of the memory allocated for pages.

**Parameters:**

*ptr* (IN) Address of the memory page to be freed.

**See also:**

[gm\\_page\\_alloc](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))



**8.1.5.34 GM\_ENTRY\_POINT void\* gm\_alloc\_pages (gm\_size\_t alloc\_len)**

[gm\\_alloc\\_pages\(\)](#) allocates a page-aligned buffer of length ALLOC\_LEN, where ALLOC\_LEN is a multiple of GM\_PAGE\_LEN. Any fractional page following the buffer is wasted.

**Return values:**

*ptr* Pointer to the allocated buffer.

0 Error occurred.

**Parameters:**

*alloc\_len* (IN) The length of buffer to be allocated.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.35 GM\_ENTRY\_POINT void gm\_free\_pages (void \* addr, gm\_size\_t alloc\_len)**

[gm\\_free\\_pages\(\)](#) frees the pages at **addr**, which were previously allocated with [gm\\_alloc\\_pages\(\)](#).

**Parameters:**

*addr* (IN) The address of the buffer to be freed.

*alloc\_len* (IN) The length (in bytes) of the buffer to be freed.

**See also:**

[gm\\_alloc\\_pages](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.36 GM\_ENTRY\_POINT gm\_size\_t gm\_max\_length\_for\_size (unsigned int size)**

[gm\\_max\\_length\\_for\\_size\(\)](#) returns the maximum length of a message that will fit in a GM buffer of size **size**.

**Return values:**

*gm\_size\_t*

**Parameters:**

*size* (IN) The size of the GM buffer.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.37 GM\_ENTRY\_POINT gm\_status\_t gm\_max\_node\_id (struct gm\_port \* port, unsigned int \* n)**

Stores the maximum GM node ID supported by the network interface card ([Details](#)).

**8.1.5.38 GM\_ENTRY\_POINT int gm\_memcmp (const void \* *a*, const void \* *b*, gm\_size\_t *len*)**

[gm\\_memcmp\(\)](#) emulates the ANSI memcmp() function.

**Return values:**

*int* Returns an integer less than, equal to, or greater than zero if the first *len* bytes of *a* is found, respectively, to be less than, to match, or be greater than the first *len* bytes of *b*.

**Parameters:**

*a* (IN) The first memory area for comparison.

*b* (IN) The second memory area for comparison.

*len* (IN) The number of bytes to compare.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.39 GM\_ENTRY\_POINT void\* gm\_memorize\_message (void \* message, void \* buffer, unsigned len)**

[gm\\_memorize\\_message\(\)](#) is a wrapper around function `memcpy()`. [gm\\_memorize\\_message\(\)](#) copies a message into a buffer if needed. If **message** and **buffer** differ, `gm_memorize_message(port,message,buffer)` copies the message pointed to by **message** into the buffer pointed to by **buffer**. [gm\\_memorize\\_message\(\)](#) returns **buffer**. This function optionally optimizes the handling of FAST receive messages as described in "See Chapter 9 [Receiving Messages]."

**Return values:**

*something*

**Parameters:**

*message* (IN) Address of the message to be copied.

*buffer* (OUT) Address where the message is to be copied.

*len* (IN) The length in bytes of the message to be copied.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.40 GM\_ENTRY\_POINT unsigned int gm\_min\_message\_size (struct gm\_port \* port)**

Returns the minimum supported message size ([Details](#)).

**8.1.5.41 GM\_ENTRY\_POINT unsigned int gm\_min\_size\_for\_length (gm\_size\_t length)**

[gm\\_min\\_size\\_for\\_length\(\)](#) returns the minimum GM message buffer size required to store a message of length **length**.

**Return values:**

*gm\_log2\_roundup*

**Parameters:**

*length* (IN) The length of the message.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.42 GM\_ENTRY\_POINT unsigned int gm\_mtu (struct gm\_port \* port)**[gm\\_mtu\(\)](#) returns the value of GM\_MTU.**Return values:***GM\_MTU***Parameters:***port* (IN) The handle to the GM port.**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.43 GM\_ENTRY\_POINT char\* gm\_node\_id\_to\_host\_name (struct gm\_port \* port, unsigned int node\_id)**

This function is deprecated. Use [gm\\_node\\_id\\_to\\_host\\_name\\_ex\(\)](#) instead. Returns a pointer to the host name of the host containing the network interface card with GM node id *node\_id*. ([Details](#)).

**8.1.5.44 GM\_ENTRY\_POINT gm\_status\_t gm\_node\_id\_to\_unique\_id (struct gm\_port \* port, unsigned int n, char unique[6])**

Stores the MAC address for the interface ([gm\\_node\\_id\\_to\\_unique\\_id](#) "Details").

**8.1.5.45 GM\_ENTRY\_POINT unsigned int gm\_num\_ports (struct gm\_port \* p)**

Returns the number of ports supported by this build. ([Details](#)).

**8.1.5.46** GM\_ENTRY\_POINT unsigned int gm\_num\_send\_tokens (struct gm\_port \* p)

Returns the number of send tokens for this port. ([Details](#)).

**8.1.5.47** GM\_ENTRY\_POINT unsigned int gm\_num\_receive\_tokens (struct gm\_port \* p)

Returns the number of receive tokens for this port. ([Details](#)).

**8.1.5.48** GM\_ENTRY\_POINT unsigned int gm\_get\_port\_id (struct gm\_port \* p)

Returns the id of the GM port ([Details](#)).

**8.1.5.49** GM\_ENTRY\_POINT gm\_status\_t gm\_open (struct gm\_port \*\* p, unsigned int unit, unsigned int port, const char \* port\_name, enum gm\_api\_version version)

Opens a GM port on an interface ([Details](#)).

**8.1.5.50** GM\_ENTRY\_POINT void gm\_provide\_receive\_buffer\_with\_tag (struct gm\_port \* p, void \* ptr, unsigned int size, unsigned int priority, unsigned int tag)

Provides GM with a buffer into which it can receive messages ([Details](#)).

**8.1.5.51** GM\_ENTRY\_POINT int gm\_receive\_pending (struct gm\_port \* p)

Returns nonzero if a receive event is pending ([gm\\_receive\\_pending](#) "Details").

**8.1.5.52** GM\_ENTRY\_POINT int gm\_next\_event\_peek (struct gm\_port \* p, gm\_u16\_t \* sender)

Returns the nonzero event type if an event is pending ([Details](#)).

**8.1.5.53** GM\_ENTRY\_POINT union gm\_rcv\_event\* gm\_receive (struct gm\_port \* p)

Returns a receive event. ([Details](#)).

**8.1.5.54** GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_register\_memory (struct gm\_port \* p, void \* ptr, gm\_size\_t length)

Registers virtual memory for DMA transfers. ([gm\\_register\\_memory](#) "Details").

**8.1.5.55** GM\_ENTRY\_POINT int gm\_send\_token\_available (struct gm\_port \* p, unsigned int priority)

Tests for the availability of a send token without allocating the send token. ([Details](#)).

**8.1.5.56** GM\_ENTRY\_POINT void gm\_send\_with\_callback (struct gm\_port \* p, void \* message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \* context)

A fully asynchronous send. ([Details](#)).

**8.1.5.57** GM\_ENTRY\_POINT void gm\_send\_to\_peer\_with\_callback (struct gm\_port \* p, void \* message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \* context)

A fully asynchronous send from/to the same GM port on the sending and receiving side. ([Details](#)).

**8.1.5.58** GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_set\_acceptable\_sizes (struct gm\_port \* p, enum [gm\\_priority](#) priority, gm\_size\_t mask)

[gm\\_set\\_acceptable\\_sizes\(\)](#) informs GM of the acceptable sizes of GM messages received on port **p** with priority **priority**. Each set bit of **mask** indicates an acceptable size. While calling this function is not required, clients should call it during program initialization to detect errors involving the reception of badly sized messages to be reported nearly instantaneously, rather than after a substantial delay of 30 seconds or more.

Note: the MASK is a long to support larger than 2GByte packets (those with size larger than 31).

**Return values:**

**GM\_SUCCESS** Operation completed successfully.

**GM\_PERMISSION\_DENIED** Port number hasn't been set.

**GM\_INTERNAL\_ERROR** LANai is not running.

**GM\_INVALID\_PARAMETER** The priority has an invalid value.

**Parameters:**

*p* (IN) The GM port in use.  
*priority* (IN) The priority of the message.  
*mask*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.59** GM\_ENTRY\_POINT void gm\_set\_alarm (struct gm\_port \* *p*, struct gm\_alarm \* *my\_alarm*, gm\_u64\_t *usecs*, void(\* *callback*)(void \*), void \* *context*)

Sets an alarm, which may already be pending. ([gm\\_set\\_alarm](#) "Details").

**8.1.5.60** GM\_ENTRY\_POINT gm\_size\_t gm\_strlen (const char \* *cptr*)

[gm\\_strlen](#)() calculates the length of a string.

**Return values:**

*gm\_size\_t* The length of the string.

**Parameters:**

*cptr* (IN) The string.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.61** GM\_ENTRY\_POINT char\* gm\_strncpy (char \* *to*, const char \* *from*, int *len*)

[gm\\_strncpy](#)() copies exactly *n* bytes, truncating *src* or adding null characters to *dst* if necessary. The result will not be null-terminated if the length of *src* is *n* or more.

**Return values:**

*char* \* Returns a pointer to a destination string.

**Parameters:**

*to* (IN) The destination string to be copied.

*from* (IN) The source string to be copied.

*len* (IN) The number of bytes to be copied.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.62 GM\_ENTRY\_POINT int gm\_strcmp (const char \* a, const char \* b)**

[gm\\_strcmp\(\)](#) reimplements strcmp().

**Return values:**

*int* Returns an integer less than, equal to, or greater than zero if a is found, respectively, to be less than, to match, or be greater than b.

**Parameters:**

*a* The first string to be compared.

*b* The second string to be compared.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.63 GM\_ENTRY\_POINT int gm\_strncmp (const char \* a, const char \* b, int len)**

[gm\\_strncmp\(\)](#) reimplements strncmp().

**Return values:**

*int* Returns an integer less than, equal to, or greater than zero if the first len bytes of a is found, respectively, to be less than, to match, or be greater than b.



**Parameters:**

- a* (IN) The first string to be compared.
- b* (IN) The second string to be compared.
- len* (IN) The length in bytes.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.64 GM\_ENTRY\_POINT int gm\_strncasecmp (const char \* *a*, const char \* *b*, int *len*)**

[gm\\_strncasecmp\(\)](#) reimplements `strncasecmp()`.

**Return values:**

*int* Returns an integer less than, equal to, or greater than zero if the first *len* bytes of *a* is found, respectively, to be less than, to match, or be greater than *b*.

**Parameters:**

- a* (IN) The first string to be compared.
- b* (IN) The second string to be compared.
- len* (IN) The number of bytes.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.65 GM\_ENTRY\_POINT gm\_u64\_t gm\_ticks (struct gm\_port \* *port*)**

[gm\\_ticks\(\)](#) returns a 64-bit extended version of the LANai real time clock (RTC). For implementation reasons, the granularity of [gm\\_ticks\(\)](#) is 50 microseconds at the application level.

**Return values:**

*gm\_u64\_t*

**Parameters:**

*port* (IN) The handle to the GM port.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.66 GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_unique\_id (struct gm\_port \* port, char unique[6])**

Returns the board id number for an interface. ([gm\\_unique\\_id](#) "Details").

**8.1.5.67 GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_unique\_id\_to\_node\_id (struct gm\_port \* port, char unique[6], unsigned int \* node\_id)**

Returns the GM node id for a specific interface. ([Details](#)).

**8.1.5.68 GM\_ENTRY\_POINT void gm\_unknown (struct gm\_port \* p, union gm\_rcv\_event \* e)**

GM Event Handler. ([Details](#)).

**8.1.5.69 GM\_ENTRY\_POINT [gm\\_status\\_t](#) \_gm\_get\_route (struct gm\_port \* p, unsigned int node\_id, char \* route, unsigned int \* len)**

[\\_gm\\_get\\_route](#) function. ([Details](#)).

**8.1.5.70 GM\_ENTRY\_POINT void gm\_dump\_buffers (void)**

[gm\\_dump\\_buffers\(\)](#)

**Author:**

??

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.71 GM\_ENTRY\_POINT void gm\_register\_buffer (void \* *addr*, int *size*)**

[gm\\_register\\_buffer\(\)](#) registered a GM buffer.

**Return values:**

*1*

*0*

**Parameters:**

*addr* (**IN**) Address of data to be registered.

*size* (**IN**) Size of buffer to be registered.

**Author:**

??

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.72 GM\_ENTRY\_POINT int gm\_unregister\_buffer (void \* *addr*, int *size*)**

[gm\\_unregister\\_buffer\(\)](#) unregistered a GM buffer.

**Return values:**

*1*

*0*

**Parameters:**

*addr* (**IN**) Address of data to be freed.

*size* (**IN**) Size of buffer to be freed.

**Author:**

??

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.73 GM\_ENTRY\_POINT struct [gm\\_lookaside](#)\* gm\_create\_lookaside  
([gm\\_size\\_t entry\\_len](#), [gm\\_size\\_t min\\_entry\\_cnt](#))**

[gm\\_create\\_lookaside](#)() returns a newly created lookaside list to be used to allocate blocks of ENTRY\_LEN bytes. MIN\_ENTRY\_CNT entries are preallocated.

**Return values:**

[gm\\_lookaside](#) Handle to the lookaside table.

**Parameters:**

*entry\_len* (IN)

*min\_entry\_cnt* (IN)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.74 GM\_ENTRY\_POINT void gm\_destroy\_lookaside (struct [gm\\_lookaside](#)  
\* *l*)**

[gm\\_destroy\\_lookaside](#)() frees a lookaside list and all associated resources, including any buffers currently allocated from the lookaside list.

**Parameters:**

*l* (IN) Handle to the lookaside table.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.75 GM\_ENTRY\_POINT void\* gm\_lookaside\_alloc (struct [gm\\_lookaside](#) \*  
*l*)**

[gm\\_lookaside\\_alloc](#)() allocates an entry from the lookaside table, with debugging. It returns a buffer of size ENTRY\_LEN specified when the entry list L was created, or '0' if the buffer could not be allocated.

**Return values:**

*ptr* Buffer of size ENTRY\_LEN.

**Parameters:**

*l* (IN) Handle to the [gm\\_lookaside](#) list.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.76 GM\_ENTRY\_POINT void\* gm\_lookaside\_zalloc (struct [gm\\_lookaside](#) \*  
*l*)**

[gm\\_lookaside\\_zalloc\(\)](#) allocates and clear an entry from the lookaside table.

**Return values:**

*ptr*

**Parameters:**

*l* (IN) Handle to the [gm\\_lookaside](#) list.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.77 GM\_ENTRY\_POINT void gm\_lookaside\_free (void \* *ptr*)**

[gm\\_lookaside\\_free\(\)](#) schedules an allocated entry to be freed, and actually performs any scheduled free. It frees a block of memory previously allocated by a call to [gm\\_lookaside\\_alloc\(\)](#). The contents of the block of memory are guaranteed to be unchanged until the next operation is performed on the lookaside list.

**Parameters:**

*ptr* (IN) Pointer to the entry to be freed.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.78 GM\_ENTRY\_POINT** struct [gm\\_hash](#)\* **gm\_create\_hash** (long(\*  
*gm\_user\_compare*)(void \*key1, void \*key2), unsigned long(\*  
*gm\_user\_hash*)(void \*key1), gm\_size\_t *key\_len*, gm\_size\_t *data\_len*,  
gm\_size\_t *gm\_min\_entries*, int *flags*)

Creates a hash table. ([Details](#)).

**8.1.5.79 GM\_ENTRY\_POINT** void **gm\_destroy\_hash** (struct [gm\\_hash](#) \* *hash*)

[gm\\_destroy\\_hash](#)() frees all resources associated with the hash table, except for any client-allocated buffers.

**Parameters:**

*hash* (IN) Handle to the hash table.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.80 GM\_ENTRY\_POINT** void\* **gm\_hash\_remove** (struct [gm\\_hash](#) \* *hash*,  
void \* *key*)

Removes an entry from the hash table. ([Details](#)).

**8.1.5.81 GM\_ENTRY\_POINT** void\* **gm\_hash\_find** (struct [gm\\_hash](#) \* *hash*, void  
\* *key*)

Finds an entry in the hash table. ([Details](#)).

**8.1.5.82 GM\_ENTRY\_POINT** [gm\\_status\\_t](#) **gm\_hash\_insert** (struct [gm\\_hash](#) \*  
*hash*, void \* *key*, void \* *datum*)

Inserts an entry in the hash table. ([Details](#)).

**8.1.5.83** `GM_ENTRY_POINT` void `gm_hash_rekey` (struct [gm\\_hash](#) \* *hash*, void \* *old\_key*, void \* *new\_key*)

Replaces a key in the hash table. ([Details](#)).

**8.1.5.84** `GM_ENTRY_POINT` long `gm_hash_compare_strings` (void \* *key1*, void \* *key2*)

[gm\\_hash\\_compare\\_strings\(\)](#) is the function used to compare two strings.

**Return values:**

*long*

**Parameters:**

*key1* (IN) The key for the first string.

*key2* (IN) The key for the second string.

**See also:**

[gm\\_hash\\_compare\\_ints](#) [gm\\_hash\\_compare\\_longs](#) [gm\\_hash\\_compare\\_ptrs](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.85** `GM_ENTRY_POINT` unsigned long `gm_hash_hash_string` (void \* *key*)

[gm\\_hash\\_hash\\_string\(\)](#) is the function used to hash keys.

**Return values:**

*long*

**Parameters:**

*key* (IN) The key for the string in the hash table.

**See also:**

[gm\\_hash\\_compare\\_string](#) [gm\\_hash\\_hash\\_int](#) [gm\\_hash\\_hash\\_long](#) [gm\\_hash\\_hash\\_ptr](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.86 GM\_ENTRY\_POINT long gm\_hash\_compare\_longs (void \* *key1*, void \* *key2*)**

[gm\\_hash\\_compare\\_longs\(\)](#) is the function used to compare two longs.

**Return values:**

*long*

**Parameters:**

*key1* (IN)

*key2* (IN)

**See also:**

[gm\\_hash\\_compare\\_ints](#) [gm\\_hash\\_compare\\_strings](#) [gm\\_hash\\_compare\\_ptrs](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.87 GM\_ENTRY\_POINT unsigned long gm\_hash\_hash\_long (void \* *key*)**

[gm\\_hash\\_hash\\_long\(\)](#) is the function used to hash keys.

**Return values:**

*long*

**Parameters:**

*key* (IN)

**See also:**

[gm\\_hash\\_compare\\_long](#) [gm\\_hash\\_hash\\_int](#) [gm\\_hash\\_hash\\_string](#) [gm\\_hash\\_hash\\_ptr](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))



**8.1.5.88** `GM_ENTRY_POINT long gm_hash_compare_ints (void * key1, void * key2)`

[gm\\_hash\\_compare\\_ints\(\)](#) is the function used to compare two ints.

**Return values:**

*long*

**Parameters:**

*key1* (IN) The key for the first int.

*key2* (IN) The key for the second int.

**See also:**

[gm\\_hash\\_compare\\_longs](#) [gm\\_hash\\_compare\\_strings](#) [gm\\_hash\\_compare\\_ptrs](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.89** `GM_ENTRY_POINT unsigned long gm_hash_hash_int (void * key)`

[gm\\_hash\\_hash\\_int\(\)](#) is the function used to hash keys.

**Return values:**

*long*

**Parameters:**

*key* (IN)

**See also:**

[gm\\_hash\\_compare\\_int](#) [gm\\_hash\\_hash\\_ptr](#) [gm\\_hash\\_hash\\_long](#) [gm\\_hash\\_hash\\_string](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.90 GM\_ENTRY\_POINT long gm\_hash\_compare\_ptrs (void \* *key1*, void \* *key2*)**

[gm\\_hash\\_compare\\_ptrs\(\)](#) is the function used to compare two ptrs.

**Return values:**

*long*

**Parameters:**

*key1* (IN) The key for the first ptr.

*key2* (IN) The key for the second ptr.

**See also:**

[gm\\_hash\\_compare\\_longs](#) [gm\\_hash\\_compare\\_strings](#) [gm\\_hash\\_compare\\_ints](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.91 GM\_ENTRY\_POINT unsigned long gm\_hash\_hash\_ptr (void \* *key*)**

[gm\\_hash\\_hash\\_ptr\(\)](#) is the function used to hash keys.

**Return values:**

*long*

**Parameters:**

*key* (IN)

**See also:**

[gm\\_hash\\_compare\\_ptr](#) [gm\\_hash\\_hash\\_int](#) [gm\\_hash\\_hash\\_long](#) [gm\\_hash\\_hash\\_string](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.92 GM\_ENTRY\_POINT unsigned long gm\_crc (void \* *ptr*, gm\_size\_t *len*)**

[gm\\_crc\(\)](#) computes a CRC-32 of the indicated range of memory.

**Return values:**

*long*

**Parameters:**

*ptr* (IN) Pointer to a range of memory.

*len* (IN) The length of the indicated range of memory.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.93 GM\_ENTRY\_POINT unsigned long gm\_crc\_str (const char \* *ptr*)**

[gm\\_crc\\_str\(\)](#) computes a CRC-32 for the indicated string.

**Return values:**

*long*

**Parameters:**

*ptr* (IN) Pointer to a string.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.94 GM\_ENTRY\_POINT int gm\_rand (void)**

[gm\\_rand\(\)](#) returns a pseudo-random integer, using a poor but fast random number generator.

**Return values:**

*RANDOM\_NUMBER* The random number that was generated.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.95 GM\_ENTRY\_POINT void gm\_srand (int seed)**

[gm\\_srand\(\)](#) returns a pseudo-random integer, and requires a seed for the random number generator.

**Parameters:**

*seed* (IN) Seed for the random number generator.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.96 GM\_ENTRY\_POINT unsigned int gm\_rand\_mod (unsigned int a)**

[gm\\_rand\\_mod\(\)](#) returns a pseudo-random number modulo **modulus**, using a poor but fast random number generator.

**Return values:**

**RANDOM\_NUMBER**

**Parameters:**

*a* (IN) The modulus bound.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.97 GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_init (void)**

[gm\\_init\(\)](#) initializes GM. It increments the GM initialization counter and initializes GM if it was uninitialized. This call must be performed before any other GM call and before any reference to a GM global variable (e.g.: [GM\\_PAGE\\_LEN](#)). Each call to [gm\\_init\(\)](#) should be matched by a call to [gm\\_finalize\(\)](#).

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

*GM\_FAILURE* Error occurred.

**See also:**

[gm\\_finalize](#) [gm\\_open](#) [gm\\_close](#) [gm\\_exit](#) [gm\\_abort](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.98 GM\_ENTRY\_POINT void gm\_finalize (void)**

[gm\\_finalize\(\)](#) decrements the GM initialization counter and if it becomes zero, frees all resources associated with GM in the current process. Each call to [gm\\_finalize\(\)](#) should be matched by a call to [gm\\_init\(\)](#).

**See also:**

[gm\\_init](#) [gm\\_open](#) [gm\\_close](#) [gm\\_abort](#) [gm\\_exit](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.99 GM\_ENTRY\_POINT unsigned long gm\_log2\_roundup (unsigned long *n*)**

[gm\\_log2\\_roundup\(\)](#) returns the logarithm, base 2, of *n*, rounding up to the next integer.

**Return values:***LOG***Parameters:***n* (IN) The integer for which the logarithm will be computed.**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.100 GM\_ENTRY\_POINT struct gm\_mutex\* gm\_create\_mutex (void)**[gm\\_create\\_mutex\(\)](#) creates a GM mutex.**Return values:***gm\_mutex* (OUT) Handle to the GM mutex.**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.101 GM\_ENTRY\_POINT void gm\_destroy\_mutex (struct gm\_mutex \* mu)**[gm\\_destroy\\_mutex\(\)](#) destroys a GM mutex.**Parameters:***mu* (IN) The handle to the GM mutex.**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.102 GM\_ENTRY\_POINT void gm\_mutex\_enter (struct gm\_mutex \* *mu*)**

[gm\\_mutex\\_enter\(\)](#)

**Parameters:**

*mu* (IN) The handle to the GM mutex.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.103 GM\_ENTRY\_POINT void gm\_mutex\_exit (struct gm\_mutex \* *mu*)**

[gm\\_mutex\\_exit\(\)](#)

**Parameters:**

*mu* (IN) The handle to the GM mutex.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.104 GM\_ENTRY\_POINT struct [gm\\_zone](#)\* gm\_zone\_create\_zone (void \*  
*base*, gm\_size\_t *length*)**

[gm\\_zone\\_create\\_zone\(\)](#)

**Return values:**

[gm\\_zone](#) Handle to the GM zone.

**Parameters:**

*base*

*length*

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.105 GM\_ENTRY\_POINT void gm\_zone\_destroy\_zone (struct [gm\\_zone](#) \* zone)**[gm\\_zone\\_destroy\\_zone\(\)](#)**Parameters:***zone* (IN) Pointer to the GM zone.**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.106 GM\_ENTRY\_POINT void\* gm\_zone\_free (struct [gm\\_zone](#) \* zone, void \* a)**[gm\\_zone\\_free\(\)](#)**Parameters:***zone* (IN) Pointer to the GM zone.*a***Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.107 GM\_ENTRY\_POINT void\* gm\_zone\_malloc (struct [gm\\_zone](#) \* zone, gm\_size\_t length)**[gm\\_zone\\_malloc\(\)](#) mallocs a GM zone.



**Parameters:**

*zone* (IN) Pointer to the GM zone.

*length*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.108 GM\_ENTRY\_POINT void\* gm\_zone\_malloc (struct [gm\\_zone](#) \* zone, gm\_size\_t count, gm\_size\_t length)**

[gm\\_zone\\_malloc\(\)](#) callocs a GM zone.

**Parameters:**

*zone* (IN) Pointer to the GM zone.

*count*

*length*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.109 GM\_ENTRY\_POINT int gm\_zone\_addr\_in\_zone (struct [gm\\_zone](#) \* zone, void \* p)**

[gm\\_zone\\_addr\\_in\\_zone\(\)](#)

**Parameters:**

*zone* (IN) Pointer to the GM zone.

*p*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.110 GM\_ENTRY\_POINT void gm\_resume\_sending (struct gm\_port \* *p*, unsigned int *priority*, unsigned int *target\_node\_id*, unsigned int *target\_port\_id*, gm\_send\_completion\_callback\_t *callback*, void \* *context*)**

gm\_resume\_sending() reenables packet transmission of messages from **port** of priority **priority** destined for target\_port\_id of **target\_node\_id**. This function should only be called after an error is reported to a send completion callback routine. The message that generated the error is not resent. The first four parameters must match those of the failed send. It should be called only once per reported error. This function requires a send token, which will be returned to the client in the callback function.

gm\_resume\_sending() and gm\_drop\_sends(), as most gm requests, require a send token, and the callback you give to them is just meant to return this token. These gm\_requests always succeed (if called in a valid manner), so the callback will **always** be called with **GM\_SUCCESS** (which here does not mean at all that something was sent successfully, just that the request has been taken into account, and the token used for that request was recycled).

**Parameters:**

*p* (IN) The handle to the GM port.

*priority* (IN) The priority of the message being sent.

*target\_node\_id* (IN) The GM node to which the message is being sent.

*target\_port\_id* (IN) The GM port on the destination GM node to which the message is being sent.

*callback* (IN) The function called when the send is complete.

*context* (IN) Pointer to an integer or to a structure that is passed to the callback function.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in gm.h)

**8.1.5.111 GM\_ENTRY\_POINT void gm\_drop\_sends (struct gm\_port \* *port*, unsigned int *priority*, unsigned int *target\_node\_id*, unsigned int *target\_port\_id*, gm\_send\_completion\_callback\_t *callback*, void \* *context*)**

gm\_drop\_sends() drops all enqueued sends for **port** of priority **priority** destined for target\_port\_id of **target\_node\_id** to be dropped, and reenables packet transmission on that connection. This function should only be called after an error is reported to a send completion callback routine. The first four parameters must match those of the failed

send. It should be called only once per reported error. This function requires a send token, which will be returned to the client in the callback function. The dropped sends will then be returned to the client with a status of GM\_SEND\_DROPPED.

[gm\\_drop\\_sends\(\)](#) and [gm\\_resume\\_sending\(\)](#), as most gm requests, require a send token, and the callback you give to them is just meant to return this token. These gm\_requests always succeed (if called in a valid manner), so the callback will **always** be called with **GM\_SUCCESS** (which here does not mean at all that something was sent successfully, just that the request has been taken into account, and the token used for that request was recycled).

**Parameters:**

*port* (IN) The handle to the GM port.

*priority* (IN) The priority of the message.

*target\_node\_id* (IN) The GM node to which the message is being sent.

*target\_port\_id* (IN) The GM port on the destination GM node to which the message was sent.

*callback* (IN) The function called when the send is complete.

*context* (IN) Pointer to an integer or to a structure that is passed to the callback function.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.112 GM\_ENTRY\_POINT [gm\\_pid\\_t](#) gm\_getpid (void)**

[gm\\_getpid\(\)](#) is a cover for the usual Unix getpid() functionality.

**Return values:**

*gm\_pid\_t* The process ID of the parent of the current process.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.113 GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_directcopy\_get (struct gm\_port \* *p*, void \* *source\_addr*, void \* *target\_addr*, gm\_size\_t *length*, unsigned int *source\_instance\_id*, unsigned int *source\_port\_id*)**

[gm\\_directcopy\\_get\(\)](#) copies data of length **length** bytes, specified at the address **source\_addr** of the local process using the port **port\_id** of the board **source\_instance\_id** to the memory area specified at the address **target\_addr** of the current process. This implementation bypasses all of the protection of the operating system to provide a memory copy from one process's memory space to another one. The memory areas must have been registered by GM prior to calling this function in order to lock memory pages at their physical locations. There are no alignment or length constraints but the maximum performance will be reached with aligned addresses on both sides. This function is supported exclusively on Linux.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

*GM\_FAILURE* Error occurred.

**Parameters:**

*p* (IN) Handle to the GM port.

*source\_addr* (IN) Address of the data to be copied.

*target\_addr* (IN) Target address of the copied data.

*length* (IN) The length (in bytes) of the area to be copied.

*source\_instance\_id* (IN) The id of the interface.

*source\_port\_id* (IN) The port id of the interface.

**Author:**

Patrick Geoffray

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.114 GM\_ENTRY\_POINT void gm\_perror (const char \* *message*, [gm\\_status\\_t](#) *error*)**

[gm\\_perror\(\)](#) is similar to ANSI perror(), but takes the error code as a parameter to allow thread safety in future implementations, and only supports GM error numbers. Prints **message** followed by a description of **errno**.

**Parameters:**

*message* (OUT) Textual description of the GM error.

*error* (IN) GM Error code.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.115 GM\_ENTRY\_POINT int gm\_sleep (unsigned int *seconds*)**

[gm\\_sleep\(\)](#) emulates the ANSI standard [sleep\(\)](#), sleeping the entire process for **seconds** seconds.

**Return values:**

*SLEEP*

**Parameters:**

*seconds* (IN) The number of seconds for which the process should sleep.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.116 GM\_ENTRY\_POINT void gm\_exit ([gm\\_status\\_t](#) *status*)**

[gm\\_exit\(\)](#) causes the current process to exit with a status appropriate to the GM status code **status**.

**Parameters:**

*status* (IN) The GM status code, as specified in [gm.h](#).

**See also:**

[gm\\_init](#) [gm\\_open](#) [gm\\_close](#) [gm\\_finalize](#) [gm\\_abort](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.117 GM\_ENTRY\_POINT int gm\_printf (const char \* *format*, ...)**

[gm\\_printf\(\)](#) emulates or invokes the ANSI standard printf() function.

**Return values:**

*0* Operation completed successfully.

**Parameters:**

*format* Specifies how the arguments are converted for output.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.118 GM\_ENTRY\_POINT char\* gm\_strerror (gm\_status\_t *error*)**

[gm\\_strerror\(\)](#) is an error function for GM. The error is only valid until next call to this function.

**Return values:**

*char*

**Parameters:**

*error* (IN) GM status code.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.119 GM\_ENTRY\_POINT gm\_status\_t gm\_set\_enable\_nack\_down (struct gm\_port \* *port*, int *flag*)**

[gm\\_set\\_enable\\_nack\\_down\(\)](#)

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

***GM\_PERMISSION\_DENIED*****Parameters:**

*port* (IN) Handle to the GM port.

*flag*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.120 GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_max\_node\_id\_in\_use (struct [gm\\_port](#) \* *port*, unsigned int \* *n*)**

Returns the maximum GM node ID that is in use by the network attached to the port. ([Details](#)).

**8.1.5.121 GM\_ENTRY\_POINT int gm\_eprintf (const char \* *format*, ...)**

[gm\\_eprintf\(\)](#) emulates or invokes the ANSI standard `vprintf()` function.

**Return values:**

*0* Operation completed successfully.

**Parameters:**

*format* Specifies how the variable-length arguments are converted for output.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.122 GM\_ENTRY\_POINT void\* gm\_memset (void \* *s*, int *c*, [gm\\_size\\_t](#) *n*)**

[gm\\_memset\(\)](#) reimplements the UNIX function `memset()`.

**Return values:**

*void* Returns a pointer to the memory area *s*.

**Parameters:**

- s* (IN) The memory area.
- c* (IN) The constant byte size.
- n* (IN) The number of bytes.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.123 GM\_ENTRY\_POINT `char * gm_strdup (const char * in)`**

[gm\\_strdup\(\)](#) reimplements the UNIX function `strdup()`.

**Return values:**

*char \** Returns a pointer to the duplicated string, or NULL if insufficient memory was available.

**Parameters:**

- in* (IN) The string to be duplicated.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.124 GM\_ENTRY\_POINT `gm_status_t gm_mark (struct gm_mark_set * set, gm_mark_t * m)`**

[gm\\_mark\(\)](#) adds ‘\*MARK’ to SET. Requires O(constant) time if the mark set has pre-allocated resources for the mark. Otherwise, requires O(constant) average time.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

**Parameters:**

- set* (IN) Handle to the GM mark set.



*m*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.125 GM\_ENTRY\_POINT int gm\_mark\_is\_valid (struct [gm\\_mark\\_set](#) \* *set*, [gm\\_mark\\_t](#) \* *m*)**

[gm\\_mark\\_is\\_valid\(\)](#) returns nonzero value if ‘\*MARK’ is in SET. Requires O(constant) time.

**Return values:**

*int*

**Parameters:**

*set* (IN) Handle to the GM mark set.

*m*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.126 GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_create\_mark\_set (struct [gm\\_mark\\_set](#) \*\* *m*sp, unsigned long *cnt*)**

[gm\\_create\\_mark\\_set\(\)](#) returns a pointer to a new mark set at SET with enough pre-located resources to support INIT\_COUNT. Returns GM\_SUCCESS on success. Requires time comparable to [malloc\(\)](#).

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

**Parameters:**

*m*sp (IN) Handle to the GM mark set.

*cnt*

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.127 GM\_ENTRY\_POINT void gm\_destroy\_mark\_set (struct [gm\\_mark\\_set](#) \* *set*)**

[gm\\_destroy\\_mark\\_set\(\)](#) frees all resources associated with mark set '\*SET'. Requires time comparable to [free\(\)](#).

**Parameters:**

*set* (IN) Handle to the GM mark set.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.1.5.128 GM\_ENTRY\_POINT void gm\_unmark (struct [gm\\_mark\\_set](#) \* *set*, [gm\\_mark\\_t](#) \* *m*)**

[gm\\_unmark\(\)](#) removes '\*MARK' from SET. Requires O(constant) time.

**Parameters:**

*set* (IN) Handle to the GM mark set.

*m*

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.129 GM\_ENTRY\_POINT void gm\_unmark\_all (struct gm\_mark\_set \* set)**

[gm\\_unmark\\_all\(\)](#) unmarks all marks for the mark set, freeing all but the initial number of preallocated mark references.

Removes all marks from SET. Requires O(constant) time.

**Parameters:**

*set* (IN) Handle to the GM mark set.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.130 GM\_ENTRY\_POINT gm\_status\_t gm\_on\_exit (gm\_on\_exit\_callback\_t callback, void \* arg)**

[gm\\_on\\_exit\(\)](#) is like Linux `on_exit()`. This function registers a callback so that 'CALLBACK(STATUS,ARG)' is called when the program exits. Callbacks are called in the reverse of the order of registration. This function is also somewhat similar to BSD 'atexit()'.

Call the callbacks in the reverse order registered inside [gm\\_exit\(\)](#), passing GM exit status and registered argument to the callback.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

*GM\_OUT\_OF\_MEMORY*

**Parameters:**

*callback*

*arg*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.1.5.131** GM\_ENTRY\_POINT void gm\_put (struct gm\_port \* p, void \* local\_buffer, gm\_remote\_ptr\_t remote\_buffer, gm\_size\_t len, enum gm\_priority priority, unsigned int target\_node\_id, unsigned int target\_port\_id, gm\_send\_completion\_callback\_t callback, void \* context)

Directed send (PUT) ([Details](#)).

**8.1.5.132** GM\_ENTRY\_POINT gm\_status\_t gm\_global\_id\_to\_node\_id (struct gm\_port \* port, unsigned int global\_id, unsigned int \* node\_id)

Stores at \*node\_id the local connection ID corresponding to the connection to global\_id. ([gm\\_global\\_id\\_to\\_node\\_id](#) "Details").

**8.1.5.133** GM\_ENTRY\_POINT gm\_status\_t gm\_node\_id\_to\_global\_id (struct gm\_port \* port, unsigned int node\_id, unsigned int \* global\_id)

Stores at \*global\_id the global node ID corresponding to the connection identified by the local connection ID. ([Details](#)).

**8.1.5.134** GM\_ENTRY\_POINT gm\_status\_t gm\_node\_id\_to\_host\_name\_ex (struct gm\_port \* port, unsigned int timeout\_usecs, unsigned int node\_id, char \* name[GM\_MAX\_HOST\_NAME\_LEN+1])

Store at \*name the host name of the host containing the network interface card with GM node id node\_id. ([Details](#)).

**8.1.5.135** GM\_ENTRY\_POINT gm\_status\_t gm\_host\_name\_to\_node\_id\_ex (struct gm\_port \* port, unsigned int timeout\_usecs, const char \* host\_name, unsigned int \* node\_id)

Store at \*node\_id the node ID of the host containing the network interface card with GM host name host\_name. ([Details](#)).

## 8.1.6 Variable Documentation

**8.1.6.1** GM\_ENTRY\_POINT const unsigned char gm\_log2\_roundup\_table[257]

Log 2 roundup table. ([Details](#)).

## 8.2 gm\_abort.c File Reference

```
#include "gm_config.h"
#include "gm_debug.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_abort](#) ()

#### 8.2.1 Detailed Description

#### 8.2.2 Function Documentation

##### 8.2.2.1 GM\_ENTRY\_POINT void gm\_abort (void)

[gm\\_abort](#)() aborts the current process, and is a wrapper around the system function [abort](#)().

**See also:**

[gm\\_init](#) [gm\\_open](#) [gm\\_close](#) [gm\\_exit](#) [gm\\_finalize](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.3 gm\_alloc\_pages.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_internal.h"
#include "gm_malloc_debug.h"
```

### Functions

- GM\_ENTRY\_POINT void \* [gm\\_alloc\\_pages](#) (gm\_size\_t alloc\_len)
- GM\_ENTRY\_POINT void [gm\\_free\\_pages](#) (void \*addr, gm\_size\_t alloc\_len)

### 8.3.1 Detailed Description

### 8.3.2 Function Documentation

#### 8.3.2.1 GM\_ENTRY\_POINT void\* gm\_alloc\_pages (gm\_size\_t alloc\_len)

[gm\\_alloc\\_pages](#)() allocates a page-aligned buffer of length ALLOC\_LEN, where ALLOC\_LEN is a multiple of GM\_PAGE\_LEN. Any fractional page following the buffer is wasted.

#### Return values:

*ptr* Pointer to the allocated buffer.

0 Error occurred.

#### Parameters:

*alloc\_len* (IN) The length of buffer to be allocated.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

**8.3.2.2 GM\_ENTRY\_POINT void gm\_free\_pages (void \* *addr*, gm\_size\_t *alloc\_len*)**

[gm\\_free\\_pages\(\)](#) frees the pages at **addr**, which were previously allocated with [gm\\_alloc\\_pages\(\)](#).

**Parameters:**

*addr* **(IN)** The address of the buffer to be freed.

*alloc\_len* **(IN)** The length (in bytes) of the buffer to be freed.

**See also:**

[gm\\_alloc\\_pages](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.4 gm\_alloc\_send\_token.c File Reference

```
#include "gm_call_trace.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_alloc\\_send\\_token](#) (gm\_port\_t \*p, unsigned int priority)

#### 8.4.1 Detailed Description

#### 8.4.2 Function Documentation

##### 8.4.2.1 GM\_ENTRY\_POINT int gm\_alloc\_send\_token (gm\_port\_t \* p, unsigned int priority)

[gm\\_alloc\\_send\\_token](#)() allocates a send token of priority **priority** previously freed with [gm\\_free\\_send\\_token](#)() and returns 0 if no token is available. Clients may choose to maintain their own send token counts without using this utility function.

#### Return values:

*send\_token\_cnt*

0 Error occurred.

#### Parameters:

*p* (IN) The GM port on the source/sender GM node from which the communication will be sent.

*priority* (IN) The priority of the message to be sent.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))



## 8.5 gm\_allow\_remote\_memory\_access.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_allow\\_remote\\_memory\\_access](#) (struct gm\_port \*port)

#### 8.5.1 Detailed Description

#### 8.5.2 Function Documentation

##### 8.5.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_allow\\_remote\\_memory\\_access](#) (struct gm\_port \* port)

[gm\\_allow\\_remote\\_memory\\_access\(\)](#) allows any remote GM port to modify the contents of any GM DMAable memory using the [gm\\_directed\\_send\(\)](#) function. This is a significant security hole, but is very useful on tightly coupled clusters on trusted networks.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

**Parameters:**

*port* (IN) Handle to the GM port.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.6 gm\_bcopy.c File Reference

```
#include "gm_internal.h"
#include "gm_bcopy_up.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_bcopy](#) (const void \*from, void \*to, gm\_size\_t len)

#### 8.6.1 Detailed Description

#### 8.6.2 Function Documentation

##### 8.6.2.1 GM\_ENTRY\_POINT void gm\_bcopy (const void \*from, void \*to, gm\_size\_t len)

[gm\\_bcopy](#)() copies len bytes starting at **from** to location **to**. This function does not handle overlapping regions.

#### Parameters:

- from** (IN) The starting location for the region to be copied.
- to** (IN) The ending location for the region to be copied.
- len** (IN) The length in bytes of the region to be copied.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.7 gm\_blocking\_receive.c File Reference

```
#include "gm_call_trace.h"
#include "gm_cmp64.h"
#include "gm_compiler.h"
#include "gm_debug.h"
#include "gm_internal.h"
#include "gm_set_alarm.h"
```

### Functions

- GM\_ENTRY\_POINT gm\_rcv\_event\_t \* [gm\\_blocking\\_receive](#) (gm\_port\_t \*p)

#### 8.7.1 Detailed Description

#### 8.7.2 Function Documentation

##### 8.7.2.1 GM\_ENTRY\_POINT gm\_rcv\_event\_t\* gm\_blocking\_receive (gm\_port\_t \* p)

[gm\\_blocking\\_receive](#)() blocks until there is a receive event and then returns a pointer to the event. If no send is immediately available, this call suspends the current process until a receive event is available. As an optimization for applications with one CPU per CPU-intensive thread, this function polls for receives for one millisecond before sleeping the process, so it is not suited for machines running more than one performance critical process or thread on the machine.

#### Return values:

**GM\_SUCCESS** Operation completed successfully.

**GM\_NO\_RECV\_EVENT** Handle all flushed alarm events, which the user never needs to know about, and which must not cause my\_alarm to be cancelled or reset.

**GM\_FLUSHED\_ALARM\_EVENT** Intercept my\_alarm.

***GM\_ALARM\_EVENT*** If our alarm went off, block. Be careful to not handle any other alarms at this time since we are not passing the `GM_ALARM_EVENT` to the user and we want to maintain the semantics that user alarms are called only inside `gm_unknown()` when it is called by the user.

**Parameters:**

*p* (IN) The GM port in use.

**Author:**

Glenn Brown

**Version:**

`GM_API_VERSION` (as defined in `gm.h`)

## 8.8 gm\_blocking\_receive\_no\_spin.c File Reference

```
#include "gm_call_trace.h"
#include "gm_cmp64.h"
#include "gm_debug.h"
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### Functions

- `GM_ENTRY_POINT gm_rcv_event_t * gm_blocking_receive_no_spin (gm_port_t *p)`

### 8.8.1 Detailed Description

### 8.8.2 Function Documentation

#### 8.8.2.1 `GM_ENTRY_POINT gm_rcv_event_t* gm_blocking_receive_no_spin (gm_port_t *p)`

`gm_blocking_receive_no_spin()` behaves just like `gm_blocking_receive()`, only it sleeps the current thread immediately if no receive is pending. It is well suited to applications with more than one CPU-intensive thread per processor.

#### Return values:

*GM\_SUCCESS* Operation completed successfully.  
*GM\_NO\_RECV\_EVENT*  
*GM\_WAKE\_REQUEST\_EVENT*  
*GM\_SLEEP\_EVENT*

#### Parameters:

*p* (IN) The GM port in use.

#### Author:

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

## 8.9 gm\_bzero.c File Reference

```
#include "gm_config.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_bzero](#) (void \*ptr, gm\_size\_t len)

#### 8.9.1 Detailed Description

#### 8.9.2 Function Documentation

##### 8.9.2.1 GM\_ENTRY\_POINT void gm\_bzero (void \* ptr, gm\_size\_t len)

[gm\\_bzero](#)() clears the **len** bytes of memory starting at **ptr**. This function does not use partword I/O unless it must (for speed, especially when doing PIO), and does not rely on the system bzero() functionality, which may not be safe for PIO mapped memory.

#### Parameters:

- ptr** (IN) The pointer to the memory location.  
**len** (IN) The number of bytes of memory to be bzero'ed.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.10 gm\_malloc.c File Reference

```
#include <stdlib.h>
#include "gm_call_trace.h"
#include "gm_internal.h"
#include "gm_malloc_debug.h"
```

### Functions

- GM\_ENTRY\_POINT void \* [gm\\_malloc](#) (gm\_size\_t len, gm\_size\_t cnt)

#### 8.10.1 Detailed Description

#### 8.10.2 Function Documentation

##### 8.10.2.1 GM\_ENTRY\_POINT void\* gm\_malloc (gm\_size\_t len, gm\_size\_t cnt)

[gm\\_malloc](#)() allocates and clears an array of **cnt** elements of length **len**.

#### Parameters:

*len* (IN) The number of bytes in each element.

*cnt* (IN) The number of elements in the array.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))



## 8.11 gm\_close.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_debug_open.h"
#include "gm_dma_malloc.h"
#include "gm_enable_trace.h"
#include "gm_internal.h"
#include "gm_ptr_hash.h"
#include "gm_trace.h"
#include <stdio.h>
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_close](#) (gm\_port\_t \*p)

#### 8.11.1 Detailed Description

#### 8.11.2 Function Documentation

##### 8.11.2.1 GM\_ENTRY\_POINT void gm\_close (gm\_port\_t \* p)

[gm\\_close](#)() closes a previously opened port **p**, and frees all resources associated with that port.

#### Parameters:

**p** (IN) The GM port to be closed.

#### See also:

[gm\\_open](#) [gm\\_init](#) [gm\\_exit](#) [gm\\_finalize](#) [gm\\_abort](#)

#### Author:

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

## 8.12 gm\_crc.c File Reference

```
#include "gm_internal.h"
#include "gm_crc32.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned long [gm\\_crc](#) (void \*\_ptr, gm\_size\_t len)
- GM\_ENTRY\_POINT unsigned long [gm\\_crc\\_str](#) (const char \*ptr)

#### 8.12.1 Detailed Description

This file contains the GM API functions, [gm\\_crc\(\)](#) and [gm\\_crc\\_str\(\)](#), which compute 32-bit CRCs on the contents of memory. These functions are not guaranteed to perform any particular variant of the CRC-32, but these functions are useful for creating robust hashing functions.

#### 8.12.2 Function Documentation

##### 8.12.2.1 GM\_ENTRY\_POINT unsigned long gm\_crc (void \* *ptr*, gm\_size\_t *len*)

[gm\\_crc\(\)](#) computes a CRC-32 of the indicated range of memory.

**Return values:**

*long*

**Parameters:**

*ptr* (IN) Pointer to a range of memory.

*len* (IN) The length of the indicated range of memory.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.12.2.2 GM\_ENTRY\_POINT unsigned long gm\_crc\_str (const char \* *ptr*)**

[gm\\_crc\\_str\(\)](#) computes a CRC-32 for the indicated string.

**Return values:**

*long*

**Parameters:**

*ptr* (IN) Pointer to a string.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.13 gm\_datagram\_send.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_enable_fast_small_send.h"
#include "gm_enable_trace.h"
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_datagram\\_send](#) (gm\_port\_t \*p, void \*message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

#### 8.13.1 Detailed Description

This file contains the GM API function [gm\\_datagram\\_send\(\)](#).

#### 8.13.2 Function Documentation

**8.13.2.1** GM\_ENTRY\_POINT void [gm\\_datagram\\_send](#) (gm\_port\_t \* p, void \* *message*, unsigned int *size*, gm\_size\_t *len*, unsigned int *priority*, unsigned int *target\_node\_id*, unsigned int *target\_port\_id*, [gm\\_send\\_completion\\_callback\\_t](#) *callback*, void \* *context*)

[gm\\_datagram\\_send\(\)](#) queues **message** of length **length** to be sent unreliably to a buffer of size **size** at **target\_port\_id** on **target\_node\_id**. **length** must be no larger than **GM-MTU**. If any network error is encountered while sending the packet, the packet is silently and immediately dropped. After the packet has been DMA'ed from host memory, [callback\(port,context,status\)](#) is called inside a user invocation of [gm\\_unknown\(\)](#), reporting the **status** of the attempted send.

**Parameters:**

*p* (**IN**) The GM port on the source/sender GM node from which the communication is being sent.

*message* (**IN**) The pointer to the data to be communicated.

*size* (**IN**) The size of the buffer.

*len* (**IN**) The length in bytes of the array.

*priority* (**IN**) The priority of the message being communicated.

*target\_node\_id* (**IN**) The GM node to which the message is being sent.

*target\_port\_id* (**IN**) The GM port on the destination GM node to which the message is being sent.

*callback* (**IN**) The function called when the send is complete.

*context* (**IN**) Pointer to an integer or to a structure that is passed to the callback function.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.14 gm\_datagram\_send\_4.c File Reference

```
#include "gm_internal.h"
#include "gm_enable_pio_sends.h"
#include "gm_enable_datagrams.h"
#include "gm_enable_trace.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_datagram\\_send\\_4](#) (gm\_port\_t \*p, gm\_u32\_t data, unsigned int size, unsigned long len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

#### 8.14.1 Detailed Description

This file contains the GM API function [gm\\_datagram\\_send\\_4\(\)](#).

#### 8.14.2 Function Documentation

**8.14.2.1** GM\_ENTRY\_POINT void [gm\\_datagram\\_send\\_4](#) (gm\_port\_t \* p, gm\_u32\_t data, unsigned int size, unsigned long len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \* context)

[gm\\_datagram\\_send\\_4\(\)](#) queues **gm\_u32\_t** message of length **length** to be sent unreliably to a buffer of size **size** at **target\_port\_id** on **target\_node\_id**. **length** must be no larger than **GM\_MTU**. If any network error is encountered while sending the packet, the packet is silently and immediately dropped. After the packet has been DMA'ed from host memory, [callback\(port,context,status\)](#) is called inside a user invocation of [gm\\_unknown\(\)](#), reporting the **status** of the attempted send.

#### Parameters:

- p** (IN) The GM port on the source/sender GM node from which the communication is being sent.

***data*** (IN) The pointer to the `gm_u32_t` data to be communicated.

***size*** (IN) The size of the buffer.

***len*** (IN) The length in bytes of the array.

***priority*** (IN) The priority of the message being communicated.

***target\_node\_id*** (IN) The GM node to which the message is being sent.

***target\_port\_id*** (IN) The GM port on the destination GM node to which the message is being sent.

***callback*** (IN) The function called when the send is complete.

***context*** (IN) Pointer to an integer or to a structure that is passed to the callback function.

**See also:**

[gm\\_datagram\\_send](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))



## 8.15 gm\_debug\_buffers.c File Reference

```
#include "gm_debug.h"
#include "gm_internal.h"
#include <stdio.h>
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_unregister\\_buffer](#) (void \*addr, int size)
- GM\_ENTRY\_POINT void [gm\\_register\\_buffer](#) (void \*addr, int size)
- GM\_ENTRY\_POINT void [gm\\_dump\\_buffers](#) (void)

### 8.15.1 Detailed Description

This file contains the GM API functions [gm\\_register\\_buffer\(\)](#), [gm\\_unregister\\_buffer\(\)](#), [gm\\_dump\\_buffers\(\)](#).

### 8.15.2 Function Documentation

#### 8.15.2.1 GM\_ENTRY\_POINT int gm\_unregister\_buffer (void \* *addr*, int *size*)

[gm\\_unregister\\_buffer\(\)](#) unregistered a GM buffer.

#### Return values:

*1*  
*0*

#### Parameters:

*addr* (IN) Address of data to be freed.  
*size* (IN) Size of buffer to be freed.

#### Author:

??

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

**8.15.2.2 GM\_ENTRY\_POINT void gm\_register\_buffer (void \* *addr*, int *size*)**

[gm\\_register\\_buffer\(\)](#) registered a GM buffer.

**Return values:**

*1*

*0*

**Parameters:**

*addr* (IN) Address of data to be registered.

*size* (IN) Size of buffer to be registered.

**Author:**

??

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.15.2.3 GM\_ENTRY\_POINT void gm\_dump\_buffers (void)**

[gm\\_dump\\_buffers\(\)](#)

**Author:**

??

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.15.3 Variable Documentation****8.15.3.1 GM\_ENTRY\_POINT char\* gm\_buf\_status\_name[ ]****Initial value:**

```
{
 "gm_in_send",
 "gm_in_recv",
 "gm_in_app",
 "gm_invalid_status",
}
```

## 8.16 gm\_deregister.c File Reference

```
#include "gm_call_trace.h"
#include "gm_compiler.h"
#include "gm_debug.h"
#include "gm_debug_mem_register.h"
#include "gm_internal.h"
#include "gm_io.h"
#include "gm_lanai_command.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_deregister\\_memory](#) (gm\_port\_t \*p, void \*ptr, gm\_size\_t length)

### 8.16.1 Detailed Description

This file contains the GM API function [gm\\_deregister\\_memory](#)().

### 8.16.2 Function Documentation

#### 8.16.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_deregister\\_memory](#) (gm\_port\_t \*p, void \*ptr, gm\_size\_t length)

[gm\\_deregister\\_memory](#)() deregisters **len** bytes of user virtual that were previously registered for DMA transfers with a matching call to [gm\\_register\\_memory](#)() using **pvma**.

#### Return values:

**GM\_SUCCESS** Operation completed successfully.  
**GM\_FAILURE**

#### Parameters:

**p** (IN) The GM port in use.

*pvma* (IN) The pvma used to register the memory to now be deregistered.

*length* (IN) The number of bytes in the array to be deregistered.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.17 gm\_directcopy.c File Reference

```
#include "gm_debug.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_directcopy\\_get](#) (struct gm\_port \*p, void \*source\_addr, void \*target\_addr, gm\_size\_t length, unsigned int source\_instance\_id, unsigned int source\_port\_id)

### 8.17.1 Detailed Description

This file contains the GM API function [gm\\_directcopy\\_get](#)().

### 8.17.2 Function Documentation

#### 8.17.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_directcopy\\_get](#) (struct gm\_port \* p, void \* source\_addr, void \* target\_addr, gm\_size\_t length, unsigned int source\_instance\_id, unsigned int source\_port\_id)

[gm\\_directcopy\\_get](#)() copies data of length **length** bytes, specified at the address **source\_addr** of the local process using the port **port\_id** of the board **source\_instance\_id** to the memory area specified at the address **target\_addr** of the current process. This implementation bypasses all of the protection of the operating system to provide a memory copy from one process's memory space to another one. The memory areas must have been registered by GM prior to calling this function in order to lock memory pages at their physical locations. There are no alignment or length constraints but the maximum performance will be reached with aligned addresses on both sides. This function is supported exclusively on Linux.

#### Return values:

**GM\_SUCCESS** Operation completed successfully.

**GM\_FAILURE** Error occurred.

#### Parameters:

**p** (IN) Handle to the GM port.

*source\_addr* (IN) Address of the data to be copied.

*target\_addr* (IN) Target address of the copied data.

*length* (IN) The length (in bytes) of the area to be copied.

*source\_instance\_id* (IN) The id of the interface.

*source\_port\_id* (IN) The port id of the interface.

**Author:**

Patrick Geoffray

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.18 gm\_directed\_send.c File Reference

```
#include "gm.h"
#include "gm_debug.h"
#include "gm_enable_put.h"
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### 8.18.1 Detailed Description

This file contains the GM API function [gm\\_directed\\_send\(\)](#).



## 8.19 gm\_dma\_alloc.c File Reference

```
#include "gm_internal.h"
#include "gm_call_trace.h"
```

### Functions

- GM\_ENTRY\_POINT void \* [gm\\_dma\\_alloc](#) (struct gm\_port \*p, gm\_size\_t count, gm\_size\_t length)

#### 8.19.1 Detailed Description

This file contains the GM API function [gm\\_dma\\_alloc\(\)](#).

#### 8.19.2 Function Documentation

##### 8.19.2.1 GM\_ENTRY\_POINT void\* gm\_dma\_alloc (struct gm\_port \* p, gm\_size\_t count, gm\_size\_t length)

[gm\\_dma\\_alloc\(\)](#) allocates and clears **count** \* **length** bytes of DMAable memory aligned on a 4-byte boundary. Memory should be freed using [gm\\_dma\\_free\(\)](#).

##### Parameters:

- p** (IN) Handle to the GM port.
- count** (IN) The number of elements to be calloc'ed.
- length** (IN) The size of each element to be calloc'ed.

##### See also:

[gm\\_dma\\_malloc](#) [gm\\_dma\\_free](#)

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.20 gm\_dma\_malloc.c File Reference

```
#include "gm_call_trace.h"
#include "gm_compiler.h"
#include "gm_debug.h"
#include "gm_debug_mem_register.h"
#include "gm_internal.h"
#include "gm_malloc_debug.h"
#include "gm_register_recv_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_dma\\_free](#) (gm\_port\_t \*p, void \*addr)
- GM\_ENTRY\_POINT void \* [gm\\_dma\\_malloc](#) (struct gm\_port \*p, gm\_size\_t length)

### 8.20.1 Detailed Description

This file includes source for the user-level API calls [gm\\_dma\\_malloc\(\)](#) and [gm\\_dma\\_free\(\)](#).

### 8.20.2 Function Documentation

#### 8.20.2.1 GM\_ENTRY\_POINT void gm\_dma\_free (gm\_port\_t \* p, void \* addr)

[gm\\_dma\\_free\(\)](#) frees **p**, which was allocated by a call to [gm\\_dma\\_malloc\(\)](#) or [gm\\_dma\\_malloc\(\)](#). Note that the memory is not necessarily unlocked and returned to the operating system, but may be reused in future calls to [gm\\_dma\\_malloc\(\)](#) or [gm\\_dma\\_malloc\(\)](#).

#### Parameters:

**p** (IN) The GM port.

**addr** (IN) The address of the memory to be freed.

**See also:**

[gm\\_dma\\_calloc](#) [gm\\_dma\\_malloc](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.20.2.2 GM\_ENTRY\_POINT void\* gm\_dma\_malloc (struct gm\_port \* p, gm\_size\_t length)**

[gm\\_dma\\_malloc\(\)](#) allocates **length** bytes of DMAable memory aligned on a 4-byte boundary. Memory should be freed using [gm\\_dma\\_free\(\)](#).

**Parameters:**

*p* (IN) Handle to the GM port.

*length* (IN) The number of bytes to be malloc'ed.

**See also:**

[gm\\_dma\\_calloc](#) [gm\\_dma\\_free](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.21 gm\_drop\_sends.c File Reference

```
#include "gm_internal.h"
#include "_gm_modsend.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_drop\\_sends](#) (struct gm\_port \*port, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

#### 8.21.1 Detailed Description

This file contains the GM API function [gm\\_drop\\_sends\(\)](#).

#### 8.21.2 Function Documentation

**8.21.2.1 GM\_ENTRY\_POINT void gm\_drop\_sends (struct gm\_port \* port, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, gm\_send\_completion\_callback\_t callback, void \* context)**

[gm\\_drop\\_sends\(\)](#) drops all enqueued sends for **port** of priority **priority** destined for target\_port\_id of **target\_node\_id** to be dropped, and reenables packet transmission on that connection. This function should only be called after an error is reported to a send completion callback routine. The first four parameters must match those of the failed send. It should be called only once per reported error. This function requires a send token, which will be returned to the client in the callback function. The dropped sends will then be returned to the client with a status of GM\_SEND\_DROPPED.

[gm\\_drop\\_sends\(\)](#) and [gm\\_resume\\_sending\(\)](#), as most gm requests, require a send token, and the callback you give to them is just meant to return this token. These gm requests always succeed (if called in a valid manner), so the callback will **always** be called with **GM\_SUCCESS** (which here does not mean at all that something was sent successfully, just that the request has been taken into account, and the token used for that request was recycled).

**Parameters:**

*port* (IN) The handle to the GM port.

*priority* (IN) The priority of the message.

*target\_node\_id* (IN) The GM node to which the message is being sent.

*target\_port\_id* (IN) The GM port on the destination GM node to which the message was sent.

*callback* (IN) The function called when the send is complete.

*context* (IN) Pointer to an integer or to a structure that is passed to the callback function.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.22 gm\_eprintf.c File Reference

```
#include "gm_config.h"
#include "gm.h"
#include "gm_debug.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_eprintf](#) (const char \*format,...)

#### 8.22.1 Detailed Description

#### 8.22.2 Function Documentation

##### 8.22.2.1 GM\_ENTRY\_POINT int gm\_eprintf (const char \* *format*, ...)

[gm\\_eprintf](#)() emulates or invokes the ANSI standard `vprintf()` function.

**Return values:**

0 Operation completed successfully.

**Parameters:**

*format* Specifies how the variable-length arguments are converted for output.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.23 gm\_exit.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_internal.h"
#include "gm_malloc_debug.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_exit](#) ([gm\\_status\\_t](#) status)

#### 8.23.1 Detailed Description

#### 8.23.2 Function Documentation

##### 8.23.2.1 GM\_ENTRY\_POINT void [gm\\_exit](#) ([gm\\_status\\_t](#) status)

[gm\\_exit](#)() causes the current process to exit with a status appropriate to the GM status code [status](#).

#### Parameters:

*status* (IN) The GM status code, as specified in [gm.h](#).

#### See also:

[gm\\_init](#) [gm\\_open](#) [gm\\_close](#) [gm\\_finalize](#) [gm\\_abort](#)

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.24 gm\_flush\_alarm.c File Reference

```
#include "gm_call_trace.h"
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_flush\\_alarm](#) (gm\_port\_t \*p)

#### 8.24.1 Detailed Description

This file contains the GM API function [gm\\_flush\\_alarm](#)().

#### 8.24.2 Function Documentation

##### 8.24.2.1 GM\_ENTRY\_POINT void gm\_flush\_alarm (gm\_port\_t \*p)

[gm\\_flush\\_alarm](#)() flushes the alarm queue.

##### Parameters:

*p* (IN) The GM port.

##### See also:

[gm\\_set\\_alarm](#) [gm\\_initialize\\_alarm](#) [gm\\_cancel\\_alarm](#)

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))



## 8.25 gm\_free.c File Reference

```
#include <stdlib.h>
#include "gm_call_trace.h"
#include "gm_debug_malloc.h"
#include "gm_internal.h"
#include "gm_malloc_debug.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_free](#) (void \*ptr)

#### 8.25.1 Detailed Description

This file contains the GM API function [gm\\_free\(\)](#).

#### 8.25.2 Function Documentation

##### 8.25.2.1 GM\_ENTRY\_POINT void gm\_free (void \* ptr)

[gm\\_free\(\)](#) frees the memory buffer at **ptr**, which was previously allocated by [gm\\_malloc\(\)](#), or [gm\\_calloc\(\)](#).

##### Parameters:

*ptr* (IN) Address of the memory to be freed.

##### See also:

[gm\\_malloc](#) [gm\\_calloc](#)

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.26 gm\_free\_send\_token.c File Reference

```
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_free\\_send\\_token](#) (gm\_port\_t \*p, unsigned priority)

### 8.26.1 Detailed Description

This file contains the GM API function [gm\\_free\\_send\\_token](#)().

### 8.26.2 Function Documentation

#### 8.26.2.1 GM\_ENTRY\_POINT void gm\_free\_send\_token (gm\_port\_t \* p, unsigned priority)

[gm\\_free\\_send\\_token](#)() increments a free send token of **priority** for **port** so that it can later be allocated using [gm\\_alloc\\_send\\_token](#)(). Clients may choose to maintain their own count of send tokens in the client's possession instead of using this utility function.

#### Parameters:

*p*  
*priority*

#### See also:

[gm\\_alloc\\_send\\_token](#) [gm\\_free\\_send\\_tokens](#)

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.27 gm\_free\_send\_tokens.c File Reference

```
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_free\\_send\\_tokens](#) (gm\_port\_t \*p, unsigned priority, unsigned count)

#### 8.27.1 Detailed Description

This file contains the GM API function [gm\\_free\\_send\\_tokens](#)().

#### 8.27.2 Function Documentation

##### 8.27.2.1 GM\_ENTRY\_POINT void gm\_free\_send\_tokens (gm\_port\_t \* p, unsigned priority, unsigned count)

[gm\\_free\\_send\\_tokens](#)() increments a **count** of free send tokens of **priority** for **port** so that it can later be allocated using [gm\\_alloc\\_send\\_token](#)(). Clients may choose to maintain their own count of send tokens in the client's possession instead of using this utility function.

##### Parameters:

*p*  
*priority*  
*count*

##### See also:

[gm\\_alloc\\_send\\_token](#) [gm\\_free\\_send\\_token](#)

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.28 gm\_get.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_enable_get.h"
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_get](#) (gm\_port\_t \*p, [gm\\_remote\\_ptr\\_t](#) remote\_buffer, void \*local\_buffer, unsigned long len, enum [gm\\_priority](#) priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

### 8.28.1 Detailed Description

This file contains the GM API function [gm\\_get\(\)](#).

### 8.28.2 Function Documentation

**8.28.2.1** GM\_ENTRY\_POINT void [gm\\_get](#) (gm\_port\_t \* p, [gm\\_remote\\_ptr\\_t](#) remote\_buffer, void \* local\_buffer, unsigned long len, enum [gm\\_priority](#) priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \* context)

[gm\\_get\(\)](#) performs an RDMA Read operation.

[gm\\_get\(\)](#) transfers the **len** bytes at **remote\_buffer** to **target\_port\_id** on **target\_node\_id** with priority **priority** and stores the data at the local virtual memory address **local\_buffer**. Call [callback\(port,context,status\)](#) when the receive completes or fails, with **status** indicating the status of the receive.

#### Parameters:

- p* (IN) The GM port on the destination GM node to which the communication is being received.

*remote\_buffer* (IN) Address of the remote buffer.

*local\_buffer* (OUT) Address of the local buffer.

*len* (IN) The length in bytes of the buffer to be received.

*priority* (IN) The priority of the data being received.

*target\_node\_id* (IN) The GM node to which the data is being received.

*target\_port\_id* (IN) The GM port on the destination GM node to which the message is being received.

*callback* (IN) The function called when the receive is complete.

*context* (IN) Pointer to an integer or to a structure that is passed to the callback function.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.29 gm\_get\_host\_name.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_host\\_name](#) (gm\_port\_t \*port, char name[GM\_MAX\_HOST\_NAME\_LEN])

### 8.29.1 Detailed Description

This file contains the GM API function [gm\\_get\\_host\\_name\(\)](#).

### 8.29.2 Function Documentation

**8.29.2.1** GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_host\\_name](#) (gm\_port\_t \*port, char name[GM\_MAX\_HOST\_NAME\_LEN])

[gm\\_get\\_host\\_name\(\)](#) copies the host name of the local node to **name**.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

**Parameters:**

*port* (IN) The GM port.

*name* (OUT) The host name of the GM node.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.30 gm\_get\_mapper\_unique\_id.c File Reference

```
#include "gm_internal.h"
```

### Functions

- [gm\\_status\\_t gm\\_get\\_mapper\\_unique\\_id](#) (struct gm\_port \*port, char \*id)

#### 8.30.1 Detailed Description

This file contains the GM API function [gm\\_get\\_mapper\\_unique\\_id\(\)](#).

#### 8.30.2 Function Documentation

##### 8.30.2.1 [gm\\_status\\_t gm\\_get\\_mapper\\_unique\\_id](#) (struct gm\_port \* port, char \* id)

[gm\\_get\\_mapper\\_unique\\_id\(\)](#) copies the 6-byte ethernet address of the network interface card associated with **port** to the buffer at **id**.

#### Return values:

*GM\_SUCCESS* Operation completed successfully.

#### Parameters:

*port* (**IN**) The GM port associated with a specific network interface card.

*id* (**OUT**) Buffer to which the 6-byte ethernet address (GM\_GET\_MAPPER\_UNIQUE\_ID) is copied.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.31 gm\_get\_node\_id.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_node\\_id](#) (gm\_port\_t \*port, unsigned int \*node\_id)

#### 8.31.1 Detailed Description

This file contains the GM API function [gm\\_get\\_node\\_id.c](#)

#### 8.31.2 Function Documentation

##### 8.31.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_get\\_node\\_id](#) (gm\_port\_t \*port, unsigned int \*node\_id)

[gm\\_get\\_node\\_id](#)() copies the GM ID of the network interface card associated with **port** to the address **node\_id**.

##### Return values:

*GM\_SUCCESS* Operation completed successfully

*GM\_NODE\_ID\_NOT\_YET\_SET* The GM node ID has not been set.

##### Parameters:

*port* (IN) The GM port.

*node\_id* (OUT) Address to copy the GM ID.

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))



## 8.32 gm\_get\_node\_type.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_get\\_node\\_type](#) ([gm\\_port\\_t](#) \*port, int \*node\_type)

#### 8.32.1 Detailed Description

This file contains the GM API function [gm\\_get\\_node\\_type\(\)](#).

#### 8.32.2 Function Documentation

##### 8.32.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_get\\_node\\_type](#) ([gm\\_port\\_t](#) \* port, int \* node\_type)

[gm\\_get\\_node\\_type\(\)](#) does something.

##### Return values:

*GM\_SUCCESS* Operation completed successfully.

##### Parameters:

*port* (IN) The GM port.

*node\_type* (OUT) GM\_GET\_NODE\_TYPE, the Big Endian, Little Endian-ness???

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.33 gm\_get\_port\_id.c File Reference

```
#include "gm_compiler.h"
#include "gm_debug.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned int [gm\\_get\\_port\\_id](#) (gm\_port\_t \*p)

#### 8.33.1 Detailed Description

#### 8.33.2 Function Documentation

##### 8.33.2.1 GM\_ENTRY\_POINT unsigned int gm\_get\_port\_id (gm\_port\_t \*p)

[gm\\_get\\_port\\_id](#)() returns the id of the GM port **p**.

#### Return values:

*port\_id* The id of the GM port.

#### Parameters:

*p* (IN) The GM port.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.34 gm\_get\_unique\_board\_id.c File Reference

```
#include "gm_internal.h"
```

### Functions

- [gm\\_status\\_t gm\\_get\\_unique\\_board\\_id](#) (gm\_port\_t \*port, char unique[6])

#### 8.34.1 Detailed Description

This file contains the GM API function [gm\\_get\\_unique\\_board\\_id\(\)](#).

#### 8.34.2 Function Documentation

##### 8.34.2.1 [gm\\_status\\_t gm\\_get\\_unique\\_board\\_id](#) (gm\_port\_t \* *port*, char *unique*[6])

[gm\\_get\\_unique\\_board\\_id\(\)](#) copies the 6-byte MAC address of the network interface card associated with **port** to the buffer at **unique**.

#### Return values:

*GM\_SUCCESS* Operation completed successfully.

#### Parameters:

*port* (**IN**) The GM port associated with a specific network interface card.

*unique* (**OUT**) Buffer to which the 6-byte MAC address (GM\_GET\_UNIQUE\_BOARD\_ID) is copied.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.35 gm\_getpid.c File Reference

```
#include "gm_config.h"
#include "gm.h"
```

### Functions

- [gm\\_pid\\_t gm\\_getpid](#) (void)

### 8.35.1 Detailed Description

This file contains the GM API function [gm\\_getpid\(\)](#).

### 8.35.2 Function Documentation

#### 8.35.2.1 [gm\\_pid\\_t gm\\_getpid](#) (void)

[gm\\_getpid\(\)](#) is a cover for the usual Unix `getpid()` functionality.

**Return values:**

*gm\_pid\_t* The process ID of the parent of the current process.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.36 gm\_handle\_sent\_tokens.c File Reference

```
#include "gm_debug.h"
#include "gm_internal.h"
```

### 8.36.1 Detailed Description

This file contains the GM API function [gm\\_handle\\_sent\\_tokens\(\)](#) and is deprecated and included only for backward compatibility. Use [gm\\_unknown\(\)](#) instead.

## 8.37 gm\_hash.c File Reference

```
#include "gm_call_trace.h"
#include "gm_compiler.h"
#include "gm_crc32.h"
#include "gm_debug.h"
#include "gm_internal.h"
```

### Data Structures

- struct [gm\\_hash\\_entry](#)
- struct [gm\\_hash\\_segment](#)
- struct [gm\\_hash](#)

### Typedefs

- typedef [gm\\_hash\\_entry](#) [gm\\_hash\\_entry\\_t](#)
- typedef [gm\\_hash\\_segment](#) [gm\\_hash\\_segment\\_t](#)
- typedef [gm\\_hash](#) [gm\\_hash\\_t](#)

### Functions

- GM\_ENTRY\_POINT struct [gm\\_hash](#) \* [gm\\_create\\_hash](#) (long(\*gm\_user\_-compare)(void \*key1, void \*key2), unsigned long(\*gm\_user\_hash)(void \*key), gm\_size\_t key\_len, gm\_size\_t data\_len, gm\_size\_t min\_cnt, int flags)
- GM\_ENTRY\_POINT void [gm\\_destroy\\_hash](#) (struct [gm\\_hash](#) \*hash)
- GM\_ENTRY\_POINT void \* [gm\\_hash\\_remove](#) ([gm\\_hash\\_t](#) \*hash, void \*key)
- GM\_ENTRY\_POINT void \* [gm\\_hash\\_find](#) ([gm\\_hash\\_t](#) \*hash, void \*key)
- GM\_ENTRY\_POINT void [gm\\_hash\\_rekey](#) ([gm\\_hash\\_t](#) \*hash, void \*old\_key, void \*new\_key)
- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_hash\\_insert](#) ([gm\\_hash\\_t](#) \*hash, void \*key, void \*data)
- GM\_ENTRY\_POINT long [gm\\_hash\\_compare\\_strings](#) (void \*key1, void \*key2)
- GM\_ENTRY\_POINT unsigned long [gm\\_hash\\_hash\\_string](#) (void \*key)
- GM\_ENTRY\_POINT long [gm\\_hash\\_compare\\_longs](#) (void \*key1, void \*key2)
- GM\_ENTRY\_POINT unsigned long [gm\\_hash\\_hash\\_long](#) (void \*key)
- GM\_ENTRY\_POINT long [gm\\_hash\\_compare\\_ints](#) (void \*key1, void \*key2)

- GM\_ENTRY\_POINT unsigned long [gm\\_hash\\_hash\\_int](#) (void \*key)
- GM\_ENTRY\_POINT long [gm\\_hash\\_compare\\_ptrs](#) (void \*key1, void \*key2)
- GM\_ENTRY\_POINT unsigned long [gm\\_hash\\_hash\\_ptr](#) (void \*key)

### 8.37.1 Detailed Description

This file contains the GM API functions [gm\\_create\\_hash\(\)](#), [gm\\_destroy\\_hash\(\)](#), [gm\\_hash\\_remove\(\)](#), [gm\\_hash\\_find\(\)](#), [gm\\_hash\\_insert\(\)](#), [gm\\_hash\\_rekey\(\)](#), [gm\\_hash\\_compare\\_strings\(\)](#), [gm\\_hash\\_hash\\_string\(\)](#), [gm\\_hash\\_compare\\_longs\(\)](#), [gm\\_hash\\_hash\\_long\(\)](#), [gm\\_hash\\_compare\\_ints\(\)](#), [gm\\_hash\\_hash\\_int\(\)](#), [gm\\_hash\\_compare\\_ptr\(\)](#), [gm\\_hash\\_hash\\_ptr\(\)](#).

This module implements a generic hash table. It uses lookaside lists to ensure efficient memory allocation in the kernel.

GM implements a generic hash table with a flexible interface. This module can automatically manage storage of fixed-size keys and/or data, or can allow the client to manage storage for keys and/or data. It allows the client to specify arbitrary hashing and comparison functions.

For example,

```
hash = gm_create_hash (gm_hash_compare_strings, gm_hash_hash_string,
 0, 0, 0, 0);
```

creates a hash table that uses null-terminated character string keys residing in client-managed storage, and returns pointers to data in client-managed storage. In this case, all pointers to hash keys and data passed by GM to the client will be the same as the pointers passed by the client to GM.

As another example,

```
hash = gm_create_hash (gm_hash_compare_ints, gm_hash_hash_int,
 sizeof (int), sizeof (struct my_big_struct),
 100, 0);
```

creates a hash table that uses 'ints' as keys and returns pointers to copies of the inserted structures. All storage for the keys and data is automatically managed by the hash table. In this case, all pointers to hash keys and data passed by GM to the client will point to GM-managed buffers. This function also preallocates enough storage for 100 hash entries, guaranteeing that at least 100 key/data pairs can be inserted in the table if the hash table creation succeeds.

The automatic storage management option of GM not only is convenient, but also is extremely space efficient for keys and data no larger than a pointer, because when

keys and data are no larger than a pointer, GM automatically stores them in the space reserved for the pointer to the key or data, rather than allocating a separate buffer.

Note that all keys and data buffers are referred to by pointers, not by value. This allows keys and data buffers of arbitrary size to be used. As a special (but common) case, however, one may wish to use pointers as keys directly, rather than use what they point to. In this special case, use the following initialization, and pass the keys (pointers) directly to the API, rather than the usual references to the keys.

```
hash = gm_create_hash (gm_hash_compare_ptrs, gm_hash_hash_ptr,
 0, DATA_LEN, MIN_CNT, FLAGS);
```

While it is possible to specify a `KEY_LEN` of `'sizeof (void *)'` during initialization and treat pointer keys just like any other keys, the API above is more efficient, more convenient, and completely architecture independent.

## 8.37.2 Typedef Documentation

### 8.37.2.1 typedef struct [gm\\_hash\\_entry](#) gm\_hash\_entry\_t

A hash table entry.

### 8.37.2.2 typedef struct [gm\\_hash\\_segment](#) gm\_hash\_segment\_t

Structure representing a segment of allocated hash table bins. To double the size of the hash table, we allocate a new segment with just enough bins to double the number of bins in the hash table and prepend it to the list of hash segments. This way, we don't have to double-buffer the hash table while growing it, and we can grow the table closer to the limits of available memory. While we sometimes have to walk the  $O(\log(N))$ -segment list to find a bin, the average lookup only looks at  $O(2)$  segments, so operations are still constant-average-time as expected for hash tables.

### 8.37.2.3 typedef struct [gm\\_hash](#) gm\_hash\_t

The state of a GM hash table, referenced by the client only using opaque pointers.

## 8.37.3 Function Documentation



**8.37.3.1 GM\_ENTRY\_POINT struct [gm\\_hash](#)\* gm\_create\_hash (long(\*  
*gm\_user\_compare*)(void \*key1, void \*key2), unsigned long(\*  
*gm\_user\_hash*)(void \*key), *gm\_size\_t* key\_len, *gm\_size\_t* data\_len,  
*gm\_size\_t* min\_cnt, int flags)**

[gm\\_create\\_hash](#)() returns a newly-created '[gm\\_hash](#)' structure or '0' if the hash table could not be created.

**Return values:**

[gm\\_hash](#) Handle to the hash table.

**Parameters:**

*gm\_user\_compare* (IN) The function used to compare keys and may be any of [gm\\_hash\\_compare\\_ints](#), [gm\\_hash\\_compare\\_longs](#), [gm\\_hash\\_compare\\_ptrs](#), [gm\\_hash\\_compare\\_strings](#), or may be a client-defined function.

*gm\_user\_hash* (IN) The function to be used to hash keys and may be any of [gm\\_hash\\_hash\\_int](#), [gm\\_hash\\_hash\\_long](#), [gm\\_hash\\_hash\\_ptr](#), [gm\\_hash\\_hash\\_string](#), or may be a client-defined function.

*key\_len* (IN) Specifies the length of the keys to be used for the hash table, or '0' if the keys should not be copied into GM-managed buffers.

*data\_len* (IN) Specifies the length of the data to be stored in the hash table, or '0' if the data should not be copied into GM-managed buffers.

*min\_cnt* (IN) Specifies the number of entries for which storage should be preallocated.

*flags* (IN) Should be '0' because no flags are currently defined.

**See also:**

[gm\\_destroy\\_hash](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.2 GM\_ENTRY\_POINT void gm\_destroy\_hash (struct [gm\\_hash](#) \* hash)**

[gm\\_destroy\\_hash](#)() frees all resources associated with the hash table, except for any client-allocated buffers.

**Parameters:**

*hash* (IN) Handle to the hash table.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.37.3.3 GM\_ENTRY\_POINT void\* gm\_hash\_remove ([gm\\_hash\\_t](#) \* *hash*, void \* *key*)**

[gm\\_hash\\_remove\(\)](#) removes an entry associated with KEY from the hash table HASH and returns a pointer to the data associated with the key, or '0' if no match exists. If the data resides in a GM-managed buffer, it is only guaranteed to be valid until the next operation on the hash table.

**Parameters:***hash* (IN) Pointer to the hash table.*key* (IN) The key for the hash table.**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.37.3.4 GM\_ENTRY\_POINT void\* gm\_hash\_find ([gm\\_hash\\_t](#) \* *hash*, void \* *key*)**

[gm\\_hash\\_find\(\)](#) finds an entry associated with KEY from the hash table HASH and returns a pointer to the data associated with the key, or '0' if no match exists.

**Parameters:***hash* (IN) Pointer to the hash table.*key* (IN) The key for the hash table.**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.5 GM\_ENTRY\_POINT void gm\_hash\_rekey (gm\_hash\_t \* hash, void \* old\_key, void \* new\_key)**

[gm\\_hash\\_rekey\(\)](#) finds each entry with key OLD\_KEY and changes the key used to store the data to NEW\_KEY. This call is guaranteed to succeed.

**Parameters:**

*hash* (IN) Pointer to the hash table.

*old\_key* (IN) The previous key for the hash table.

*new\_key* (IN) The new key for the hash table.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.6 GM\_ENTRY\_POINT gm\_status\_t gm\_hash\_insert (gm\_hash\_t \* hash, void \* key, void \* data)**

[gm\\_hash\\_insert\(\)](#) stores the association of KEY and DATA in the hash table HASH. The key '\*'KEY (or data '\*'DATA) is copied into the hash table unless the table was initialized with a KEY\_LEN (or DATA\_LEN) of 0.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

**Parameters:**

*hash* (IN) Pointer to the hash table.

*key* (IN) The key for the hash table.

*data* (IN) Data to be inserted.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.7 GM\_ENTRY\_POINT long gm\_hash\_compare\_strings (void \* *key1*, void \* *key2*)**

[gm\\_hash\\_compare\\_strings\(\)](#) is the function used to compare two strings.

**Return values:**

*long*

**Parameters:**

*key1* (IN) The key for the first string.

*key2* (IN) The key for the second string.

**See also:**

[gm\\_hash\\_compare\\_ints](#) [gm\\_hash\\_compare\\_longs](#) [gm\\_hash\\_compare\\_ptrs](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.8 GM\_ENTRY\_POINT unsigned long gm\_hash\_hash\_string (void \* *key*)**

[gm\\_hash\\_hash\\_string\(\)](#) is the function used to hash keys.

**Return values:**

*long*

**Parameters:**

*key* (IN) The key for the string in the hash table.

**See also:**

[gm\\_hash\\_compare\\_string](#) [gm\\_hash\\_hash\\_int](#) [gm\\_hash\\_hash\\_long](#) [gm\\_hash\\_hash\\_ptr](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.9 GM\_ENTRY\_POINT long gm\_hash\_compare\_longs (void \* *key1*, void \* *key2*)**

[gm\\_hash\\_compare\\_longs\(\)](#) is the function used to compare two longs.

**Return values:**

*long*

**Parameters:**

*key1* (IN)

*key2* (IN)

**See also:**

[gm\\_hash\\_compare\\_ints](#) [gm\\_hash\\_compare\\_strings](#) [gm\\_hash\\_compare\\_ptrs](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.10 GM\_ENTRY\_POINT unsigned long gm\_hash\_hash\_long (void \* *key*)**

[gm\\_hash\\_hash\\_long\(\)](#) is the function used to hash keys.

**Return values:**

*long*

**Parameters:**

*key* (IN)

**See also:**

[gm\\_hash\\_compare\\_long](#) [gm\\_hash\\_hash\\_int](#) [gm\\_hash\\_hash\\_string](#) [gm\\_hash\\_hash\\_ptr](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.11 GM\_ENTRY\_POINT long gm\_hash\_compare\_ints (void \* *key1*, void \* *key2*)**

[gm\\_hash\\_compare\\_ints\(\)](#) is the function used to compare two ints.

**Return values:**

*long*

**Parameters:**

*key1* (IN) The key for the first int.

*key2* (IN) The key for the second int.

**See also:**

[gm\\_hash\\_compare\\_longs](#) [gm\\_hash\\_compare\\_strings](#) [gm\\_hash\\_compare\\_ptrs](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.12 GM\_ENTRY\_POINT unsigned long gm\_hash\_hash\_int (void \* *key*)**

[gm\\_hash\\_hash\\_int\(\)](#) is the function used to hash keys.

**Return values:**

*long*

**Parameters:**

*key* (IN)

**See also:**

[gm\\_hash\\_compare\\_int](#) [gm\\_hash\\_hash\\_ptr](#) [gm\\_hash\\_hash\\_long](#) [gm\\_hash\\_hash\\_string](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.13 GM\_ENTRY\_POINT long gm\_hash\_compare\_ptrs (void \* *key1*, void \* *key2*)**

[gm\\_hash\\_compare\\_ptrs\(\)](#) is the function used to compare two ptrs.

**Return values:**

*long*

**Parameters:**

*key1* (IN) The key for the first ptr.

*key2* (IN) The key for the second ptr.

**See also:**

[gm\\_hash\\_compare\\_longs](#) [gm\\_hash\\_compare\\_strings](#) [gm\\_hash\\_compare\\_ints](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.37.3.14 GM\_ENTRY\_POINT unsigned long gm\_hash\_hash\_ptr (void \* *key*)**

[gm\\_hash\\_hash\\_ptr\(\)](#) is the function used to hash keys.

**Return values:**

*long*

**Parameters:**

*key* (IN)

**See also:**

[gm\\_hash\\_compare\\_ptr](#) [gm\\_hash\\_hash\\_int](#) [gm\\_hash\\_hash\\_long](#) [gm\\_hash\\_hash\\_string](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.38 gm\_hex\_dump.c File Reference

```
#include "gm_config.h"
#include "gm_compiler.h"
#include "gm.h"
#include "gm_debug.h"
#include "gm_internal_funcs.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_hex\\_dump](#) (const void \*\_ptr, gm\_size\_t len)

#### 8.38.1 Detailed Description

This file contains the GM API function [gm\\_hex\\_dump\(\)](#).

#### 8.38.2 Function Documentation

##### 8.38.2.1 GM\_ENTRY\_POINT void gm\_hex\_dump (const void \* *ptr*, gm\_size\_t *len*)

[gm\\_hex\\_dump\(\)](#) prints the hex equivalent of data at *ptr*.

##### Parameters:

- ptr* (IN) Address of the data.
- len* (IN) The length (in bytes) of the data.

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))



## 8.39 gm\_host\_name\_to\_node\_id.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_debug_yp.h"
#include "gm_internal.h"
#include "gm_lanai_command.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_host\\_name\\_to\\_node\\_id\\_ex](#) (gm\_port\_t \*port, unsigned int timeout\_usecs, const char \*host\_name, unsigned int \*node\_id)
- GM\_ENTRY\_POINT unsigned int [gm\\_host\\_name\\_to\\_node\\_id](#) (gm\_port\_t \*port, char \*host\_name)

### 8.39.1 Detailed Description

The file containing the GM API function [gm\\_host\\_name\\_to\\_node\\_id\(\)](#).

### 8.39.2 Function Documentation

**8.39.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_host\\_name\\_to\\_node\\_id\\_ex](#)**  
(gm\_port\_t \*port, unsigned int timeout\_usecs, const char \*host\_name, unsigned int \*node\_id)

[gm\\_host\\_name\\_to\\_node\\_id\\_ex\(\)](#) translates a GM host name to a node ID.

#### Return values:

*GM\_SUCCESS* Operation completed successfully.

#### Parameters:

*port* (IN) The GM port associated with the specific GM ID.

*timeout\_usecs* (IN) The maximum length of time (in microseconds) to query, or 0 to use the default (long) timeout.

*host\_name* (IN)

*node\_id* (OUT) Where to store the node ID.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.39.2.2 GM\_ENTRY\_POINT unsigned int gm\_host\_name\_to\_node\_id  
(gm\_port\_t \* port, char \* host\_name)**

This function is deprecated. Use [gm\\_host\\_name\\_to\\_node\\_id\\_ex\(\)](#) instead.

[gm\\_host\\_name\\_to\\_node\\_id\(\)](#) returns the GM ID associated **host\_name** or GM\_NO\_SUCH\_NODE\_ID in case of an error.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

*GM\_NO\_SUCH\_NODE\_ID* Error.

**Parameters:**

*port* (IN) The GM port associated with the specific GM ID.

*host\_name* (OUT) Where to store the host name. At most GM\_MAX\_HOST\_NAME\_LEN+1 bytes will be written.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.40 gm\_init.c File Reference

```
#include "gm_call_trace.h"
#include "gm_crc32.h"
#include "gm_debug.h"
#include "gm_enable_security.h"
#include "gm_internal.h"
```

### Functions

- [gm\\_status\\_t gm\\_init\(\)](#)
- [void gm\\_finalize\(\)](#)

#### 8.40.1 Detailed Description

This file contains the GM API functions [gm\\_init\(\)](#) and [gm\\_finalize\(\)](#).

#### 8.40.2 Function Documentation

##### 8.40.2.1 [gm\\_status\\_t gm\\_init \(void\)](#)

[gm\\_init\(\)](#) initializes GM. It increments the GM initialization counter and initializes GM if it was uninitialized. This call must be performed before any other GM call and before any reference to a GM global variable (e.g.: `GM_PAGE_LEN`). Each call to [gm\\_init\(\)](#) should be matched by a call to [gm\\_finalize\(\)](#).

##### Return values:

- GM\_SUCCESS* Operation completed successfully.
- GM\_FAILURE* Error occurred.

##### See also:

[gm\\_finalize](#) [gm\\_open](#) [gm\\_close](#) [gm\\_exit](#) [gm\\_abort](#)

##### Author:

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.40.2.2 void gm\_finalize (void)**

[gm\\_finalize\(\)](#) decrements the GM initialization counter and if it becomes zero, frees all resources associated with GM in the current process. Each call to [gm\\_finalize\(\)](#) should be matched by a call to [gm\\_init\(\)](#).

**See also:**[gm\\_init](#) [gm\\_open](#) [gm\\_close](#) [gm\\_abort](#) [gm\\_exit](#)**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

## 8.41 gm\_isprint.c File Reference

```
#include "gm_internal.h"
```

### Functions

- int [gm\\_isprint](#) (int *c*)

#### 8.41.1 Detailed Description

The file containing the GM API function [gm\\_isprint](#)().

#### 8.41.2 Function Documentation

##### 8.41.2.1 int gm\_isprint (int *c*)

[gm\\_isprint](#)() is just like ANSI `isprint`(), only it works in the kernel and MCP.

##### Return values:

*int* The return value is nonzero if the character is a printable character including a space, and a zero value if not.

##### Parameters:

*c* (OUT) Printable character.

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.42 gm\_log2.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned long [gm\\_log2\\_roundup](#) (unsigned long n)

### Variables

- GM\_ENTRY\_POINT const unsigned char [gm\\_log2\\_roundup\\_table](#) [257]

### 8.42.1 Detailed Description

This file contains the GM API functions related to base 2 logarithmic computations. [gm\\_log2\\_roundup\\_table](#), [gm\\_log2\\_roundup\(\)](#).

### 8.42.2 Function Documentation

#### 8.42.2.1 GM\_ENTRY\_POINT unsigned long gm\_log2\_roundup (unsigned long *n*)

[gm\\_log2\\_roundup\(\)](#) returns the logarithm, base 2, of **n**, rounding up to the next integer.

**Return values:**

*LOG*

**Parameters:**

*n* (IN) The integer for which the logarithm will be computed.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

### 8.42.3 Variable Documentation

#### 8.42.3.1 GM\_ENTRY\_POINT const unsigned char gm\_log2\_roundup\_table[257]

Initial value:

```
{
0,
0, 1, 2, 2, 3, 3, 3, 3, 4, 4, 4, 4, 4, 4, 4, 4, 4,
5, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5, 5,
6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6,
6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6,
7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7,
7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7,
7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7,
7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8, 8,
}
```

Log 2 roundup table. ([Details](#)).

## 8.43 gm\_lookaside.c File Reference

```
#include "gm.h"
#include "gm_cache_line.h"
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_internal.h"
#include "gm_malloc_debug.h"
#include "gm_struct_lock.h"
```

### Data Structures

- struct [gm\\_lookaside\\_segment](#)
- struct [gm\\_lookaside](#)
- struct [gm\\_lookaside\\_segment\\_list](#)

### Functions

- GM\_ENTRY\_POINT void \* [gm\\_lookaside\\_alloc](#) (struct [gm\\_lookaside](#) \*l)
- GM\_ENTRY\_POINT void \* [gm\\_lookaside\\_zalloc](#) (struct [gm\\_lookaside](#) \*l)
- GM\_ENTRY\_POINT void [gm\\_lookaside\\_free](#) (void \*ptr)
- GM\_ENTRY\_POINT struct [gm\\_lookaside](#) \* [gm\\_create\\_lookaside](#) (gm\_size\_t entry\_len, gm\_size\_t min\_entry\_cnt)
- GM\_ENTRY\_POINT void [gm\\_destroy\\_lookaside](#) (struct [gm\\_lookaside](#) \*l)

#### 8.43.1 Detailed Description

This file contains the GM API functions [gm\\_create\\_lookaside\(\)](#), [gm\\_destroy\\_lookaside\(\)](#), [gm\\_lookaside\\_alloc\(\)](#), [gm\\_lookaside\\_zalloc\(\)](#), and [gm\\_lookaside\\_free\(\)](#).

This file implements a lookaside list. It is mainly intended to allow efficient memory allocation of small structures in kernels where the minimum memory allocation is a page, but is not kernel-specific.

GM implements a lookaside list, which may be used to manage small fixed-length blocks more efficiently than [gm\\_malloc\(\)](#) and [gm\\_free\(\)](#). Lookaside lists can also be



used to ensure that at least a minimum number of blocks are available for allocation at all times.

[gm\\_lookaside\\_alloc\(\)](#) returns cache-line-aligned buffers, in an attempt to minimize the cost of accessing data stored in the buffers.

## 8.43.2 Function Documentation

### 8.43.2.1 GM\_ENTRY\_POINT void\* gm\_lookaside\_alloc (struct gm\_lookaside \* l)

[gm\\_lookaside\\_alloc\(\)](#) allocates an entry from the lookaside table, with debugging. It returns a buffer of size `ENTRY_LEN` specified when the entry list `L` was created, or `'0'` if the buffer could not be allocated.

**Return values:**

*ptr* Buffer of size `ENTRY_LEN`.

**Parameters:**

*l* (IN) Handle to the [gm\\_lookaside](#) list.

**Author:**

Glenn Brown

**Version:**

`GM_API_VERSION` (as defined in [gm.h](#))

### 8.43.2.2 GM\_ENTRY\_POINT void\* gm\_lookaside\_zalloc (struct gm\_lookaside \* l)

[gm\\_lookaside\\_zalloc\(\)](#) allocates and clear an entry from the lookaside table.

**Return values:**

*ptr*

**Parameters:**

*l* (IN) Handle to the [gm\\_lookaside](#) list.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.43.2.3 GM\_ENTRY\_POINT void gm\_lookaside\_free (void \* ptr)**

[gm\\_lookaside\\_free\(\)](#) schedules an allocated entry to be freed, and actually performs any scheduled free. It frees a block of memory previously allocated by a call to [gm\\_lookaside\\_alloc\(\)](#). The contents of the block of memory are guaranteed to be unchanged until the next operation is performed on the lookaside list.

**Parameters:**

*ptr* (IN) Pointer to the entry to be freed.

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.43.2.4 GM\_ENTRY\_POINT struct [gm\\_lookaside\\*](#) gm\_create\_lookaside (gm\_size\_t entry\_len, gm\_size\_t min\_entry\_cnt)**

[gm\\_create\\_lookaside\(\)](#) returns a newly created lookaside list to be used to allocate blocks of ENTRY\_LEN bytes. MIN\_ENTRY\_CNT entries are preallocated.

**Return values:**

[gm\\_lookaside](#) Handle to the lookaside table.

**Parameters:**

*entry\_len* (IN)

*min\_entry\_cnt* (IN)

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

**8.43.2.5 GM\_ENTRY\_POINT** void gm\_destroy\_lookaside (struct [gm\\_lookaside](#) \* *l*)

[gm\\_destroy\\_lookaside\(\)](#) frees a lookaside list and all associated resources, including any buffers currently allocated from the lookaside list.

**Parameters:**

*l* (**IN**) Handle to the lookaside table.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.44 gm\_malloc.c File Reference

```
#include <stdlib.h>
#include "gm_call_trace.h"
#include "gm_compiler.h"
#include "gm_internal.h"
#include "gm_malloc_debug.h"
#include "gm_debug_malloc.h"
```

### Functions

- GM\_ENTRY\_POINT void \* [gm\\_malloc](#) (gm\_size\_t len)

#### 8.44.1 Detailed Description

This file contains the GM API function [gm\\_malloc\(\)](#).

#### 8.44.2 Function Documentation

##### 8.44.2.1 GM\_ENTRY\_POINT void\* [gm\\_malloc](#) (gm\_size\_t len)

[gm\\_malloc\(\)](#) returns a pointer to the specified amount of memory. In the kernel, the memory will be nonpageable.

#### Return values:

*ptr* Pointer to the specified amount of memory.

#### Parameters:

*len* (IN) The length in bytes to be malloc'ed.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.45 gm\_mark.c File Reference

```
#include "gm.h"
```

### Data Structures

- union [gm\\_mark\\_reference](#)
- struct [gm\\_mark\\_set](#)

### Typedefs

- typedef [gm\\_mark\\_reference](#) [gm\\_mark\\_reference\\_t](#)

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_mark](#) (struct [gm\\_mark\\_set](#) \*set, [gm\\_mark\\_t](#) \*m)
- GM\_ENTRY\_POINT int [gm\\_mark\\_is\\_valid](#) (struct [gm\\_mark\\_set](#) \*set, [gm\\_mark\\_t](#) \*m)
- GM\_ENTRY\_POINT void [gm\\_unmark](#) (struct [gm\\_mark\\_set](#) \*set, [gm\\_mark\\_t](#) \*m)
- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_create\\_mark\\_set](#) (struct [gm\\_mark\\_set](#) \*\*msp, unsigned long cnt)
- GM\_ENTRY\_POINT void [gm\\_destroy\\_mark\\_set](#) (struct [gm\\_mark\\_set](#) \*set)
- GM\_ENTRY\_POINT void [gm\\_unmark\\_all](#) (struct [gm\\_mark\\_set](#) \*set)

#### 8.45.1 Detailed Description

This file implements a constant-time system for marking memory locations and verifying that they have not been corrupted. All operations have average execution time of  $O(1)$ . If enough marks were specified in the call to [gm\\_create\\_mark\\_set\(\)](#), then all operations have worst-case performance of  $O(1)$ ; otherwise, the worst-case performance is  $O(N)$  where  $N$  is the number of marks in the mark set.

We do not simply use a magic number as a mark, or a mark set identifier as a mark, since this technique can lead to false positives where uninitialized marks appear valid. With this system, any localized corruption of a mark or the mark set database will result in the mark being determined to be invalid. False positives in this system require that both the mark and mark set database be corrupted in a consistent way, which is basically impossible.

This implementation was inspired by a challenge from John Regehr, who was forwarding a challenge from a UVA CS professor, paraphrased as follows: Can you implement the functionality of an array where initialization, insertion, removal, and test for presence of an array element each has execution time bounded by a constant, assuming you are initially provided with a pointer to a sufficient amount of uninitialized memory. You are only allowed  $O(N)$  storage, where  $N$  is the max number of entries." The answer is "Yes" and a lot of the same tricks are used here.

The GM "mark" API is introduced in GM-1.4. It allows the creation and destruction of mark sets, which allow mark addition, mark removal, and test for mark in mark set operations to be performed in constant time. Marks may be members of only one mark set at a time. Marks have the very unusual property that they need not be initialized before use.

All operations on marks are extremely efficient. Mark initialization requires zero time. Removing a mark from a mark set and testing for mark inclusion in a mark set take constant time. Addition of a mark to a mark set takes  $O(\text{constant})$  time, assuming the marks set was created with support for a sufficient number of marks; otherwise, it requires  $O(\text{constant})$  average time. Finally, creation and destruction of a mark set take time comparable to the time required for a single call to `malloc()` and `free()`, respectively.

Because marks need not be initialized before use, they can actually be used to determine if other objects have been initialized. This is done by putting a mark in the object, and adding the mark to a "mark set of marks in initialized objects" once the object has been initialized. This is similar to one common use of "magic numbers" for debugging purposes, except that it is immune to the possibility that the uninitialized magic number contained the magic number before initialization, so such marks can be used for non-debugging purposes. Therefore, marks can be used in ways that magic numbers cannot.

Marks have a nice set of properties that each mark in a mark set has a unique value and if this value is corrupted, then the mark is implicitly removed from the mark set. This makes marks useful for detecting memory corruption, and are less prone to false negatives than are magic numbers, which proliferate copies of a single value.

Finally, marks are location-dependent. This means that if a mark is copied, the copy will not be a member of the mark set.

## 8.45.2 Typedef Documentation

### 8.45.2.1 typedef union [gm\\_mark\\_reference](#) gm\_mark\_reference.t

a reference to a mark, or a free entry. The reference is used to validate the mark. That mark is considered valid if and only if `(set->reference[mark->tag] == mark)`.

### 8.45.3 Function Documentation

#### 8.45.3.1 GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_mark (struct [gm\\_mark\\_set](#) \* *set*, [gm\\_mark\\_t](#) \* *m*)

[gm\\_mark\(\)](#) adds '\*MARK' to SET. Requires O(constant) time if the mark set has pre-allocated resources for the mark. Otherwise, requires O(constant) average time.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

**Parameters:**

*set* (IN) Handle to the GM mark set.

*m*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

#### 8.45.3.2 GM\_ENTRY\_POINT [int](#) gm\_mark\_is\_valid (struct [gm\\_mark\\_set](#) \* *set*, [gm\\_mark\\_t](#) \* *m*)

[gm\\_mark\\_is\\_valid\(\)](#) returns nonzero value if '\*MARK' is in SET. Requires O(constant) time.

**Return values:**

*int*

**Parameters:**

*set* (IN) Handle to the GM mark set.

*m*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

### 8.45.3.3 GM\_ENTRY\_POINT void gm\_unmark (struct gm\_mark\_set \* set, gm\_mark\_t \* m)

gm\_unmark() removes '\*MARK' from SET. Requires O(constant) time.

#### Parameters:

*set* (IN) Handle to the GM mark set.

*m*

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in gm.h)

### 8.45.3.4 GM\_ENTRY\_POINT gm\_status\_t gm\_create\_mark\_set (struct gm\_mark\_set \*\* msp, unsigned long cnt)

gm\_create\_mark\_set() returns a pointer to a new mark set at SET with enough pre-located resources to support INIT\_COUNT. Returns GM\_SUCCESS on success. Requires time comparable to malloc().

#### Return values:

*GM\_SUCCESS* Operation completed successfully.

#### Parameters:

*msp* (IN) Handle to the GM mark set.

*cnt*

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in gm.h)

### 8.45.3.5 GM\_ENTRY\_POINT void gm\_destroy\_mark\_set (struct gm\_mark\_set \* set)

gm\_destroy\_mark\_set() frees all resources associated with mark set '\*SET'. Requires time comparable to free().



**Parameters:**

*set* (IN) Handle to the GM mark set.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.45.3.6 GM\_ENTRY\_POINT void gm\_unmark\_all (struct [gm\\_mark\\_set](#) \* *set*)**

[gm\\_unmark\\_all\(\)](#) unmarks all marks for the mark set, freeing all but the initial number of preallocated mark references.

Removes all marks from SET. Requires O(constant) time.

**Parameters:**

*set* (IN) Handle to the GM mark set.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.46 gm\_max\_length\_for\_size.c File Reference

```
#include "gm_internal.h"
#include "gm_types.h"
```

### Functions

- GM\_ENTRY\_POINT gm\_size\_t [gm\\_max\\_length\\_for\\_size](#) (unsigned int size)

#### 8.46.1 Detailed Description

This file contains the GM API function [gm\\_max\\_length\\_for\\_size\(\)](#).

#### 8.46.2 Function Documentation

##### 8.46.2.1 GM\_ENTRY\_POINT gm\_size\_t gm\_max\_length\_for\_size (unsigned int size)

[gm\\_max\\_length\\_for\\_size\(\)](#) returns the maximum length of a message that will fit in a GM buffer of size *size*.

#### Return values:

*gm\_size\_t*

#### Parameters:

*size* (IN) The size of the GM buffer.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.47 gm\_max\_node\_id.c File Reference

```
#include "gm_call_trace.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_max\\_node\\_id](#) (gm\_port\_t \*port, unsigned int \*n)

#### 8.47.1 Detailed Description

This file contains the GM API function [gm\\_max\\_node\\_id\(\)](#).

#### 8.47.2 Function Documentation

##### 8.47.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_max\\_node\\_id](#) (gm\_port\_t \* port, unsigned int \* n)

[gm\\_max\\_node\\_id\(\)](#) stores the maximum GM node ID supported by the network interface card corresponding to **port** at the address of **n**.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

**Parameters:**

*port* (IN)

*n* (OUT) Buffer to store the maximum GM node ID.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.48 gm\_max\_node\_id\_in\_use.c File Reference

```
#include "gm_call_trace.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_max\\_node\\_id\\_in\\_use](#) (gm\_port\_t \*port, unsigned int \*n)

#### 8.48.1 Detailed Description

This file contains the GM API function [gm\\_max\\_node\\_id\\_in\\_use\(\)](#).

This function exists to allow gm\_board\_info to have a guess for an upper limit to the actual number of nodes in the route table. Without it, gm\_board\_info calls thousands of gm\_get\_route() calls. When debugging, this is really a mess.

#### 8.48.2 Function Documentation

##### 8.48.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_max\\_node\\_id\\_in\\_use](#) (gm\_port\_t \*port, unsigned int \*n)

[gm\\_max\\_node\\_id\\_in\\_use\(\)](#) returns the maximum GM node ID that is in use by the network attached to the port.

##### Return values:

*GM\_SUCCESS* Operation completed successfully.

*GM\_UNATTACHED*

##### Parameters:

*port* (IN) The GM port.

*n* (OUT) Buffer containing the maximum GM node ID.

##### Author:

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.49 gm\_memcmp.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_memcmp](#) (const void \*\_a, const void \*\_b, gm\_size\_t len)

#### 8.49.1 Detailed Description

This file contains the GM API function [gm\\_memcmp](#)().

#### 8.49.2 Function Documentation

##### 8.49.2.1 GM\_ENTRY\_POINT int gm\_memcmp (const void \* *a*, const void \* *b*, gm\_size\_t *len*)

[gm\\_memcmp](#)() emulates the ANSI memcmp() function.

#### Return values:

*int* Returns an integer less than, equal to, or greater than zero if the first len bytes of *a* is found, respectively, to be less than, to match, or be greater than the first len bytes of *b*.

#### Parameters:

*a* (IN) The first memory area for comparison.  
*b* (IN) The second memory area for comparison.  
*len* (IN) The number of bytes to compare.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.50 gm\_memorize\_message.c File Reference

```
#include "gm_debug.h"
#include "gm_internal.h"
```

### Functions

- void \* [gm\\_memorize\\_message](#) (void \*message, void \*buffer, unsigned len)

#### 8.50.1 Detailed Description

This file contains the GM API function [gm\\_memorize\\_message\(\)](#).

#### 8.50.2 Function Documentation

##### 8.50.2.1 void\* gm\_memorize\_message (void \* message, void \* buffer, unsigned len)

[gm\\_memorize\\_message\(\)](#) is a wrapper around function `memcpy()`. [gm\\_memorize\\_message\(\)](#) copies a message into a buffer if needed. If **message** and **buffer** differ, `gm_memorize_message(port,message,buffer)` copies the message pointed to by **message** into the buffer pointed to by **buffer**. [gm\\_memorize\\_message\(\)](#) returns **buffer**. This function optionally optimizes the handling of FAST receive messages as described in "See Chapter 9 [Receiving Messages]."

##### Return values:

*something*

##### Parameters:

*message* (IN) Address of the message to be copied.

*buffer* (OUT) Address where the message is to be copied.

*len* (IN) The length in bytes of the message to be copied.

##### Author:

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))



## 8.51 gm\_memset.c File Reference

```
#include "gm_config.h"
#include "gm.h"
```

### Functions

- GM\_ENTRY\_POINT void \* [gm\\_memset](#) (void \*s, int c, gm\_size\_t n)

#### 8.51.1 Detailed Description

#### 8.51.2 Function Documentation

##### 8.51.2.1 GM\_ENTRY\_POINT void\* gm\_memset (void \* s, int c, gm\_size\_t n)

[gm\\_memset\(\)](#) reimplements the UNIX function `memset()`.

#### Return values:

*void* Returns a pointer to the memory area *s*.

#### Parameters:

- s* (IN) The memory area.
- c* (IN) The constant byte size.
- n* (IN) The number of bytes.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.52 gm\_min\_message\_size.c File Reference

```
#include "gm_compiler.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned int [gm\\_min\\_message\\_size](#) (gm\_port\_t \*p)

### 8.52.1 Detailed Description

This file contains the GM API function [gm\\_min\\_message\\_size](#)().

### 8.52.2 Function Documentation

#### 8.52.2.1 GM\_ENTRY\_POINT unsigned int [gm\\_min\\_message\\_size](#) (gm\_port\_t \*p)

[gm\\_min\\_message\\_size](#)() returns the minimum supported message size. This value is accessed through the function API to avoid hard-coding it in user applications, allowing dynamic library upgrades.

**Return values:**

*GM\_MIN\_MESSAGE\_SIZE*

**Parameters:**

*p*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.53 gm\_min\_size\_for\_length.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned int [gm\\_min\\_size\\_for\\_length](#) (gm\_size\_t length)

#### 8.53.1 Detailed Description

This file contains the GM API function [gm\\_min\\_size\\_for\\_length](#)().

#### 8.53.2 Function Documentation

##### 8.53.2.1 GM\_ENTRY\_POINT unsigned int gm\_min\_size\_for\_length (gm\_size\_t length)

[gm\\_min\\_size\\_for\\_length](#)() returns the minimum GM message buffer size required to store a message of length **length**.

**Return values:**

*gm\_log2\_roundup*

**Parameters:**

*length* (IN) The length of the message.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.54 gm\_mtu.c File Reference

```
#include "gm.h"
#include "gm-types.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned int [gm\\_mtu](#) (struct gm\_port \*port)

#### 8.54.1 Detailed Description

This file contains the GM API function [gm\\_mtu\(\)](#).

#### 8.54.2 Function Documentation

##### 8.54.2.1 GM\_ENTRY\_POINT unsigned int gm\_mtu (struct gm\_port \* port)

[gm\\_mtu\(\)](#) returns the value of GM\_MTU.

**Return values:**

*GM\_MTU*

**Parameters:**

*port* (IN) The handle to the GM port.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.55 gm\_mutex.c File Reference

```
#include "gm_call_trace.h"
#include "gm_compiler.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT struct gm\_mutex \* [gm\\_create\\_mutex](#) ()
- GM\_ENTRY\_POINT void [gm\\_destroy\\_mutex](#) (struct gm\_mutex \*mu)
- GM\_ENTRY\_POINT void [gm\\_mutex\\_enter](#) (struct gm\_mutex \*mu)
- GM\_ENTRY\_POINT void [gm\\_mutex\\_exit](#) (struct gm\_mutex \*mu)

### 8.55.1 Detailed Description

This file contains the GM API functions [gm\\_create\\_mutex\(\)](#), [gm\\_destroy\\_mutex\(\)](#), [gm\\_mutex\\_enter\(\)](#), [gm\\_mutex\\_exit\(\)](#).

### 8.55.2 Function Documentation

#### 8.55.2.1 GM\_ENTRY\_POINT struct gm\_mutex\* gm\_create\_mutex (void)

[gm\\_create\\_mutex\(\)](#) creates a GM mutex.

**Return values:**

*gm\_mutex* (OUT) Handle to the GM mutex.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.55.2.2 GM\_ENTRY\_POINT void gm\_destroy\_mutex (struct gm\_mutex \* mu)**

[gm\\_destroy\\_mutex\(\)](#) destroys a GM mutex.

**Parameters:**

*mu* (IN) The handle to the GM mutex.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.55.2.3 GM\_ENTRY\_POINT void gm\_mutex\_enter (struct gm\_mutex \* mu)**

[gm\\_mutex\\_enter\(\)](#)

**Parameters:**

*mu* (IN) The handle to the GM mutex.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.55.2.4 GM\_ENTRY\_POINT void gm\_mutex\_exit (struct gm\_mutex \* mu)**

[gm\\_mutex\\_exit\(\)](#)

**Parameters:**

*mu* (IN) The handle to the GM mutex.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.56 gm\_next\_event\_peek.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_next\\_event\\_peek](#) (gm\_port\_t \*p, gm\_u16\_t \*sender)

### 8.56.1 Detailed Description

This file contains the GM API function [gm\\_next\\_event\\_peek](#)().

### 8.56.2 Function Documentation

#### 8.56.2.1 GM\_ENTRY\_POINT int gm\_next\_event\_peek (gm\_port\_t \* p, gm\_u16\_t \* sender)

[gm\\_next\\_event\\_peek](#)() returns the nonzero event type if an event is pending. If the event is a message receive event, then the sender parameter will be filled with the gmID of the message sender. If an event is pending a call to any [gm\\_receive\\*](#)() function will return the event immediately, although [gm\\_receive](#)() is preferred in this case for efficiency.

#### Return values:

*gm\_recv\_event\_t* Receive event type.

#### Parameters:

*p* (IN) The GM port for which the communication is received.  
*sender*

#### See also:

[gm\\_receive\\_pending](#) [gm\\_receive](#)

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.57 gm\_node\_id\_to\_host\_name.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_debug_yp.h"
#include "gm_internal.h"
#include "gm_lanai_command.h"
```

### Functions

- GM\_ENTRY\_POINT char \* [gm\\_node\\_id\\_to\\_host\\_name](#) (gm\_port\_t \*port, unsigned int node\_id)
- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_node\\_id\\_to\\_host\\_name\\_ex](#) (gm\_port\_t \*port, unsigned int timeout\_usec, unsigned int node\_id, char(\*name)[GM\_MAX\_HOST\_NAME\_LEN+1])

### 8.57.1 Detailed Description

This file contains the GM API functions [gm\\_node\\_id\\_to\\_host\\_name\(\)](#) and [gm\\_node\\_id\\_to\\_host\\_name\\_ex\(\)](#).

### 8.57.2 Function Documentation

#### 8.57.2.1 GM\_ENTRY\_POINT char\* [gm\\_node\\_id\\_to\\_host\\_name](#) (gm\_port\_t \*port, unsigned int node\_id)

This function is deprecated.

[gm\\_node\\_id\\_to\\_host\\_name\(\)](#) returns a pointer to the host name of the host containing the network interface card with GM node id node\_id. The name referenced by the returned pointer is only valid until the next GM API call.

#### Return values:

*host\_name* The host name of the host containing the GM node id node\_id.



**Parameters:**

*port* (IN) The GM port associated with the network interface card.

*node\_id* (IN) The GM node id associated with the network interface card.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.57.2.2 GM\_ENTRY\_POINT [gm\\_status.t](#) gm\_node\_id\_to\_host\_name.ex**  
([gm\\_port.t](#) \* *port*, unsigned int *timeout\_usecs*, unsigned int *node\_id*,  
char \* *name*[GM\_MAX\_HOST\_NAME\_LEN+1])

Store at \*name the host name of the host containing the network interface card with GM node id node\_id. The buffer pointed to by name must have at least GM\_MAX\_HOST\_NAME\_LEN+1 bytes of storage.

**Return values:**

*host\_name* The host name of the host containing the GM node id node\_id.

**Parameters:**

*timeout\_usecs* (IN) The maximum length of time (in microseconds) to query, or 0 to use the default (long) timeout.

*port* (IN) The GM port associated with the network interface card.

*node\_id* (IN) The GM node id associated with the network interface card.

*name* (OUT) Where to store the pointer to the host name.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.58 gm\_node\_id\_to\_unique\_id.c File Reference

```
#include "gm_call_trace.h"
#include "gm_internal.h"
#include "gm_global_id.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_node\\_id\\_to\\_unique\\_id](#) (gm\_port\_t \*port, unsigned int node\_id, char \*uid)

### 8.58.1 Detailed Description

This file contains the GM API function gm\_global\_id\_to\_unique\_id().

### 8.58.2 Function Documentation

#### 8.58.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_node\\_id\\_to\\_unique\\_id](#) (gm\_port\_t \*port, unsigned int node\_id, char \*uid)

gm\_global\_id\_to\_unique\_id() stores the MAC address for the interface with GM ID **n** at **unique**.

#### Return values:

*GM\_SUCCESS* Operation completed successfully.

*GM\_INVALID\_PARAMETER*

#### Parameters:

*port* (IN) A GM port.

*node\_id* (IN) The node ID to convert.

*uid* (OUT) Where to store the unique ID corresponding to node\_id.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.59 gm\_num\_ports.c File Reference

```
#include "gm_compiler.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned int [gm\\_num\\_ports](#) (gm\_port\_t \*p)

#### 8.59.1 Detailed Description

This file contains the GM API function [gm\\_num\\_ports\(\)](#).

#### 8.59.2 Function Documentation

##### 8.59.2.1 GM\_ENTRY\_POINT unsigned int gm\_num\_ports (gm\_port\_t \* p)

[gm\\_num\\_ports\(\)](#) returns the number of ports supported by this build. Note that this is a build-time value, and the 'p' parameter is actually meaningless at the present time.

This value is accessed through the function API to avoid hard-coding it in user applications, thus allowing dynamic library upgrades.

#### Return values:

*GM\_NUM\_PORTS*

#### Parameters:

*p* (IN) The GM port.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.60 gm\_num\_receive\_tokens.c File Reference

```
#include "gm_compiler.h"
#include "gm_debug.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned int [gm\\_num\\_receive\\_tokens](#) (gm\_port\_t \*p)

### 8.60.1 Detailed Description

This file contains the GM API function [gm\\_num\\_receive\\_tokens](#)().

### 8.60.2 Function Documentation

#### 8.60.2.1 GM\_ENTRY\_POINT unsigned int gm\_num\_receive\_tokens (gm\_port\_t \*p)

[gm\\_num\\_receive\\_tokens](#)() returns the number of receive tokens for this port. This value is accessed through the function API to avoid hard-coding it in user applications, allowing dynamic library upgrades.

#### Return values:

*GM\_NUM\_RECV\_TOKENS*

#### Parameters:

*p* (IN) The GM port.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.61 gm\_num\_send\_tokens.c File Reference

```
#include "gm_compiler.h"
#include "gm_debug.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned int [gm\\_num\\_send\\_tokens](#) (gm\_port\_t \*p)

#### 8.61.1 Detailed Description

This file contains the GM API function [gm\\_num\\_send\\_tokens](#)().

#### 8.61.2 Function Documentation

##### 8.61.2.1 GM\_ENTRY\_POINT unsigned int gm\_num\_send\_tokens (gm\_port\_t \**p*)

[gm\\_num\\_send\\_tokens](#)() returns the number of send tokens for this port. This value is accessed through the function API to avoid hard-coding it in user applications, allowing dynamic library upgrades.

#### Return values:

*GM\_NUM\_SEND\_TOKENS*

#### Parameters:

*p* (IN) The GM port.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.62 gm\_on\_exit.c File Reference

```
#include "gm.h"
#include "gm_call_trace.h"
#include "gm_config.h"
#include "gm_internal_funcs.h"
```

### Data Structures

- struct [gm\\_on\\_exit\\_record](#)

### Typedefs

- typedef [gm\\_on\\_exit\\_record](#) [gm\\_on\\_exit\\_record\\_t](#)

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t](#) [gm\\_on\\_exit](#) ([gm\\_on\\_exit\\_callback\\_t](#) callback, void \*arg)

#### 8.62.1 Detailed Description

This file contains the GM API function [gm\\_on\\_exit\(\)](#).

#### 8.62.2 Typedef Documentation

##### 8.62.2.1 typedef struct [gm\\_on\\_exit\\_record](#) [gm\\_on\\_exit\\_record\\_t](#)

List element storing the details of a callback that should be called upon exit.

#### 8.62.3 Function Documentation

### 8.62.3.1 GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_on\_exit ([gm\\_on\\_exit\\_callback\\_t](#) *callback*, void \* *arg*)

[gm\\_on\\_exit\(\)](#) is like Linux `on_exit()`. This function registers a callback so that 'CALLBACK(STATUS,ARG)' is called when the program exits. Callbacks are called in the reverse of the order of registration. This function is also somewhat similar to BSD 'atexit()'.

Call the callbacks in the reverse order registered inside [gm\\_exit\(\)](#), passing GM exit status and registered argument to the callback.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

*GM\_OUT\_OF\_MEMORY*

**Parameters:**

*callback*

*arg*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.63 gm\_open.c File Reference

```
#include "gm_call_trace.h"
#include "gm_compiler.h"
#include "gm_debug.h"
#include "gm_debug_open.h"
#include "gm_internal.h"
#include "gm_internal_funcs.h"
#include "gm_enable_ethernet.h"
#include "gm_enable_security.h"
#include "gm_enable_trace.h"
#include "gm_ptr_hash.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_open](#) ([gm\\_port\\_t \\*\\*port\\_p](#), unsigned *unit*, unsigned *port\_id*, const char \**client\_type*, enum [gm\\_api\\_version api\\_version](#))

### 8.63.1 Detailed Description

The file containing the GM API function [gm\\_open\(\)](#).

### 8.63.2 Function Documentation

**8.63.2.1** GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_open](#) ([gm\\_port\\_t \\*\\* port\\_p](#), unsigned *unit*, unsigned *port\_id*, const char \* *client\_type*, enum [gm\\_api\\_version api\\_version](#))

[gm\\_open\(\)](#) opens a GM port **port\_p** for LANai interface **unit**, a pointer to the port's state at **\*port\_p**. The pointer must be passed to all subsequent functions that operate on the opened port. **port\_id** is a null-terminated ASCII string that is used to identify the port client for debugging (and potentially other) purposes; pass in the name of your



program. Note that unit and port numbers start at 0, and that ports 0 (internal use) and 1 (mapper) and 3 (ethernet emulation) are reserved, so clients should use port 2 and ports 4-7.

[gm\\_open\(\)](#) is to be called by clients other than the daemon and mapper.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

*GM\_INVALID\_PARAMETER* Invalid parameter passed.

*GM\_BUSY* Could not open device.

*GM\_NO\_SUCH\_DEVICE* `port_id > GM_NUM_PORTS - 1`.

*GM\_INCOMPATIBLE\_LIB\_AND\_DRIVER* The GM user library linked with this program may not be compatible with the installed driver.

*GM\_OUT\_OF\_MEMORY* Could not allocate storage for user port.

**Parameters:**

*port\_p* (IN) Pointer to the handle of the GM port.

*unit* (IN) The device for the Myrinet interface. = 0 if device is myri0.

*port\_id* (OUT) The id of the GM port that is opened. = 2, 4, 5, 6, or 7 (Ports 0, 1, and 3 are for privileged use only.)

*client\_type* (IN) Unused.

*api\_version* (IN) `GM_API_VERSION = GM_API_VERSION` as defined in [gm.h](#)

**See also:**

[gm\\_close](#) [gm\\_init](#) [gm\\_abort](#) [gm\\_finalize](#) [gm\\_exit](#)

**Author:**

Glenn Brown

**Version:**

`GM_API_VERSION` (as defined in [gm.h](#))

## 8.64 gm\_page\_alloc.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_internal.h"
#include "gm_malloc_debug.h"
```

### Data Structures

- struct [gm\\_free\\_page](#)
- struct [gm\\_page\\_allocation\\_record](#)

### Functions

- GM\_ENTRY\_POINT void \* [gm\\_page\\_alloc](#) ()
- GM\_ENTRY\_POINT void [gm\\_page\\_free](#) (void \*ptr)

#### 8.64.1 Detailed Description

This file contains [gm\\_page\\_alloc\(\)](#) and [gm\\_page\\_free\(\)](#).

These GM API functions allow pages to be allocated and freed.

These modules implement a platform-independent page allocation interface, with automatic initialization and finalization for use in the kernel.

In some cases, this file uses platform-specific features to implement page allocation efficiently, but by default this module allocates blocks of pages and keeps a page free list to recycle freed pages, but the module never frees the blocks of pages until all pages have been freed.

#### 8.64.2 Function Documentation

##### 8.64.2.1 GM\_ENTRY\_POINT void\* gm\_page\_alloc (void)

[gm\\_page\\_alloc\(\)](#) returns a ptr to a newly allocated page-aligned buffer of length GM\_PAGE\_LEN.

**Return values:**

*ptr* Page-aligned buffer of length GM\_PAGE\_LEN.

**See also:**

[gm\\_page\\_free](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.64.2.2 GM\_ENTRY\_POINT void gm\_page\_free (void \* ptr)**

[gm\\_page\\_free\(\)](#) frees the page of memory at **ptr** previously allocated by [gm\\_page\\_alloc\(\)](#). If all pages have been freed, free all of the memory allocated for pages.

**Parameters:**

*ptr* (IN) Address of the memory page to be freed.

**See also:**

[gm\\_page\\_alloc](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.65 gm\_perror.c File Reference

```
#include "gm_debug.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_perror](#) (const char \*message, [gm\\_status\\_t](#) error)

#### 8.65.1 Detailed Description

This file contains the GM API function [gm\\_perror\(\)](#).

#### 8.65.2 Function Documentation

##### 8.65.2.1 GM\_ENTRY\_POINT void [gm\\_perror](#) (const char \* *message*, [gm\\_status\\_t](#) *error*)

[gm\\_perror\(\)](#) is similar to ANSI perror(), but takes the error code as a parameter to allow thread safety in future implementations, and only supports GM error numbers. Prints *message* followed by a description of **errno**.

##### Parameters:

*message* (OUT) Textual description of the GM error.

*error* (IN) GM Error code.

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.66 gm\_printf.c File Reference

```
#include "gm_config.h"
#include "gm.h"
#include <stdio.h>
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_printf](#) (const char \*format,...)

#### 8.66.1 Detailed Description

#### 8.66.2 Function Documentation

##### 8.66.2.1 GM\_ENTRY\_POINT int gm\_printf (const char \*format, ...)

[gm\\_printf](#)() emulates or invokes the ANSI standard printf() function.

**Return values:**

- 0 Operation completed successfully.

**Parameters:**

*format* Specifies how the arguments are converted for output.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.67 gm\_provide\_receive\_buffer.c File Reference

```
#include "gm_debug.h"
#include "gm_debug_recv_tokens.h"
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_provide\\_receive\\_buffer\\_with\\_tag](#) (gm\_port\_t \*p, void \*ptr, unsigned size, unsigned priority, unsigned int tag)

### 8.67.1 Detailed Description

This file contains the GM API function [gm\\_provide\\_receive\\_buffer\\_with\\_tag](#)().

### 8.67.2 Function Documentation

#### 8.67.2.1 GM\_ENTRY\_POINT void gm\_provide\_receive\_buffer\_with\_tag (gm\_port\_t \*p, void \*ptr, unsigned size, unsigned priority, unsigned int tag)

[gm\\_provide\\_receive\\_buffer\\_with\\_tag](#)() provides GM with a buffer into which it can receive messages with matching **size** and **priority** fields. It is the client software's responsibility to provide buffers of each **size** and **priority** that might be received; not doing so can cause program deadlock, which will eventually result in a port being closed after a timeout. This timeout is a function of the number of packets sent.

The client software may provide up to [gm\\_num\\_receive\\_tokens](#)() different receive buffers into which messages may be received.

Each buffer provided by the client software to GM via this function will be used only once to receive a message. In other words, calling [gm\\_provide\\_receive\\_buffer\\_with\\_tag](#)(port,buffer,size,priority,tag) provides GM a token to receive a single message of size **size** and priority **priority** into the receive buffer **buffer**. When a message is eventually received into this buffer, [gm\\_receive](#)(port) stores the buffer pointer **buffer** and

**tag** in the returned event, returning control of the buffer (token) to the client software. If the client software wishes for the buffer to be reused for a similar receive, it must call [gm\\_provide\\_receive\\_buffer\\_with\\_tag\(\)](#) again with the same or similar parameters.

Once a buffer has been provided to GM, its content should not be changed until control of the buffer has been returned to the client software via [gm\\_receive\(\)](#).

The **tag** parameter must be in the range [0,255], and is returned in the receive event describing a receive into a buffer. It may be used in any way the client desires, and need not be unique.

**Parameters:**

- p* (IN) The GM port.
- ptr* (IN) The address of the message communicated.
- size* (IN) The size of the message.
- priority* (IN) The priority of the message.
- tag* (OUT) The tag for a receive event queue.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.68 gm\_put.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_enable_put.h"
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_put](#) (gm\_port\_t \*p, void \*source\_buffer, [gm\\_remote\\_ptr\\_t](#) target\_buffer, unsigned long len, enum [gm\\_priority](#) priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

### 8.68.1 Detailed Description

This file contains the GM API function [gm\\_put\(\)](#).

### 8.68.2 Function Documentation

**8.68.2.1** GM\_ENTRY\_POINT void [gm\\_put](#) (gm\_port\_t \*p, void \*source\_buffer, [gm\\_remote\\_ptr\\_t](#) target\_buffer, unsigned long len, enum [gm\\_priority](#) priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

[gm\\_put\(\)](#) transfers the **len** bytes at **source\_buffer** to **target\_port\_id** on **target\_node\_id** with priority **priority** and stores the data at the remote virtual memory address **target\_buffer**. Call [callback\(port,context,status\)](#) when the send completes or fails, with **status** indicating the status of the send. The order of the transfer is preserved relative to messages of the same priority sent using [gm\\_send\(\)](#) or [gm\\_send\\_to\\_peer\(\)](#).

#### Parameters:

- p** (IN) The GM port on the source/sender GM node from which the communication is being sent.



*source\_buffer* (IN) Address of the send buffer.

*target\_buffer* (OUT) Address of the receive buffer.

*len* (IN) The length in bytes of the buffer to be sent.

*priority* (IN) The priority of the data being sent.

*target\_node\_id* (IN) The GM node to which the data is being sent.

*target\_port\_id* (IN) The GM port on the destination GM node to which the message is being sent.

*callback* (IN) The function called when the send is complete.

*context* (IN) Pointer to an integer or to a structure that is passed to the callback function.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.69 gm\_rand.c File Reference

```
#include "gm_internal.h"
#include "gm_crc32.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_rand](#) ()
- GM\_ENTRY\_POINT void [gm\\_srand](#) (int seed)

### 8.69.1 Detailed Description

This file contains the GM API functions [gm\\_rand](#)() and [gm\\_srand](#)().

### 8.69.2 Function Documentation

#### 8.69.2.1 GM\_ENTRY\_POINT int gm\_rand (void)

[gm\\_rand](#)() returns a pseudo-random integer, using a poor but fast random number generator.

**Return values:**

*RANDOM\_NUMBER* The random number that was generated.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

#### 8.69.2.2 GM\_ENTRY\_POINT void gm\_srand (int seed)

[gm\\_srand](#)() returns a pseudo-random integer, and requires a seed for the random number generator.

**Parameters:**

*seed* (IN) Seed for the random number generator.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.70 gm\_rand\_mod.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT unsigned int [gm\\_rand\\_mod](#) (unsigned int *a*)

#### 8.70.1 Detailed Description

#### 8.70.2 Function Documentation

##### 8.70.2.1 GM\_ENTRY\_POINT unsigned int gm\_rand\_mod (unsigned int *a*)

[gm\\_rand\\_mod](#)() returns a pseudo-random number modulo **modulus**, using a poor but fast random number generator.

**Return values:**

*RANDOM\_NUMBER*

**Parameters:**

*a* (IN) The modulus bound.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.71 gm\_receive.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_debug_recv_tokens.h"
#include "gm_enable_trace.h"
#include "gm_internal.h"
#include "gm_tick.h"
```

### Functions

- GM\_ENTRY\_POINT gm\_recv\_event\_t \* [gm\\_receive](#) (gm\_port\_t \*p)
- GM\_ENTRY\_POINT gm\_recv\_event\_t \* [gm\\_receive\\_debug\\_buffers](#) (gm\_port\_t \*port)

### 8.71.1 Detailed Description

This file contains the GM API functions [gm\\_receive\(\)](#) and [gm\\_receive\\_debug\\_buffers\(\)](#).

### 8.71.2 Function Documentation

#### 8.71.2.1 GM\_ENTRY\_POINT gm\_recv\_event\_t\* gm\_receive (gm\_port\_t \* p)

[gm\\_receive\(\)](#) returns a receive event. If no significant receive event is pending, then an event of type GM\_NO\_RECV\_EVENT is immediately returned.

#### Return values:

*gm\_recv\_event\_t*

#### Parameters:

*p* (IN) The GM port for which the communication is received.

#### See also:

[gm\\_receive\\_pending](#)

**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))**8.71.2.2 GM\_ENTRY\_POINT gm\_recv\_event\_t\* gm\_receive\_debug\_buffers  
(gm\_port\_t \* port)**[gm\\_receive\\_debug\\_buffers\(\)](#)**Return values:***gm\_recv\_event\_t***Parameters:***port* (IN) The GM port for which the communication is received.**See also:**[gm\\_receive\\_pending](#) [gm\\_receive](#)**Author:**

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

## 8.72 gm\_receive\_pending.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_receive\\_pending](#) (gm\_port\_t \*p)

#### 8.72.1 Detailed Description

This file contains the GM API function [gm\\_receive\\_pending\(\)](#).

#### 8.72.2 Function Documentation

##### 8.72.2.1 GM\_ENTRY\_POINT int gm\_receive\_pending (gm\_port\_t \* p)

[gm\\_receive\\_pending\(\)](#) returns nonzero if a receive event is pending. If a receive event is pending, a call to any [gm\\_receive\\*\(\)](#) function will return the event immediately, although [gm\\_receive\(\)](#) is preferred in this case for efficiency.

##### Return values:

*event.recv.type*

##### Parameters:

*p* (IN) The GM port for which the received communication is pending.

##### See also:

[gm\\_receive](#)

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.73 gm\_register.c File Reference

```
#include "gm_call_trace.h"
#include "gm_compiler.h"
#include "gm_debug.h"
#include "gm_internal.h"
#include "gm_io.h"
#include "gm_debug_mem_register.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_register\\_memory\\_ex](#) (gm\_port\_t \*p, void \*ptr, gm\_size\_t length, void \*pvma)
- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_register\\_memory](#) (gm\_port\_t \*p, void \*ptr, gm\_size\_t length)

### 8.73.1 Detailed Description

This file contains the GM API function [gm\\_register\\_memory\(\)](#).

### 8.73.2 Function Documentation

#### 8.73.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_register\\_memory\\_ex](#) (gm\_port\_t \*p, void \* ptr, gm\_size\_t length, void \* pvma)

[gm\\_register\\_memory\\_ex\(\)](#) registers **length** bytes of user virtual memory starting at **ptr** for DMA transfers, associating the memory with port virtual address **pvma**. The memory is locked down (made nonpageable) and DMAs on the region of memory are enabled. Memory may be registered multiple times. Memory may be deregistered using matching calls to [gm\\_deregister\\_memory\(\)](#). Note that memory registration is an expensive operation relative to sending and receiving packets, so you should use persistent memory registrations wherever possible. Also note that memory registration is not supported on Solaris 2.7 and earlier due to operating system limitations.



Note that **pvma** must be used in all subsequent GM API calls to refer to the registered memory region.

`gm_register_memory(p,ptr,len)` is equivalent to `gm_register_memory_ex(p,ptr,len,ptr)`.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

*GM\_FAILURE*

*GM\_PERMISSION\_DENIED*

*GM\_INVALID\_PARAMETER*

*GM\_OUT\_OF\_MEMORY*

**Parameters:**

*p* (IN) Handle to the GM port.

*ptr* (IN) The address of the memory region to be registered.

*length* (IN) The length in bytes of the memory region to be registered.

*pvma* (IN) The port virtual memory address with which to associate this region.

**See also:**

[gm\\_register\\_memory](#) , [gm\\_deregister\\_memory](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION\_2\_0\_6

**8.73.2.2 GM\_ENTRY\_POINT [gm\\_status\\_t](#) gm\_register\_memory (gm\_port\_t \* p, void \* ptr, gm\_size\_t length)**

[gm\\_register\\_memory](#)() registers **length** bytes of user virtual memory starting at **ptr** for DMA transfers. The memory is locked down (made nonpageable) and DMAs on the region of memory are enabled. Memory may be registered multiple times. Memory may be deregistered using matching calls to [gm\\_deregister\\_memory](#)(). Note that memory registration is an expensive operation relative to sending and receiving packets, so you should use persistent memory registrations wherever possible. Also note that memory registration is not supported on Solaris 2.7 and earlier due to operating system limitations.

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

*GM\_FAILURE*

*GM\_PERMISSION\_DENIED*

*GM\_INVALID\_PARAMETER*

*GM\_OUT\_OF\_MEMORY*

**Parameters:**

*p* (**IN**) Handle to the GM port.

*ptr* (**OUT**) The address of the memory location to be registered.

*length* (**IN**) The length in bytes of the memory location to be registered.

**See also:**

[gm\\_deregister\\_memory](#) , [gm\\_register\\_memory\\_ex](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.74 gm\_resume\_sending.c File Reference

```
#include "gm_internal.h"
```

```
#include "_gm_modsend.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_resume\\_sending](#) (struct gm\_port \*p, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

#### 8.74.1 Detailed Description

The file contains the GM API function [gm\\_resume\\_sending\(\)](#).

#### 8.74.2 Function Documentation

**8.74.2.1 GM\_ENTRY\_POINT void gm\_resume\_sending (struct gm\_port \* p, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, gm\_send\_completion\_callback\_t callback, void \* context)**

[gm\\_resume\\_sending](#)(0 reenables packet transmission of messages from **port** of priority **priority** destined for target\_port\_id of **target\_node\_id**. This function should only be called after an error is reported to a send completion callback routine. The message that generated the error is not resent. The first four parameters must match those of the failed send. It should be called only once per reported error. This function requires a send token, which will be returned to the client in the callback function.

[gm\\_resume\\_sending\(\)](#) and [gm\\_drop\\_sends\(\)](#), as most gm requests, require a send token, and the callback you give to them is just meant to return this token. These gm\_requests always succeed (if called in a valid manner), so the callback will **always** be called with **GM.SUCCESS** (which here does not mean at all that something was sent successfully, just that the request has been taken into account, and the token used for that request was recycled).

#### Parameters:

*p* (IN) The handle to the GM port.

***priority*** (IN) The priority of the message being sent.

***target\_node\_id*** (IN) The GM node to which the message is being sent.

***target\_port\_id*** (IN) The GM port on the destination GM node to which the message is being sent.

***callback*** (IN) The function called when the send is complete.

***context*** (IN) Pointer to an integer or to a structure that is passed to the callback function.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.75 gm\_send.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_enable_fast_small_send.h"
#include "gm_enable_trace.h"
#include "gm_internal.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_send\\_with\\_callback](#) (gm\_port\_t \*p, void \*message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

#### 8.75.1 Detailed Description

This file contains the GM API functions [gm\\_send\\_with\\_callback\(\)](#).

#### 8.75.2 Function Documentation

**8.75.2.1** GM\_ENTRY\_POINT void [gm\\_send\\_with\\_callback](#) (gm\_port\_t \*p, void \*message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, unsigned int target\_port\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

[gm\\_send\\_with\\_callback\(\)](#) is a fully asynchronous send. It queues the **message** of length **len** to be sent with priority **priority** to node **target\_node\_id**. As GM is event-based, the effective completion of the send is notified to the client software by the execution of the callback function specified by **callback**. Before calling [gm\\_send\\_with\\_callback\(\)](#), the client software must first possess a send token of the same priority, and by calling [gm\\_send\\_with\\_callback\(\)](#) the client implicitly relinquishes this send token. After a call to [gm\\_send\\_with\\_callback\(..., message, len, ...\)](#), the memory specified by **message** and

**len** must not be modified until the send completes. After the send completes, `callback(port,context,status)` will be called inside `gm_unknown()`, with **status** indicating the status of the completed send.

**Note:** The order of messages with different priorities or with different destination ports is not preserved. Only the order of messages with the same priority and to the same destination port is preserved.

In the special case that the **target\_port\_id** is the same as the sending port ID (as is often the case), the streamlined `gm_send_to_peer_with_callback()` function may be used instead of `gm_send_with_callback()`, allowing the **target\_port\_id** parameter to be omitted, and slightly improving small-message performance.

**Parameters:**

***p*** (IN) A pointer to the GM port on the source/sender GM node over which the message is to be sent.

***message*** (IN) A pointer to the data to be sent.

***size*** (IN) The size receive buffer in which to store the message on the remote node.

***len*** (IN) The length (in bytes) of the message to be sent.

***priority*** (IN) The priority with which to send the message ('GM\_HIGH\_PRIORITY' or 'GM\_LOW\_PRIORITY').

***target\_node\_id*** (IN) The ID of the GM node to which the message should be sent.

***target\_port\_id*** (IN) The ID of the GM port to which the message should be sent.

***callback*** (IN) The client function to call when the send completes.

***context*** (IN) A pointer to pass to the CALLBACK function when it is called.

**See also:**

[gm\\_send\\_to\\_peer\\_with\\_callback](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.76 gm\_send\_to\_peer.c File Reference

```
#include "gm_internal.h"
#include "gm_enable_fast_small_send.h"
#include "gm_send_queue.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_send\\_to\\_peer\\_with\\_callback](#) (gm\_port\_t \*p, void \*message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

#### 8.76.1 Detailed Description

This file contains the GM API function [gm\\_send\\_to\\_peer\\_with\\_callback](#)().

#### 8.76.2 Function Documentation

**8.76.2.1** GM\_ENTRY\_POINT void [gm\\_send\\_to\\_peer\\_with\\_callback](#) (gm\_port\_t \*p, void \*message, unsigned int size, gm\_size\_t len, unsigned int priority, unsigned int target\_node\_id, [gm\\_send\\_completion\\_callback\\_t](#) callback, void \*context)

[gm\\_send\\_to\\_peer\\_with\\_callback](#)() is an asynchronous send like [gm\\_send\\_with\\_callback](#)(), only with the **target\_node\_id** implicitly set to the same ID as **port** (sending to the same port). This function is marginally faster than [gm\\_send\\_with\\_callback](#)().

##### Parameters:

- p** (IN) A pointer to the GM port over which the message is to be sent.
- message** (IN) A pointer to the data to be sent.
- size** (IN) The size receive buffer in which to store the message on the remote node.
- len** (IN) The length (in bytes) of the message to be sent.
- priority** (IN) The priority with which to send the message ('GM.HIGH-PRIORITY' or 'GM.LOW-PRIORITY').

*target\_node\_id* (IN) The ID of the GM node to which the message should be sent.

*callback* (IN) The client function to call when the send completes.

*context* (IN) A pointer to pass to the CALLBACK function when it is called.

**See also:**

[gm\\_send\\_with\\_callback](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))



## 8.77 gm\_send\_token\_available.c File Reference

```
#include "gm_internal.h"
```

### Functions

- int [gm\\_send\\_token\\_available](#) (gm\_port\_t \*p, unsigned priority)

#### 8.77.1 Detailed Description

This file contains the GM API function [gm\\_send\\_token\\_available\(\)](#).

#### 8.77.2 Function Documentation

##### 8.77.2.1 int gm\_send\_token\_available (gm\_port\_t \* p, unsigned priority)

[gm\\_send\\_token\\_available\(\)](#) tests for the availability of a send token without allocating the send token. It is similar to the function [gm\\_alloc\\_send\\_token\(port,priority\)](#) without the allocation.

##### Return values:

*int*

##### Parameters:

*p* (IN) The GM port on the source/sender GM node from which the communication would be sent.

*priority* (IN) The priority of the message to be sent.

##### See also:

[gm\\_alloc\\_send\\_token](#)

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.78 gm\_set\_acceptable\_sizes.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_set\\_acceptable\\_sizes](#) (struct gm\_port \*p, enum [gm\\_priority](#) priority, gm\_size\_t mask)

#### 8.78.1 Detailed Description

This file contains the GM API function [gm\\_set\\_acceptable\\_sizes\(\)](#).

#### 8.78.2 Function Documentation

##### 8.78.2.1 GM\_ENTRY\_POINT [gm\\_status\\_t gm\\_set\\_acceptable\\_sizes](#) (struct gm\_port \*p, enum [gm\\_priority](#) priority, gm\_size\_t mask)

[gm\\_set\\_acceptable\\_sizes\(\)](#) informs GM of the acceptable sizes of GM messages received on port **p** with priority **priority**. Each set bit of **mask** indicates an acceptable size. While calling this function is not required, clients should call it during program initialization to detect errors involving the reception of badly sized messages to be reported nearly instantaneously, rather than after a substantial delay of 30 seconds or more.

Note: the MASK is a long to support larger than 2GByte packets (those with size larger than 31).

##### Return values:

**GM\_SUCCESS** Operation completed successfully.

**GM\_PERMISSION\_DENIED** Port number hasn't been set.

**GM\_INTERNAL\_ERROR** LANai is not running.

**GM\_INVALID\_PARAMETER** The priority has an invalid value.

##### Parameters:

**p** (IN) The GM port in use.

*priority* (IN) The priority of the message.

*mask*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.79 gm\_set\_alarm.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_cmp64.h"
#include "gm_internal.h"
#include "gm_send_queue.h"
#include "gm_set_alarm.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_cancel\\_alarm](#) (gm\_alarm\_t \*my\_alarm)
- GM\_ENTRY\_POINT void [gm\\_initialize\\_alarm](#) (gm\_alarm\_t \*my\_alarm)
- GM\_ENTRY\_POINT void [gm\\_set\\_alarm](#) (gm\_port\_t \*port, gm\_alarm\_t \*my\_alarm, gm\_u64\_t usecs, void(\*callback)(void \*), void \*context)

### 8.79.1 Detailed Description

This file includes source for the user-level alarm API calls [gm\\_initialize\\_alarm\(\)](#), [gm\\_set\\_alarm\(\)](#), and [gm\\_cancel\\_alarm\(\)](#).

The alarm API allows the GM client to schedule a callback function to be called after a delay, specified in microseconds. An unbounded number of alarms may be set, although alarm overhead increases linearly in the number of set alarms, and the client must provide storage for each set alarm.

These source for all the functions is included here, because if the user needs one, the user needs them all.

This code is a bit tricky, since the LANai provides only one alarm per port. This alarm is used to time the first pending alarm, but if any new alarm is set, then the LANai alarm must be flushed to find out what time it is and potentially to allow the alarm to be rescheduled for an earlier time. Therefore, unless no other alarm is set, alarms are set by adding them to an "unset alarm" queue, flushing any pending alarm, and then adding the alarms to the sorted alarm queue only once the pending LANai alarm has been flushed or triggered.

### 8.79.2 Function Documentation

**8.79.2.1 GM\_ENTRY\_POINT void gm\_cancel\_alarm (gm\_alarm\_t \* my\_alarm)**

[gm\\_cancel\\_alarm\(\)](#) cancels a scheduled alarm, or does nothing if an alarm is not scheduled. Alarms are cancelled by simply removing them from the alarm list so that when the next GM\_ALARM\_EVENT arrives from the LANai [\\_gm\\_handle\\_alarm\(\)](#) will not find the alarm.

**Parameters:**

*my\_alarm* (IN) Alarm to be cancelled.

**See also:**

[gm\\_initialize\\_alarm](#) [gm\\_set\\_alarm](#) [gm\\_flush\\_alarm](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.79.2.2 GM\_ENTRY\_POINT void gm\_initialize\_alarm (gm\_alarm\_t \* my\_alarm)**

[gm\\_initialize\\_alarm\(\)](#) initializes a client-allocated **gm\_alarm\_t** structure for use with [gm\\_set\\_alarm\(\)](#). This function should be called after the structure is allocated but before a pointer to it is passed to [gm\\_set\\_alarm\(\)](#) or [gm\\_cancel\\_alarm\(\)](#).

**Parameters:**

*my\_alarm* (IN) Alarm to be initialized.

**See also:**

[gm\\_cancel\\_alarm](#) [gm\\_set\\_alarm](#) [gm\\_flush\\_alarm](#)

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

### 8.79.2.3 GM\_ENTRY\_POINT void gm\_set\_alarm (gm\_port\_t \* port, gm\_alarm\_t \* my\_alarm, gm\_u64\_t usecs, void(\* callback)(void \*), void \* context)

[gm\\_set\\_alarm\(\)](#) sets an alarm, which may already be pending. If it is pending, it is rescheduled. The user supplies MY\_ALARM, a pointer to storage for the alarm state, which has been initialized with gm\_init\_alarm\_my\_alarm(). When the alarm occurs, CALLBACK(CONTEXT) is called.

[gm\\_set\\_alarm\(\)](#) schedules CALLBACK(CONTEXT) to be called after USEC microseconds (or later), or reschedule the alarm if it has already been scheduled and has not yet triggered. CALLBACK must be non-NULL. CONTEXT is treated as an opaque pointer by GM, and will be passed as the single parameter to the client-supplied CALLBACK function.

GM clients will also be able to take advantage of the fact that an application is guaranteed to receive a single GM\_ALARM\_EVENT for each call to a client-supplied callback, with the corresponding callback occurring during the call to [gm\\_unknown\(\)](#) that processes that alarm. This means that a case statement like the following in the client's event loops can be used to significantly reduce the overhead of polling for any effect of a client supplied alarm callback:

```
case GM_ALARM_EVENT:
 gm_unknown (event);

 poll for the effect of alarm callbacks only here

break;
```

#### Parameters:

*port* (IN) A pointer to the GM Port.  
*my\_alarm* (IN) Alarm to be initialized.  
*usecs*  
*callback*  
*context*

#### See also:

[gm\\_cancel\\_alarm](#) [gm\\_initialize\\_alarm](#) [gm\\_flush\\_alarm](#)

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.80 gm\_set\_enable\_nack\_down.c File Reference

```
#include "gm_internal.h"
#include "gm_debug.h"
```

### Functions

- [gm\\_status\\_t gm\\_set\\_enable\\_nack\\_down](#) (struct gm\_port \*port, int flag)

#### 8.80.1 Detailed Description

This file contains the GM API function [gm\\_set\\_enable\\_nack\\_down\(\)](#).

#### 8.80.2 Function Documentation

##### 8.80.2.1 [gm\\_status\\_t gm\\_set\\_enable\\_nack\\_down](#) (struct gm\_port \* port, int flag)

[gm\\_set\\_enable\\_nack\\_down\(\)](#)

**Return values:**

*GM\_SUCCESS* Operation completed successfully.

*GM\_PERMISSION\_DENIED*

**Parameters:**

*port* (IN) Handle to the GM port.

*flag*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.81 gm\_simple\_example.h File Reference

```
#include "gm.h"
```

### 8.81.1 Detailed Description

Include file for gm\_simple\_example\_send.c and gm\_simple\_example\_recv.c.



## 8.82 gm\_sleep.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_sleep](#) (unsigned int seconds)

#### 8.82.1 Detailed Description

This file contains the GM API function [gm\\_sleep](#)().

#### 8.82.2 Function Documentation

##### 8.82.2.1 GM\_ENTRY\_POINT int gm\_sleep (unsigned int *seconds*)

[gm\\_sleep](#)() emulates the ANSI standard [sleep](#)(), sleeping the entire process for **seconds** seconds.

**Return values:**

*SLEEP*

**Parameters:**

*seconds* (**IN**) The number of seconds for which the process should sleep.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.83 gm\_strcmp.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_strcmp](#) (const char \*a, const char \*b)

#### 8.83.1 Detailed Description

This file contains the GM API function [gm\\_strcmp](#)().

#### 8.83.2 Function Documentation

##### 8.83.2.1 GM\_ENTRY\_POINT int gm\_strcmp (const char \* a, const char \* b)

[gm\\_strcmp](#)() reimplements strcmp().

##### Return values:

*int* Returns an integer less than, equal to, or greater than zero if a is found, respectively, to be less than, to match, or be greater than b.

##### Parameters:

- a* The first string to be compared.
- b* The second string to be compared.

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.84 gm\_strdup.c File Reference

```
#include "gm.h"
```

### Functions

- GM\_ENTRY\_POINT char \* [gm\\_strdup](#) (const char \*in)

#### 8.84.1 Detailed Description

#### 8.84.2 Function Documentation

##### 8.84.2.1 GM\_ENTRY\_POINT char\* gm\_strdup (const char \* in)

[gm\\_strdup](#)() reimplements the UNIX function `strdup()`.

#### Return values:

*char* \* Returns a pointer to the duplicated string, or NULL if insufficient memory was available.

#### Parameters:

*in* (IN) The string to be duplicated.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.85 gm\_strerror.c File Reference

```
#include "gm.h"
```

### Functions

- GM\_ENTRY\_POINT char \* [gm\\_strerror](#) ([gm\\_status\\_t](#) error)

#### 8.85.1 Detailed Description

This file contains the GM API function [gm\\_strerror](#)().

#### 8.85.2 Function Documentation

##### 8.85.2.1 GM\_ENTRY\_POINT char\* [gm\\_strerror](#) ([gm\\_status\\_t](#) error)

[gm\\_strerror](#)() is an error function for GM. The error is only valid until next call to this function.

**Return values:**

*char*

**Parameters:**

*error* (IN) GM status code.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.86 gm\_strlen.c File Reference

```
#include "gm_internal.h"
#include "gm_debug.h"
```

### Functions

- GM\_ENTRY\_POINT gm\_size\_t [gm\\_strlen](#) (const char \*cptr)

#### 8.86.1 Detailed Description

This file contains the GM API function [gm\\_strlen](#)().

#### 8.86.2 Function Documentation

##### 8.86.2.1 GM\_ENTRY\_POINT gm\_size\_t gm\_strlen (const char \* cptr)

[gm\\_strlen](#)() calculates the length of a string.

**Return values:**

*gm\_size\_t* The length of the string.

**Parameters:**

*cptr* (IN) The string.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.87 gm\_strncasecmp.c File Reference

```
#include "gm_internal.h"
#include "gm_debug.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_strncasecmp](#) (const char \*a, const char \*b, int len)

### 8.87.1 Detailed Description

This file contains the GM API function [gm\\_strncasecmp](#)().

### 8.87.2 Function Documentation

#### 8.87.2.1 GM\_ENTRY\_POINT int gm\_strncasecmp (const char \* a, const char \* b, int len)

[gm\\_strncasecmp](#)() reimplements strncasecmp().

#### Return values:

*int* Returns an integer less than, equal to, or greater than zero if the first len bytes of a is found, respectively, to be less than, to match, or be greater than b.

#### Parameters:

*a* (IN) The first string to be compared.  
*b* (IN) The second string to be compared.  
*len* (IN) The number of bytes.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.88 gm\_strncmp.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT int [gm\\_strncmp](#) (const char \*a, const char \*b, int len)

### 8.88.1 Detailed Description

This file contains the GM API function [gm\\_strncmp](#)().

### 8.88.2 Function Documentation

#### 8.88.2.1 GM\_ENTRY\_POINT int gm\_strncmp (const char \* a, const char \* b, int len)

[gm\\_strncmp](#)() reimplements strcmp().

#### Return values:

*int* Returns an integer less than, equal to, or greater than zero if the first len bytes of a is found, respectively, to be less than, to match, or be greater than b.

#### Parameters:

- a* (IN) The first string to be compared.
- b* (IN) The second string to be compared.
- len* (IN) The length in bytes.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.89 gm\_strncpy.c File Reference

```
#include "gm_config.h"
#include "gm.h"
```

### Functions

- GM\_ENTRY\_POINT char \* [gm\\_strncpy](#) (char \*to, const char \*from, int len)

### 8.89.1 Detailed Description

This file contains the GM API function [gm\\_strncpy\(\)](#).

### 8.89.2 Function Documentation

#### 8.89.2.1 GM\_ENTRY\_POINT char\* [gm\\_strncpy](#) (char \*to, const char \*from, int len)

[gm\\_strncpy\(\)](#) copies exactly n bytes, truncating src or adding null characters to dst if necessary. The result will not be null-terminated if the length of src is n or more.

#### Return values:

*char* \* Returns a pointer to a destination string.

#### Parameters:

*to* (IN) The destination string to be copied.

*from* (IN) The source string to be copied.

*len* (IN) The number of bytes to be copied.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))



## 8.90 gm\_ticks.c File Reference

```
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT gm\_u64\_t [gm\\_ticks](#) (struct gm\_port \*port)

#### 8.90.1 Detailed Description

This file contains the GM API function [gm\\_ticks\(\)](#).

#### 8.90.2 Function Documentation

##### 8.90.2.1 GM\_ENTRY\_POINT gm\_u64\_t gm\_ticks (struct gm\_port \* port)

[gm\\_ticks\(\)](#) returns a 64-bit extended version of the LANai real time clock (RTC). For implementation reasons, the granularity of [gm\\_ticks\(\)](#) is 50 microseconds at the application level.

##### Return values:

*gm\_u64\_t*

##### Parameters:

*port* (IN) The handle to the GM port.

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.91 gm\_unique\_id.c File Reference

```
#include "gm_internal.h"
```

### Functions

- [gm\\_status\\_t gm\\_unique\\_id](#) (gm\_port\_t \*port, char \*id)

#### 8.91.1 Detailed Description

This file contains the GM API function [gm\\_unique\\_id\(\)](#).

#### 8.91.2 Function Documentation

##### 8.91.2.1 [gm\\_status\\_t gm\\_unique\\_id](#) (gm\_port\_t \* *port*, char \* *id*)

[gm\\_unique\\_id\(\)](#) returns the board id number for an interface.

##### Return values:

*GM\_SUCCESS* Operation completed successfully.

##### Parameters:

*port* (IN) The GM port.

*id* (OUT) Contains the LANai board id number.

##### Author:

Glenn Brown

##### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.92 gm\_unique\_id\_to\_node\_id.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_internal.h"
#include "gm_global_id.h"
```

### Functions

- [gm\\_status\\_t gm\\_unique\\_id\\_to\\_node\\_id](#) (gm\_port\_t \*port, char \*unique, unsigned int \*node\_id)

#### 8.92.1 Detailed Description

This file contains the GM API function [gm\\_unique\\_id\\_to\\_node\\_id\(\)](#).

#### 8.92.2 Function Documentation

##### 8.92.2.1 [gm\\_status\\_t gm\\_unique\\_id\\_to\\_node\\_id](#) (gm\_port\_t \*port, char \*unique, unsigned int \*node\_id)

[gm\\_unique\\_id\\_to\\_node\\_id\(\)](#) takes the MAC address and returns the GM node id for a specific port.

#### Return values:

*GM\_SUCCESS* Operation completed successfully.

*GM\_INVALID\_PARAMETER*

#### Parameters:

*port* (IN) The GM port.

*unique* (IN) The unique ID to translate.

*node\_id* (OUT) Where to store the node ID corresponding to unique.

#### Author:

Glenn Brown

**Version:**GM\_API\_VERSION (as defined in [gm.h](#))

## 8.93 gm\_unknown.c File Reference

```
#include "gm_call_trace.h"
#include "gm_debug.h"
#include "gm_debug_blocking.h"
#include "gm_enable_put.h"
#include "gm_internal.h"
```

### Functions

- GM\_ENTRY\_POINT void [gm\\_unknown](#) (gm\_port\_t \*p, gm\_rcv\_event\_t \*e)

#### 8.93.1 Detailed Description

This file contains the GM API function [gm\\_unknown\(\)](#).

#### 8.93.2 Function Documentation

##### 8.93.2.1 GM\_ENTRY\_POINT void [gm\\_unknown](#) (gm\_port\_t \* p, gm\_rcv\_event\_t \* e)

[gm\\_unknown\(\)](#) handles all GM events not recognized or processed by the client software, allowing the GM library and network interface card firmware to interact. This function also catches and reports several common client program errors, and converts some unrecognizable events into recognizable form for the client.

#### Parameters:

- p* (IN) The GM port associated with a specific event.
- e* (IN) The GM receive event.

#### Author:

Glenn Brown

#### Version:

GM\_API\_VERSION (as defined in [gm.h](#))

## 8.94 gm\_zone.c File Reference

```
#include "gm_call_trace.h"
#include "gm_internal.h"
#include "gm_malloc_debug.h"
```

### Data Structures

- struct [gm\\_zone\\_area](#)
- struct [gm\\_zone](#)

### Typedefs

- typedef [gm\\_zone\\_area](#) [gm\\_zone\\_area\\_t](#)
- typedef [gm\\_zone](#) [gm\\_zone\\_t](#)

### Functions

- GM\_ENTRY\_POINT struct [gm\\_zone](#) \* [gm\\_zone\\_create\\_zone](#) (void \*base, gm\_size\_t length)
- GM\_ENTRY\_POINT void [gm\\_zone\\_destroy\\_zone](#) (struct [gm\\_zone](#) \*zone)
- GM\_ENTRY\_POINT void \* [gm\\_zone\\_free](#) (struct [gm\\_zone](#) \*zone, void \*a)
- GM\_ENTRY\_POINT void \* [gm\\_zone\\_malloc](#) (struct [gm\\_zone](#) \*zone, gm\_size\_t length)
- GM\_ENTRY\_POINT void \* [gm\\_zone\\_calloc](#) (struct [gm\\_zone](#) \*zone, gm\_size\_t count, gm\_size\_t length)
- GM\_ENTRY\_POINT int [gm\\_zone\\_addr\\_in\\_zone](#) (struct [gm\\_zone](#) \*zone, void \*p)

#### 8.94.1 Detailed Description

This file contains the GM API functions [gm\\_zone\\_create\\_zone\(\)](#), [gm\\_zone\\_destroy\\_zone\(\)](#), [gm\\_zone\\_free\(\)](#), [gm\\_zone\\_malloc\(\)](#), and [gm\\_zone\\_calloc\(\)](#), [gm\\_zone\\_addr\\_in\\_zone\(\)](#).

This file provides alloc, calloc, and free routines to manage an externally specified "zone" of memory.

This package use buddy-system memory allocation to allocate ( $2^n$ )-byte regions of memory, where "n" is referred to as the "size" of the allocated area of memory.

All allocated (or freed) areas are maximally aligned. A zone is a chunk of memory starting and ending on page boundaries. The size and state of each area are encoded in a pair of bit-arrays. One array has a bit set for each position corresponding to a buddy-boundary. The second array has a bit set for each position corresponding to an area that is not free.

In the code, variables named with "size" are in logarithmic units and those named "length" are in real units.

#### DEBUGGING

If debugging is turned on, we "mark" all valid areas in the zone, using the gm\_mark API, and we check each area passed to a function. This should catch any DMA overruns type corruption of the data structures stored in the managed memory.

### 8.94.2 Define Documentation

#### 8.94.2.1 #define GM\_AREA\_FOR\_PTR(ptr)

Value:

```
((gm_zone_area_t *)
 ((char *) (ptr) - GM_ZONE_AREA_PTR_OFFSET))
```

### 8.94.3 Typedef Documentation

#### 8.94.3.1 typedef struct gm\_zone\_area gm\_zone\_area\_t

Zones are divided into managed buffers called "areas", which may either represent free buffers or buffers holding user data.

#### 8.94.3.2 typedef struct gm\_zone gm\_zone\_t

State of a zone, which is a region of memory from which one can allocate memory using the zone allocation functions.

### 8.94.4 Function Documentation

**8.94.4.1 GM\_ENTRY\_POINT struct [gm\\_zone](#)\* gm\_zone\_create\_zone (void \*  
*base*, gm\_size\_t *length*)**

[gm\\_zone\\_create\\_zone\(\)](#)

**Return values:**

[gm\\_zone](#) Handle to the GM zone.

**Parameters:**

*base*

*length*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.94.4.2 GM\_ENTRY\_POINT void gm\_zone\_destroy\_zone (struct [gm\\_zone](#) \*  
*zone*)**

[gm\\_zone\\_destroy\\_zone\(\)](#)

**Parameters:**

*zone* (IN) Pointer to the GM zone.

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.94.4.3 GM\_ENTRY\_POINT void\* gm\_zone\_free (struct [gm\\_zone](#) \* *zone*, void  
\* *a*)**

[gm\\_zone\\_free\(\)](#)

**Parameters:**

*zone* (IN) Pointer to the GM zone.



*a*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.94.4.4 GM\_ENTRY\_POINT void\* gm\_zone\_malloc (struct [gm\\_zone](#) \* zone, gm\_size\_t length)**

[gm\\_zone\\_malloc\(\)](#) mallocs a GM zone.

**Parameters:**

*zone* (IN) Pointer to the GM zone.

*length*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.94.4.5 GM\_ENTRY\_POINT void\* gm\_zone\_calloc (struct [gm\\_zone](#) \* zone, gm\_size\_t count, gm\_size\_t length)**

[gm\\_zone\\_calloc\(\)](#) callocs a GM zone.

**Parameters:**

*zone* (IN) Pointer to the GM zone.

*count*

*length*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

**8.94.4.6** `GM_ENTRY_POINT int gm_zone_addr_in_zone (struct gm\_zone *  
zone, void * p)`

`gm\_zone\_addr\_in\_zone()`

**Parameters:**

*zone* (**IN**) Pointer to the GM zone.

*p*

**Author:**

Glenn Brown

**Version:**

GM\_API\_VERSION (as defined in [gm.h](#))

---

## Chapter 9

# GM Page Documentation

### 9.1 XI. Alarms

GM provides the following simple alarm API. The alarm API allows the GM client to schedule a callback function to be called after a delay, specified in microseconds. An unbounded number of alarms may be set, although alarm overhead increases linearly in the number of set alarms, and the client must provide storage for each set alarm.

- [gm\\_initialize\\_alarm\(\)](#)
- [gm\\_cancel\\_alarm\(\)](#)
- [gm\\_set\\_alarm\(\)](#)
- [gm\\_flush\\_alarm\(\)](#)

GM clients will also be able to take advantage of the fact that an application is guaranteed to receive a single `GM_ALARM_EVENT` for each call to a client-supplied callback, with the corresponding callback occurring during the call to [gm\\_unknown\(\)](#) that processes that alarm. This means that a case statement like the following in the client's event loops can be used to significantly reduce the overhead of polling for any effect of a client supplied alarm callback:

```
case GM_ALARM_EVENT:
 gm_unknown (event);
 /* poll for effect of alarm callbacks only here */
 break;
```

---

## 9.2 VII. Page Allocation

The following GM API allows pages to be allocated and freed.

- [gm\\_page\\_alloc\(\)](#)
- [gm\\_page\\_free\(\)](#)
- [gm\\_alloc\\_pages\(\)](#)
- [gm\\_free\\_pages\(\)](#)

## 9.3 XV. GM Constants, Macros, and Enumerated Types

A number of constants, macros, and enumerated types are defined in the [gm.h](#) include file.

- Enum: GM\_HIGH\_PRIORITY The priority of high priority messages (1).
- Enum: GM\_LOW\_PRIORITY The priority of low priority messages (0).

/\* Maximum length of GM host name



1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## 9.5 II. About This Document

This document describes the GM message passing system and the GM-2.0 API. Several GM-1.0 API functions have been deprecated; however, the 1.0 API will continue to be supported by the GM libraries for the foreseeable future. New programs should use the GM API as described in this document.



## 9.6 X. Endian Conversion

GM receive events are delivered to the user in network byte order. This enhances the performance of GM programs, but is a minor inconvenience to developers using the GM API. The client must call a special function to convert each field read from the `gm_recv_event_t` union to host byte order. Neglecting this conversion will result in undefined program behavior in most cases.

In the absence of automatic checks, endian conversion is typically an error-prone programming task. Therefore, support has been added to GM-1.4 '[gm.h](#)' to ensure that no conversion is missing. Note, however, the support is incompatible with the deprecated [gm\\_send\(\)/GM\\_SENT\\_EVENT](#) mechanism in GM. All you need to do to activate the checking is add the line

```
#define GM_STRONG_TYPES 1
```

before the line

```
\#include "gm.h"
```

in your source code. Once the feature is activated, the compiler will report errors if any type conversion is missing. The error messages can be a bit cryptic and are platform specific, but they generally indicate some sort of type mismatch.

Endian conversion of fields in receive events from network to host order is achieved with the following functions:

Network to host conversion routines.

[gm\\_ntoh\\_u8\(\)](#) unsigned 8-bit

[gm\\_ntoh\\_u16\(\)](#) unsigned 16-bit

[gm\\_ntoh\\_u32\(\)](#) unsigned 32-bit

[gm\\_ntoh\\_u64\(\)](#) unsigned 64-bit

[gm\\_ntoh\\_s8\(\)](#) signed 8-bit

[gm\\_ntoh\\_s16\(\)](#) signed 16-bit

[gm\\_ntoh\\_s32\(\)](#) signed 32-bit

[gm\\_ntoh\\_s64\(\)](#) signed 64-bit

Host to network conversion routines.

[gm\\_hton\\_u8\(\)](#) unsigned 8-bit

[gm\\_hton\\_u16\(\)](#) unsigned 16-bit

[gm\\_hton\\_u32\(\)](#) unsigned 32-bit

[gm\\_hton\\_u64\(\)](#) unsigned 64-bit

[gm\\_hton\\_s8\(\)](#) signed 8-bit

[gm\\_hton\\_s16\(\)](#) signed 16-bit

[gm\\_hton\\_s32\(\)](#) signed 32-bit

[gm\\_hton\\_s64\(\)](#) signed 64-bit

(1) On 64-bit Solaris machines, the `GM_STRONG_TYPES` feature can be used during compilation to check for missing conversion, but if the resulting programs will not run and must be recompiled without this feature.

## 9.7 XIV. Example Programs

These GM example programs are intended primarily as illustrations and models of GM API usage. As such, they supplement the GM API documentation, and it is probably more useful to read their source than to execute them.

These programs illustrate the use of simple message sending ([gm\\_send\\_to\\_peer\\_with\\_callback\(\)](#)) and receiving ([gm\\_receive\(\)](#)), and also the use of [gm\\_directed\\_send\\_with\\_callback\(\)](#).

- [gm\\_simple\\_example.h](#)
- [gm\\_simple\\_example\\_send.c](#)
- [gm\\_simple\\_example\\_recv.c](#)

## 9.8 XII. High Availability Extensions

While GM automatically handles transient network errors such as dropped, corrupted, or misrouted packets, and while the GM mapper automatically reconfigures the network if links or nodes appear or disappear, GM cannot automatically handle catastrophic errors such as crashed hosts or loss of network connectivity without the cooperation of the client program.

When GM detects a catastrophic error, it temporarily disables the delivery of all messages with the same sender port, target port, and priority as the message that experienced the error, and GM informs the client of catastrophic network errors by passing a status other than `GM_SUCCESS` to the client's send completion callback routine. The client program is then expected to call either `gm_resume_sending()` or `gm_drop_sends()`, which re-enable the delivery of messages with the same sender port, target port, and priority. This mechanism preserves the message order over the prioritized connection between the sending and receiving ports, while allowing the client to decide if the other packets that it has already enqueued over the same connection should be transmitted or dropped.

Simpler GM programs, such as MPI programs, will typically consider GM send errors to be fatal and will typically exit when they see a send error. This is reasonable for applications running on small or physically robust clusters where errors are rare and when users can tolerate restarting jobs in the rare event of a network error. Poorly written GM programs may simply ignore the error codes, which will cause the program to eventually hang with no error indication when catastrophic errors are encountered. This poor programming practice is strongly discouraged: Developers should always check the send completion status. More sophisticated applications, such as high availability database applications, will respond to the network faults, which appear to the client as send completion status codes other than `GM_SUCCESS`.

A complete list of send completion status codes can be found in [gm.h](#) and section [VIII. Sending Messages](#).

When the send completion status code indicates an error a sophisticated client program may respond by calling `gm_resume_sending()` or `gm_drop_sends()`. Calling `gm_resume_sending()` causes GM to simply re-enable delivery of subsequent messages over the connection, including those that have already been enqueued. This would be the typical response of a distributed database that assumes the underlying network is unreliable and layers its own reliability protocol over GM. Calling `gm_drop_sends()` causes GM to drop all enqueued sends over the disabled connection, return them to the client with status `GM_SEND_DROPPED`, and re-enable the connection. This would be the typical response of a program that wishes to reorder subsequent communication over the connection in response to the error.

Note that each of the fault response functions (`gm_drop_sends()` and

`gm_resume_sending()` requires a send token. This send token is implicitly returned to the caller when the callback function passed to `gm_drop_sends()` or `gm_resume_sending()` is called by GM.

Here is an example program demonstrating the use of `gm_drop_sends()`. In this example, there are no messages queued after the message that has just been discarded, and `gm_drop_sends()` and `gm_resume_sending()` are equivalent. They just re-enable the target subport for further `gm_send_with_callback()` calls. If the send callback returns you an error, that means the corresponding message has been definitely discarded, both `gm_resume_sending()` and `gm_drop_sends()` only impact messages that have been queued **after** the message that has just been discarded.

```
#include <stdio.h>
#include <assert.h>

#include "gm.h"

unsigned int received,sent;
unsigned int my_gm_node_id;

static void test_send_callback (struct gm_port *port, void *context,
 gm_status_t status);

static void
drop_send_callback (struct gm_port *port,
 void *context,
 gm_status_t status)
{
 fprintf(stderr, "Got gm_drop_send notification, start resending\n");
 gm_send_with_callback(port, context, 20, 1, GM_LOW_PRIORITY,
 my_gm_node_id, 7, test_send_callback, context);
}

static void
test_send_callback (struct gm_port *port,
 void *context,
 gm_status_t status)
{
 switch (status)
 {
 {
 case GM_SUCCESS:
 fprintf(stderr, "Send successfully delivered\n");
 sent += 1;
 break;

 case GM_SEND_TIMED_OUT:
 fprintf(stderr, "Send timeout, provide buffers and initiate resend...\n");
 gm_provide_receive_buffer(port, context, 20, GM_LOW_PRIORITY);
 gm_provide_receive_buffer(port, context, 20, GM_LOW_PRIORITY);
 gm_drop_sends (port, GM_LOW_PRIORITY, my_gm_node_id, 7,
 drop_send_callback, context);

 break;
 }
 }
}
```

```
 case GM_SEND_DROPPED:
 fprintf(stderr, "Got DROPPED_SEND notification, resend\n");
 gm_send_with_callback(port, context, 20, 1, GM_LOW_PRIORITY,
 my_gm_node_id, 7, test_send_callback, context);

 break;

 default:
 fprintf(stderr, "Something bad happen\n");
 assert (0);
 }
}

int main (void)
{
 struct gm_port *port;
 gm_rcv_event_t *event;
 char *token;

 assert (gm_open (&port, 0, 7, "Test Resume", GM_API_VERSION) == GM_SUCCESS);
 token = gm_dma_malloc (port, sizeof (char));
 assert (token != NULL);

 received = 0;
 gm_get_node_id(port, &my_gm_node_id);
 gm_send_with_callback (port, token, 20, sizeof (char),
 GM_LOW_PRIORITY, my_gm_node_id, 7,
 test_send_callback, token);
 gm_send_with_callback (port, token, 20, sizeof (char),
 GM_LOW_PRIORITY, my_gm_node_id, 7,
 test_send_callback, token);

 while (received < 2 || sent < 2)
 {
 event = gm_receive (port);
 switch (gm_ntoh_u8 (event->rcv.type))
 {
 case GM_NO_RECV_EVENT:
 break;
 case GM_RECV_EVENT:
 fprintf(stderr, "received message\n");
 received += 1;
 break;
 default:
 gm_unknown (port, event);
 }
 }

 gm_dma_free (port, token);
 gm_close (port);
}
```

## 9.9 V. Initialization

Before calling any other GM function, `gm_init()` should be called. `gm_finalize()` should be called after all other GM calls and before your program exits. Each call to `gm_init()` should be balanced by a call to `gm_finalize()` before the program exits. Although GM automatically handles ungraceful program termination without such balanced calls on operating systems with memory protection, developers are strongly discouraged from relying on this feature because on some systems, such as those using the VxWorks embedded runtime system, the calls to `gm_finalize()` are required for proper shutdown of GM to allow ports to be reused without rebooting VxWorks.

A GM port is initialized by calling `gm_open()(struct gm_port**PORT, unsigned int UNIT, unsigned int PORT_ID, char *PORT_NAME, enum gm_api_version VERSION)` to open port number `PORT_ID` of Myrinet interface number `UNIT`. The pointer returned at `*PORT` must be passed to subsequent GM API calls. `PORT_NAME` is a character string of up to `gm_max_port_name_length()` bytes describing the client. The name is currently used for debugging purposes only, but this information will eventually be available to all GM clients on the network through a mechanism TBD. `VERSION` should be `'GM_API_VERSION_2_0'` as defined in `gm.h`.

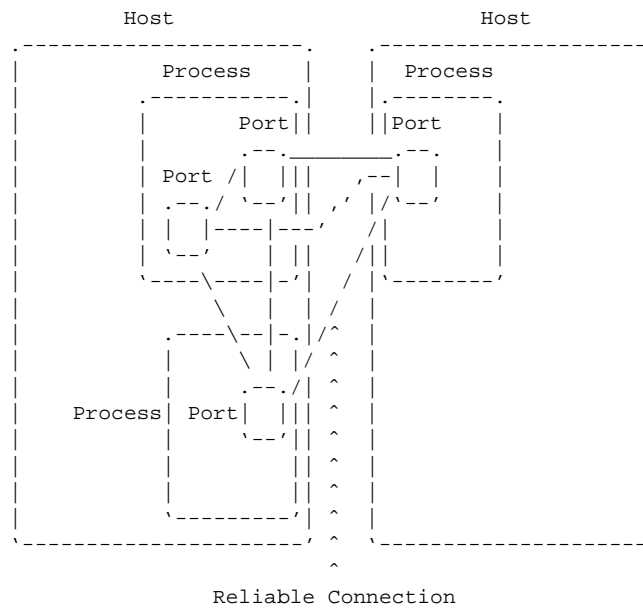
Note that while the GM API uses `'struct gm_port *'` pointers throughout, these pointers are opaque to the client. The client should not attempt to dereference these pointers.

After opening a port, the client implicitly possesses `gm_num_send_tokens()` send tokens and `gm_num_receive_tokens()` receive tokens. Most GM programs will use most or all of the `gm_num_receive_tokens()` immediately after opening a port to pass receive buffers to GM using `gm_provide_receive_buffer()`.

After the client has provided all receive buffers that it will provide during port initialization, the client should call `gm_set_acceptable_sizes()` for each priority (`GM_LOW_PRIORITY` and `GM_HIGH_PRIORITY`) to indicate what GM receive sizes the client expects to receive on the port. While this call is not strictly required, calling it allows GM to immediately reject any contradictory sends, immediately generating a send error at the sender. If these calls to `gm_set_acceptable_sizes()` are not made, then the error will not be reported until the sender experiences a GM long-period timeout, which takes about a minute to be generated by default. Therefore, calling `gm_set_acceptable_sizes()` can save much time during application development.

## 9.10 IV. Programming Model

The GM communication system provides reliable, ordered delivery between communication endpoints, called **ports**, with two levels of priority. This model is *connectionless* in that there is no need for client software to establish a connection with a remote port in order to communicate with it: the client software simply builds a message and sends it to any port in the network. (This apparently paradoxical *connectionless reliability* is achieved by GM maintaining reliable connections between each pair of hosts in the network and multiplexing the traffic between ports over these reliable connections.)



### 9.10.1 1. GM Endpoints (Ports)

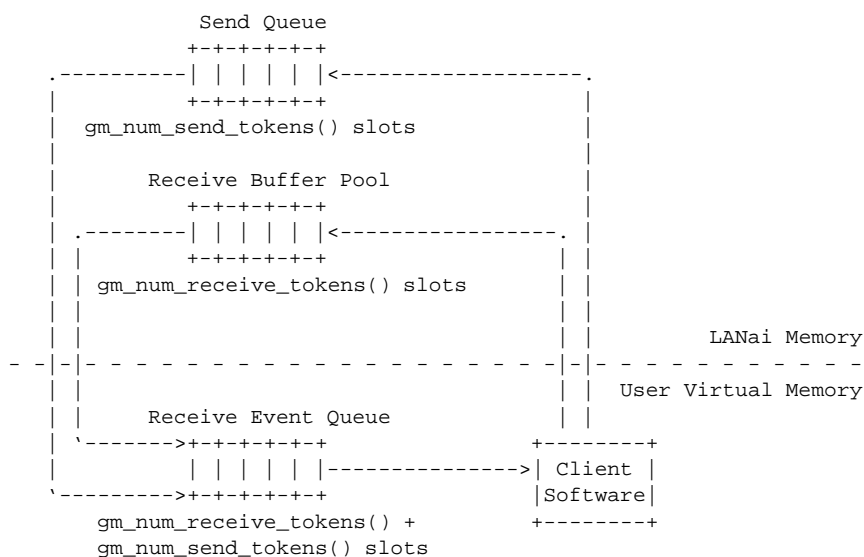
Under operating systems that provide memory protection, GM provides memory protected network access. It should be impossible for any non-privileged GM client application to use GM to access any memory other than the application's own memory, except as explicitly allowed by the GM API. The unforgeable source of each received message is available to the receiver, allowing the receiver to discard messages from untrusted sources.



The largest message GM can send or receive is limited to  $(2^{*31})-1$  bytes. However, because send and receive buffers must reside in DMAable memory, the maximum message size is limited by the amount of DMAable memory the GM driver is allowed to allocate by the operating system. Most GM applications obtain DMAable memory using the straightforward `gm_dma_malloc()` and `gm_dma_free()` calls, but sophisticated applications with large memory requirements may perform DMA memory management using `gm_register_memory()` and `gm_deregister_memory()` to pin and unpin memory on operating systems that support memory registration.

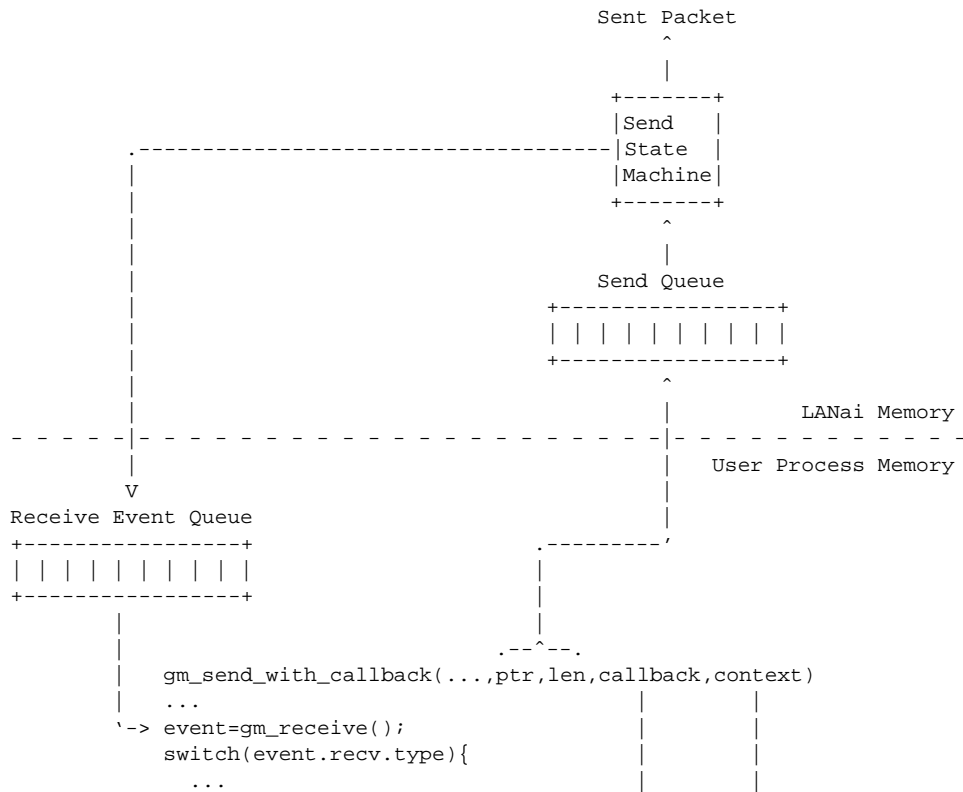
Message order is preserved only for messages of the same priority, from the same sending port, and directed to the same receiving port. Messages with differing priority never block each other. Consequently, low priority messages may pass high priority messages, unlike in some other communication systems. Typical GM applications will either use only one GM priority, or use the high priority channel for control messages (such as client-to-client acks) or for single-hop message forwarding.

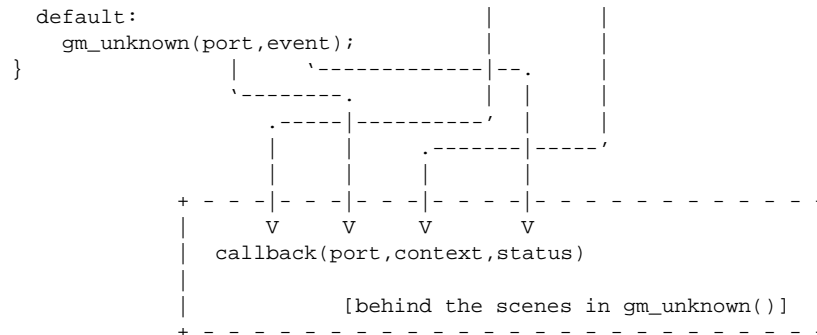
Both sends and receives in GM are regulated by implicit tokens, representing space allocated to the client in various internal GM queues, as depicted in the following figure. At initialization, the client implicitly possesses `gm_num_send_tokens()` send tokens, and `gm_num_receive_tokens()` receive tokens. The client may call certain functions only when possessing an implicit send or receive token, and in calling that function, the client implicitly relinquishes the token(1). The client program is responsible for tracking the number of tokens of each type that it possesses, and must not call any GM function requiring a token when the client does not possess the appropriate token. Calling a GM API function without the required tokens has undefined results, but GM usually reports such errors, and such errors will not cause system security to be violated.



### 9.10.2 2. User Token Flow (Sending)

As stated above, sends are token regulated. A client of a port may send a message only when it possesses a send token for that port. By calling a GM API send functions, the client implicitly relinquishes that send token. The client passes a `callback` and `context` pointer to the send function. When the send completes, GM calls `callback`, passing a pointer to the GM port, the client-supplied `context` pointer, and status code indicating if the send completed successfully or with an error. When GM calls the client's `callback` function, the send token is implicitly passed back to the client. Most GM programs, which rely on GM's fault tolerance to handle transient network faults, should consider a send completing with a status other than `GM_SUCCESS` to be a fatal error. However, more sophisticated programs may use the GM fault tolerance API extensions to handle such non-transient errors. These extensions are described in an appendix. It is important to note that the client-supplied `callback` function will be called only within a client's call to `gm_unknown()`, the GM unknown event handler function that the client must call when it receives an unrecognized event. The `gm_unknown()` function is described in more detail below.





The following functions require send tokens:

- [gm\\_datagram\\_send\(\)](#)
- [gm\\_directed\\_send\\_with\\_callback\(\)](#)
- [gm\\_drop\\_sends\(\)](#)
- [gm\\_resume\\_sending\(\)](#)
- [gm\\_send\\_to\\_peer\\_with\\_callback\(\)](#)
- [gm\\_send\\_with\\_callback\(\)](#)

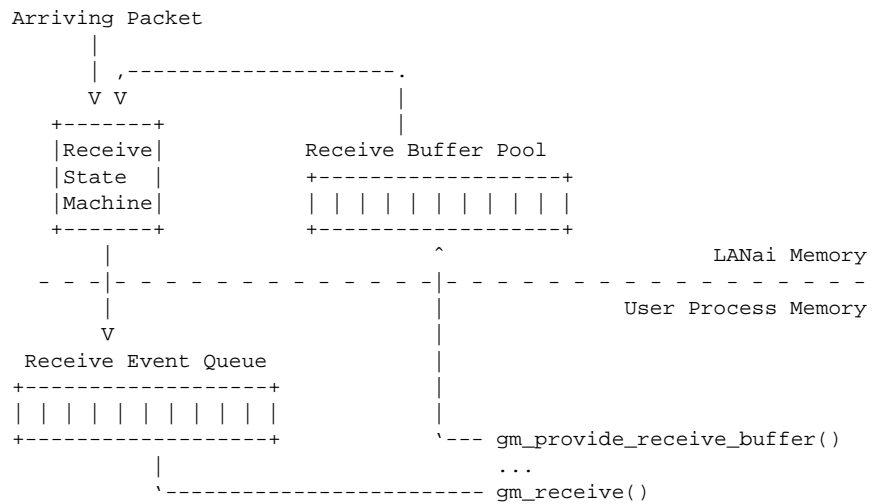
The send token is implicitly returned to the client when the function's callback is called or, for the GM-1.0 functions [gm\\_send\(\)](#) and [gm\\_send\\_to\\_peer\(\)](#), a send token is implicitly passed to the client with each pointer returned in a `GM_SENT_EVENT`. (The legacy `GM_SENT_EVENTS` are generated if and only if the legacy [gm\\_send\(\)](#) and [gm\\_send\\_to\\_peer\(\)](#) functions are called.)

### 9.10.3 3. User Token Flow (Receiving)

GM receives are also token regulated. After a port is opened, the client implicitly possesses [gm\\_num\\_receive\\_tokens\(\)](#) receive tokens, allowing it to provide GM with up to this many receive buffers using [gm\\_provide\\_receive\\_buffer\(\)](#). With each call to [gm\\_provide\\_receive\\_buffer\(\)](#), the client implicitly relinquishes a receive token. With each buffer passed to [gm\\_provide\\_receive\\_buffer\(\)](#), the client passes a corresponding integer `SIZE` indicating that the length of the receive buffer is at least [gm\\_max\\_length\\_for\\_size\(\)](#) bytes.

Before a client of a port can receive a message of a particular size and priority, the client software must provide GM with a receive token of matching size and priority. The receive token specifies the buffer in which to store the matching receive. When a message of matching size and priority is received, that message will be transferred into the receive buffer specified in the receive token. Note that multiple receive tokens of the same size and priority **may** be provided to the port.

After providing receive buffers with sizes matching the sizes of all packets that potentially could be received, the client must poll for receive events using a `gm_receive()`, `gm_blocking_receive()`, or `gm_blocking_receive_no_spin()` function. The `gm_receive()`, `gm_blocking_receive()`, or `gm_blocking_receive_no_spin()` function will return a `gm_receive_event`. The receipt of events of type `GM_RECV_EVENT` and `GM_HIGH_RECV_EVENT` describe received packets of low and high priority, respectively. All other events should be simply passed to `gm_unknown()`. Such events are used internally by GM for sundry purposes, and the client need not be concerned with the contents of unrecognized receive events unless otherwise stated in this document.



To avoid deadlock of the port, the client software must ensure that the port is never without a receive token for any acceptable combination of size and priority for more than a bounded amount of time, that the port is informed which combinations of size and priority are not acceptable for receives, and that the client not send to any remote port that does not do likewise.

By convention, when a port runs out of **low** priority receive tokens for any combination of sizes, the client may defer replacing the receive tokens pending the completion of a bounded number of **high** priority sends, but must always replace exhausted types of high priority receive tokens without waiting for any sends to complete. Using this technique, reliable, deadlock-free, single-hop forwarding can be achieved.

The following functions require a receive token:

- [gm\\_provide\\_receive\\_buffer\(\)](#)
- [gm\\_provide\\_receive\\_buffer\\_with\\_tag\(\)](#)

A single receive token is passed to the client with each of the following events:

- GM\_RAW\_RECV\_EVENT
- GM\_RECV\_EVENT
- GM\_HIGH\_RECV\_EVENT
- GM\_HIGH\_PEER\_RECV\_EVENT
- GM\_FAST\_HIGH\_RECV\_EVENT
- GM\_FAST\_HIGH\_PEER\_RECV\_EVENT

(However, if the client passes these events to [gm\\_unknown\(\)](#), then the token is implicitly returned to GM.) Any of the GM receive functions can generate these types of events. These functions are:

- [gm\\_receive\(\)](#)
- [gm\\_blocking\\_receive\(\)](#)
- [gm\\_blocking\\_receive\\_no\\_spin\(\)](#)

## 9.11 III. Overview

GM is a message-based communication system for Myrinet. Like many messaging systems, GM's design objectives included low CPU overhead, portability, low latency, and high bandwidth. Additionally, GM has several distinguishing characteristics:

- GM has extremely low overhead of about 1 microsecond per packet on all architectures.
- GM can provide simultaneous memory-protected user-level OS-bypass network interface access to several user-level applications simultaneously. (On systems that do not support memory protection, such as VxWorks, no memory protection is provided.)
- GM provides reliable ordered delivery between hosts in the presence of network faults. GM will detect and retransmit lost and corrupted packets. GM will also reroute packets around network faults when alternate routes exist. Catastrophic network errors, such as crashed hosts or disconnected links, are nonfatal; the undeliverable packets are returned to the client with an error indication, although most client programs are unable to adapt in the presence of such severe errors.
- GM supports clusters of over 10,000 nodes.
- GM provides two levels of message priority to allow efficient deadlock-free bounded-memory forwarding.
- GM allows clients to send messages up to  $2^{31} - 1$  bytes long, under operating systems that support sufficient amounts of DMAable memory to be allocated.
- GM automatically maps Myrinet networks.

GM is a light-weight communication layer, and as such has limitations that can be addressed by layering a heavier-weight interface over GM. Some such limitations are the following:

- GM is unable to send messages from or receive messages into nonDMAable memory.
- The GM API does not yet support any gather or scatter operations directly.

From the client's point of view, GM consists of a library, `libgm.a`, and a header file, `gm.h`. All externally visible GM identifiers in these files match the regular expression `^_*[Gg][Mm]_` to minimize name space pollution.

Additionally, GM has other parts that system administrators need to be concerned about:

- **gm** The GM driver provides systems services. It is called "gm" under Unix, and is the 'Myricom Myrinet Adapter' driver implemented in "gm.sys" under Windows NT.
- **mapper** The Myrinet mapper daemon maps the network. It is called "sbin/mapper" under Unix, and is the 'Myricom Myrinet Mapper Daemon' service implemented in "gm\_mapper\_service.exe" under Windows NT.

### 9.11.1 1. Definitions

This document attaches special meaning to a few commonly used words. The meaning of each of these words in the context of this document is defined here. In particular, please note the special meanings of the words `size` and `length`. Understanding the special meaning of these terms is critical to understanding this document.

- **aligned** A value is said to be aligned if it is a multiple of the required GM alignment. The required GM alignment is 1 on LANai7 hardware, 4 on LANai4 hardware, and 8 on LANai5 hardware. Pointers to memory allocated by GM are always automatically aligned.
- **client software**
- **client** The **client software** or simply **client** is the non-GM software that uses GM to provide a reliable ordered message delivery service. It can be an application, or a higher level networking layer, such as MPI or VI.
- **message** A **message** is an aligned array of bytes in DMAable memory.
- **buffer** A **buffer** is a contiguous region of DMAable memory into which a message may be copied. All GM buffers must be aligned.
- **length** The **length** of a message is the number of bytes of data that comprise the message. There is no alignment restriction on the length of any GM message.

The **length** of a receive buffer is the number of bytes that may be safely copied into the buffer.

- **packet** A **packet** is an aggregation of bytes sent over the network. Packet lengths are limited to just over `gm_mtu()(port)` (usually 4096 bytes) to bound the time any packet can monopolize network resource. Note that multiple packets are required to send large messages over the network, but the segmentation of messages into packets and reassembly of packets into messages is performed automatically by GM.

- **size** The **size** of the message is any integer greater than or equal to

$$\frac{\log_2 (\text{LENGTH} + 8)}{2}$$

where LENGTH is the length of the message.

The **size** of a receive buffer is any positive integer **less** than or equal to

$$\frac{\log_2 (\text{LENGTH} + 8)}{2}$$

where LENGTH is the length of the buffer. Consequently, a buffer of size SIZE must have a LENGTH of at least  $2^{2 \cdot \text{SIZE}} - 8$ .

A buffer having a longer length serves no useful purpose in GM, but is allowed.

The function `gm_min_size_for_length()(LENGTH)` can be used to compute the minimum size for any length, and the function `gm_max_length_for_size()(SIZE)` can be used to compute the maximum length for any size.

- **port** A **port** is a GM communication endpoint, and serves as the interface between the client software and the network.
- **user** A human using an application that uses GM.
- **user virtual memory** Memory directly accessible by software running in a user application.
- **kernel virtual memory** Memory directly accessible by the GM driver.



### 9.11.2 2. Notation

The following terms abbreviations are used in the GM documentation and source code. Some of these abbreviations are obvious to speakers of English, but are included for speakers of other languages. This section does not include architecture-specific abbreviations used in the architecture-specific GM driver code, as those are documented by the architecture's vendor and are not of interest to most GM developers.

- **accum** accumulator
- **addr** address
- **alloc** allocate
- **arch(s)** architecture(s)
- **buf**
- **buff** buffer
- **cnt** count
- **create** allocate and then initialize
- **destroy** finalize and then free
- **dma** direct memory access
- **hash** hash table
- **hp** host pointer (a pointer of the appropriate size for the host architecture in question)
- **insn(s)** instruction(s)
- **intr** interrupt
- **KVMA** kernel virtual memory address
- **lookaside** lookaside list
- **LSB(s)** least significant byte(s)

- **lsb**(’s) least significant bit(s)
- **MAC** Media Access Control. This is a commonly referred to sublayer of the datalink layer in the ISO network reference model.
- **MAC Address** A 6-byte address unique to a Myrinet interface. It is equivalent to an ethernet address.
- **minor** device minor number
- **num** number
- **phys** physical
- **pre** prefetch or precompute
- **PTE** page table entry
- **recv** receive
- **ref** reference
- **seg** segment
- **sema** semaphore
- **UVMA** user virtual memory address
- **virt** virtual
- **VM** virtual memory
- **VMA** virtual memory address
- **zalloc** allocate and clear

## 9.12 IX. Receiving Messages

Similarly to message sends, message receives in GM are regulated by a simple token-passing mechanism: Before a message can be received, the client software must provide GM a receive token that allows the message to be received and specifies a buffer to hold the received data.

After initialization, the client implicitly possesses all `gm_num_receive_tokens()` receive tokens. The client software grants receive tokens to GM by calling `gm_provide_receive_buffer()(PORT,BUFFER,SIZE,PRIORITY)`, indicating that GM may receive any message into `BUFFER` as long as the `size` and `priority` fields of the received message exactly match the `SIZE` and `PRIORITY` fields passed to `gm_provide_receive_buffer()`. Eventually, GM will use the buffer indicated by `MESSAGE` and `SIZE` to receive a message of the indicated `SIZE` and `PRIORITY`. Unlike some messaging systems, GM requires that the `SIZE` of the received message match the token size exactly. GM will **not** use the next larger sized receive buffer when a receive buffer of the correct size is not available. All receive buffers passed to `gm_provide_receive_buffer()` must DMAable. They must also be aligned or be within memory allocated using `gm_dma_calloc()` or `gm_dma_malloc()` to ensure that messages can be DMAed into the buffer, and must be at least `gm_max_length_for_size()(SIZE)` bytes long.

Typical GM clients will provide at least 2 receive buffers for each size and priority of message that might be received to maximize performance by allowing one buffer to be processed and replaced while the network is filling the other. However, 1 receive buffer for each size-priority combination is sufficient for correct operation. Additionally, it is almost always a good idea to provide additional buffers for the smallest sizes, so that many small messages may be received while the host is busy computing. There is no need to provide tokens for receives smaller than `gm_min_message_size()`.

After providing receive tokens, code may poll for pending events using `gm_receive_pending()(port)`, which returns a nonzero value if a receive is pending or zero if no event has been received. `gm_next_event_peek()(struct gm_port *P, gm_u16_t *SENDER)` can also be used to peek at the event at the head of the queue. The return value is the event type (zero if no event is pending). The `SENDER` parameter will be filled with the sender of the message if the event is a message receive event. The client may also poll for receives using `gm_receive()(PORT)`, which returns a pointer to an event structure of type `gm_event_t`. If no `recv` event is in the receive queue, a pointer to a fake receive event of `GM_NO_RECV_EVENT` will be returned. The event returned by `gm_receive()` is only guaranteed to be valid until the next call to `gm_receive()`.

There are several variants of `gm_receive()` available, all of which can safely be used in the same program.

`gm_receive()` returns the first pending receive event or `GM_NO_RECV_EVENT` if none is

pending.

[gm\\_blocking\\_receive\(\)](#) returns the first pending receive, blocking if necessary. This function polls for receives for 1 millisecond before sleeping, so it should generally be used only if the polling thread has a dedicated processor.

[gm\\_blocking\\_receive\\_no\\_spin\(\)](#) returns the first pending receive, blocking if necessary. This function sleeps immediately if no receive is pending. It should be generally used in environments with more than one thread per processor.

Once the client has obtained a receive event from one of these three functions, the client should either process the event if the client recognizes the event, or pass the event to [gm\\_unknown\(\)](#) if the event is unrecognized. All fields in the receive event are in network byte order, and must be converted to host byte order as specified in section [X. Endian Conversion](#).

The client is not required to handle any receive events, and may simply pass all events to [gm\\_unknown\(\)](#), but any useful GM program will handle `GM_RECV_EVENT` or `GM_HIGH_RECV_EVENT` in order to access the received data. The receive event types that the client software may choose to recognize are as follows (GM internal events are not listed):

- **GM\_NO\_RECV\_EVENT** No event is in the event queue.
- **GM\_ALARM\_EVENT** `GM_ALARM_EVENT` should be treated as an unknown event and passed to [gm\\_unknown\(\)](#). However, because client alarm handlers are called within [gm\\_unknown\(\)](#) when [gm\\_unknown\(\)](#) receives a `GM_ALARM_EVENT`, it can be useful for a program to perform alarm polling only after passing `GM_ALARM_EVENT` to [gm\\_unknown\(\)](#), as in the `test/gm_allsize.c` example program. See the documentation for [gm\\_set\\_alarm\(\)](#) for more information.
- **GM\_SENT\_EVENT** This type indicates that one or more sends completed. \*Developers using the GM-1.5 API should never see this event type\*, as it is generated only if the client calls the GM-1.0 [gm\\_send\(\)](#) function, which is deprecated in favor of the superior [gm\\_send\\_with\\_callback\(\)](#) functions.

`event->sent.message_list` points to a null-terminated array of 'void' pointers, which are message pointers from earlier [gm\\_send\(\)](#) calls that have completed successfully. For each pointer in this array, a send token is implicitly returned to the client.

- **\_GM\_SLEEP\_EVENT**
- **GM\_RAW\_RECV\_EVENT** This type is for internal use by the GM mapper process and will never be received by normal GM clients. It provides the following information in the `event → recv` structure:

- `length` the number of bytes received
- `buffer` the location of the received bytes

- **GM\_RAW\_RECV\_EVENT**
- **GM\_BAD\_SEND\_DETECTED\_EVENT**
- **GM\_SEND\_TOKEN\_VIOLATION\_EVENT**
- **GM\_RECV\_TOKEN\_VIOLATION\_EVENT**
- **GM\_BAD\_RECV\_TOKEN\_EVENT**
- **GM\_ALARM\_VIOLATION\_EVENT**
- **GM\_RECV\_EVENT**
- **GM\_HIGH\_RECV\_EVENT** This event indicates that a normal receive has occurred. The following information is available in the `event` → `recv` structure.
  - `length` the number of bytes of received data
  - `size` the size of the buffer into which the message was received
  - `buffer` a pointer to the buffer passed in a call to [gm\\_provide\\_receive\\_buffer\(\)](#), which allowed this receive to occur
  - `sender_node_id` the GM identifier for the node that sent the message
  - `sender_port_id` the GM identifier for the port that sent the message
  - `tag` the tag passed to [gm\\_provide\\_receive\\_buffer\\_with\\_tag\(\)](#) or 0 if [gm\\_provide\\_receive\\_buffer\(\)](#) was used instead
  - `type` `GM_HIGH_RECV_EVENT` indicates the receipt of a high-priority packet. `GM_RECV_EVENT` indicates the receipt of a low-priority packet.
- **GM\_PEER\_RECV\_EVENT**
- **GM\_HIGH\_PEER\_RECV\_EVENT** These events may be safely ignored (passed to [gm\\_unknown\(\)](#)), in which case the event will be converted to a normal `GM_RECV_EVENT` and passed to the client in the next call to a `gm_*receive*()` function.

These events are just like the normal `GM_RECV_EVENT` and `GM_HIGH_RECV_EVENT` events, but indicate that the sender port id is the same as the receiver port id. Most GM programs should handle these events directly just like they handle normal receive events.

- `length` the number of bytes of received data
- `size` the size of the buffer into which the message was received

- `buffer` a pointer to the buffer passed in a call to `gm_provide_receive_buffer()`, which allowed this receive to occur
  - `sender_node_id` the GM identifier for the node that sent the message
  - `sender_port_id` the GM identifier for the port that sent the message
  - `tag` the tag passed to `gm_provide_receive_buffer_with_tag()` or 0 if `gm_provide_receive_buffer()` was used instead.
  - `type` The PEER event types indicate that the sender port number is the same as the port number. The HIGH event types indicate that the message was sent with high priority.
- **GM\_FAST\_RECV\_EVENT**
  - **GM\_FAST\_HIGH\_RECV\_EVENT**
  - **GM\_FAST\_PEER\_RECV\_EVENT**
  - **GM\_FAST\_HIGH\_PEER\_RECV\_EVENT** These events may be safely ignored (passed to `gm_unknown()`), in which case the event will be converted to a normal `GM_RECV_EVENT` and passed to the client in the next call to a `gm_*receive*`() function. The conversion process will copy the receive message from the receive queue into the receive buffer.

These types indicate that a small-message receive occurred with the small message stored in the receive queue for improved small-message performance. The **PEER** event types indicate that the sender port number is the same as the port number. The **HIGH** event types indicate that the message was sent with high priority.

If your program uses any small messages that are immediately processed and discarded upon receipt, then your program can improve performance by processing these messages directly. If after examining the message your program determines that it needs the data copied into the buffer, it can either call `gm_memorize_message()` to do so or can pass the event to `gm_unknown()`.

- `message` a pointer to the received message, which is stored in the receive queue and is only guaranteed to be valid until the next call to `gm_receive()`.
- `length` the number of bytes of received data
- `size` the size of the buffer into which the message was received
- `buffer` a pointer to the buffer passed in a call to `gm_provide_receive_buffer()`, which allowed this receive to occur
- `sender_node_id` the GM identifier for the node that sent the message
- `sender_port_id` the GM identifier for the port that sent the message
- `tag` the tag passed to `gm_provide_receive_buffer_with_tag()` or 0 if `gm_provide_receive_buffer()` was used instead.
- `type` The PEER types indicate that the sender port number is the same as the port number. The HIGH types indicate that the message was sent with high priority.

Note that although the receive data is in the receive queue and no receive buffer was used to store the received message, the client must have provided an appropriate receive buffer before the receive could take place, and this buffer is passed back to the client in the fast receive event. If the client needs to store the data `*message` past the next call to `gm_receive()`, then the client should copy `*message` into `*buffer` using `gm_memorize_message()`, which is simply a version of `bcopy()` optimized for copying aligned messages. After calling `gm_memorize_message()`, the fast receive event becomes equivalent to a normal receive event.

Although the number of receive events may seem daunting at first glance, almost all of the event types can be ignored. The following receive dispatch loop is fully functional for a nontrivial application that accepts messages ports, accepts only small control messages sent with high priority, and accepts low priority messages of any size:

```
{
 struct gm_port *my_port;
 gm_rcv_event_t *e;
 void *some_buffer;
 ...
 while (1) {
 e = gm_receive (my_port);
 switch (gm_htohc (e->rcv.type))
 {
 case GM_HIGH_RECV_EVENT:
 /* Handle high-priority control messages here in bounded time */
 gm_provide_rcv_buffer (my_port,
 gm_ntohp (e->rcv.buffer),
 gm_ntohc (e->rcv.size),
 GM_HIGH_PRIORITY);

 break;

 case GM_RECV_EVENT:
 /* Handle data messages here in bounded time */
 gm_provide_rcv_buffer (my_port, some_buffer,
 gm_ntohc (e->rcv.size),
 GM_LOW_PRIORITY);

 break;

 case GM_NO_RECV_EVENT:
 /* Do bounded-time processing here, if desired. */
 break;

 default:
 gm_unknown (my_port, e);
 }
 }
}
```

However, the following implementation is slightly faster because it handles control messages without copying them into the receive buffer:

```
{
```

```

struct gm_port *my_port;
gm_recv_event_t *e;
void *some_buffer;
...
while (1) {
 e = gm_receive (my_port);
 switch (gm_ntohc (e->recv.type))
 {
 case GM_FAST_HIGH_PEER_RECV_EVENT:
 case GM_FAST_HIGH_RECV_EVENT:
 /* Handle high-priority control messages here in bounded time */
 gm_provide_recv_buffer (my_port,
 gm_ntohp (e->recv.buffer),
 gm_ntohc (e->recv.size),
 GM_HIGH_PRIORITY);

 break;

 case GM_FAST_PEER_RECV_EVENT:
 case GM_FAST_RECV_EVENT:
 gm_memorize_message (gm_ntohp (e->recv.message),
 gm_ntohp (e->recv.buffer),
 gm_ntohl (e->recv.length));

 case GM_PEER_RECV_EVENT:
 /* Handle data messages here in bounded time */
 gm_provide_recv_buffer (my_port, some_buffer,
 gm_ntohc (e->recv.size),
 GM_LOW_PRIORITY);

 break;

 case GM_NO_RECV_EVENT:
 /* Do bounded-time processing here, if desired. */
 break;

 default:
 gm_unknown (my_port, e);
 }
}

```

Any receive event not recognized by an application must be passed immediately to [gm\\_unknown\(\)](#), as in the example above. The function [gm\\_unknown\(\)](#) will free any resources associated with the event that the client application would normally be expected to free if it recognized the type. Also, additional, undocumented event types will be received by an application and are handled by [gm\\_unknown\(\)](#). These messages can be used for supporting features such as GM alarms and blocking receives.

The motivation for putting small messages in the receive queue despite the fact that doing so might require a receive-side copy is the following set of observations:

A large fraction of small receive messages are control messages that can be processed immediately upon reception, and consequently do not need to be copied into the more permanent buffer to survive calls to [gm\\_receive\(\)](#).

The cost of performing an additional DMA to place the message in the buffer, rather



than in the receive queue, is actually more expensive for very small messages than having the host perform the copy.

Therefore, placing small received messages in the receive command queue rather than in the more permanent receive buffer enhances performance and is worth the added complexity.

To prevent program deadlock, the client software must ensure that GM is never without a receive token (buffer) for any potentially received message for more than a bounded amount of time. Generally, except for the case of message **forwarding** described in the next chapter, this means that after each successful call to `gm_receive()` the client will call `gm_provide_receive_buffer()` to replace the receive token (buffer) with one of the same SIZE and PRIORITY before the next call to `gm_receive()` or `gm_send()`. If such a deadlock condition exists for too long (on the order of a minute) or too often (a significant fraction of a one-minute interval), then remote sends directed at the receiving port will time out.

## 9.13 VIII. Sending Messages

In GM, message sends are regulated by a simple token-passing mechanism to prevent GM's bounded-size internal queues from overflowing. The client software must possess a send token before calling `gm_send_with_callback()`. After initialization, the client software implicitly possesses all `gm_num_send_tokens()` send tokens, and implicitly passes one token to the GM library with each call to `gm_send_with_callback()` or `gm_send_to_peer_with_callback()`. The token is retained by GM until the send completes, at which time GM calls the client-supplied callback, implicitly returning the send token to the client. The contents of the send message should not be modified in the interval between the call to `gm_send_with_callback()` and the send completion, because doing so will cause undefined data to be delivered to the receiver.

The order of messages with different priorities or with different destination ports is not preserved. Only the order of messages with the same priority and to the same destination port is preserved.

In the special case that the `TARGET_PORT_ID` is the same as the sending port ID (as is often the case), the streamlined `gm_send_to_peer_with_callback()` function may be used instead of `gm_send_with_callback()`, allowing the `TARGET_PORT_ID` parameter to be omitted, and slightly improving small-message performance on 32-bit Myrinet interfaces.

The send completion status codes (listed in `gm.h`) are as follows:

- **GM\_SUCCESS** The send succeeded. This status code does not indicate an error.
- **GM\_FAILURE** Operation failed.
- **GM\_INPUT\_BUFFER\_TOO\_SMALL** Input buffer is too small.
- **GM\_OUTPUT\_BUFFER\_TOO\_SMALL** Output buffer is too small.
- **GM\_TRY\_AGAIN** Try again.
- **GM\_BUSY** GM Port is Busy.
- **GM\_MEMORY\_FAULT** Memory Fault.
- **GM\_INTERRUPTED** Interrupted.
- **GM\_INVALID\_PARAMETER** Invalid input parameter.

- **GM\_OUT\_OF\_MEMORY** Out of Memory.
- **GM\_INVALID\_COMMAND** Invalid Command.
- **GM\_PERMISSION\_DENIED** Permission Denied.
- **GM\_INTERNAL\_ERROR** Internal Error.
- **GM\_UNATTACHED** Unattached.
- **GM\_UNSUPPORTED\_DEVICE** Unsupported Device.
  
- **GM\_SEND\_TIMED\_OUT** The target port is open and responsive and the message is of an acceptable size, but the receiver failed to provide a matching receive buffer within the timeout period. This error can be caused by the receiver neglecting its responsibility to provide receive buffers in a timely fashion or crashing. It can also be caused by severe congestion at the receiving node where many senders are contending for the same receive buffers on the target port for an extended period. This error indicates a programming error in the client software.
- **GM\_SEND\_REJECTED** The receiver indicated (in a call to [gm\\_set\\_acceptable\\_sizes\(\)](#)) the size of the message was unacceptable. This error indicates a programming error in the client software.
- **GM\_SEND\_TARGET\_PORT\_CLOSED** The message cannot be delivered because the destination port has been closed.
- **GM\_SEND\_TARGET\_NODE\_UNREACHABLE** The target node could not be reached over the Myrinet. This error can be caused by the network becoming disconnected for too long, the remote node being powered off, or by network links being rearranged when the Myrinet mapper is not running.
- **GM\_SEND\_DROPPED** The send was dropped at the client's request. (The client called [gm\\_drop\\_sends\(\)](#).) This status code does not indicate an error.
- **GM\_SEND\_PORT\_CLOSED** Clients should never see this internal error code.
- **GM\_NODE\_ID\_NOT\_YET\_SET** Node ID is not yet set.
- **GM\_STILL\_SHUTTING\_DOWN** GM Port is still shutting down.
- **GM\_CLONE\_BUSY** GM Clone Busy.

- **GM\_NO\_SUCH\_DEVICE** No such device.
- **GM\_ABORTED** Aborted.
- **GM\_INCOMPATIBLE\_LIB\_AND\_DRIVER** Incompatible GM library and driver.
- **GM\_UNTRANSLATED\_SYSTEM\_ERROR** Untranslated System Error.
- **GM\_ACCESS\_DENIED** Access Denied.
- **GM\_NO\_DRIVER\_SUPPORT** No Driver Support.
- **GM\_PTE\_REF\_CNT\_OVERFLOW** PTE Ref Cnt Overflow.
- **GM\_NOT\_SUPPORTED\_IN\_KERNEL** Not supported in the kernel.
- **GM\_NOT\_SUPPORTED\_ON\_ARCH** Not supported for this architecture.
- **GM\_NO\_MATCH** No match.
- **GM\_USER\_ERROR** User error.
- **GM\_TIMED\_OUT** Timed out.
- **GM\_DATA\_CORRUPTED** Data has been corrupted.
- **GM\_HARDWARE\_FAULT** Hardware fault.
- **GM\_SEND\_ORPHANED** Send orphaned.
- **GM\_MINOR\_OVERFLOW** Minor overflow.
- **GM\_PAGE\_TABLE\_FULL** Page Table is Full.
- **GM\_UC\_ERROR** UC Error.
- **GM\_INVALID\_PORT\_NUMBER** Invalid Port Number.

## 9.14 VI. Memory Setup

GM will only send messages from memory allocated with a [gm\\_dma\\_alloc\(\)](#) or [gm\\_dma\\_malloc\(\)](#) function, or memory that has been registered for DMA transfers using [gm\\_register\\_memory\(\)](#). If the client attempts to send data from nonDMAable memory, GM will send bytes of value 0xaa instead. If the client attempts to receive data into nonDMAable memory, the data will be silently discarded, and a BOGUS send or receive will appear in the kernel log.

Note that some operating systems (e.g.: Solaris) do not support [gm\\_register\\_memory\(\)](#) due to operating system limitations, so the [gm\\_dma\\_alloc\(\)](#) or [gm\\_dma\\_malloc\(\)](#) functions must be used instead to obtain DMA memory.

Unless explicitly enabled using [gm\\_allow\\_remote\\_memory\\_access\(\)\(PORT\)](#), GM will not allow remote processes to use [gm\\_directed\\_send\\_with\\_callback\(\)](#) ([gm\\_put\(\)](#)) to modify the memory of the process. If remote memory access has been enabled, then this protection is disabled, and **any** remote GM port may modify the contents of **any** DMAable memory associated with that port. GM developers should be aware of this potential security risk, although it is usually not a concern.

## 9.15 XVI. Function Summary

We have subdivided this GM function summary into a listing of "Basic GM API" functions and "Advanced GM API" functions.

The following "basic" functions should suffice for most GM API applications.

### Basic GM API

#### Initialization

- [gm\\_init\(\)](#)
- [gm\\_finalize\(\)](#)
- [gm\\_open\(\)](#)
- [gm\\_close\(\)](#)
- [gm\\_exit\(\)](#)
- [gm\\_abort\(\)](#)

Currently, [gm\\_open\(\)](#) implicitly calls [gm\\_init\(\)](#) for the caller and [gm\\_close\(\)](#) implicitly calls [gm\\_finalize\(\)](#), but developers should not rely on this.

#### Memory Setup

- [gm\\_dma\\_calloc\(\)](#)
- [gm\\_dma\\_malloc\(\)](#)
- [gm\\_dma\\_free\(\)](#)

#### Sending Messages

- [gm\\_send\\_to\\_peer\\_with\\_callback\(\)](#)
- [gm\\_send\\_with\\_callback\(\)](#)

#### Receiving Messages

- [gm\\_provide\\_receive\\_buffer\\_with\\_tag\(\)](#)
- [gm\\_receive\(\)](#)
- [gm\\_unknown\(\)](#)

#### Endian Conversion

- [gm\\_ntoh\\_u8\(\)](#)
- [gm\\_ntoh\\_u16\(\)](#)
- [gm\\_ntoh\\_u32\(\)](#)
- [gm\\_ntoh\\_u64\(\)](#)
- [gm\\_ntoh\\_s8\(\)](#)
- [gm\\_ntoh\\_s16\(\)](#)
- [gm\\_ntoh\\_s32\(\)](#)
- [gm\\_ntoh\\_s64\(\)](#)

#### Error Handling

- [gm\\_perror\(\)](#)
- [gm\\_eprintf\(\)](#)
- [gm\\_printf\(\)](#)

#### Utility Functions

- [gm\\_max\\_length\\_for\\_size\(\)](#)
- [gm\\_min\\_message\\_size\(\)](#)
- [gm\\_min\\_size\\_for\\_length\(\)](#)
- [gm\\_set\\_acceptable\\_sizes\(\)](#)

### Advanced GM API

The following "advanced" functions supplement the Basic GM API functions; they may be useful for more complex applications.

#### Memory Setup

- [gm\\_allow\\_remote\\_memory\\_access\(\)](#)
- [gm\\_register\\_memory\(\)](#)
- [gm\\_deregister\\_memory\(\)](#)
- [gm\\_alloc\\_pages\(\)](#)
- [gm\\_free\\_pages\(\)](#)
- [gm\\_page\\_alloc\(\)](#)
- [gm\\_page\\_free\(\)](#)

#### Sending Messages

- [gm\\_alloc\\_send\\_token\(\)](#)
- [gm\\_datagram\\_send\(\)](#)
- [gm\\_datagram\\_send\\_4\(\)](#)
- [gm\\_directcopy\\_get\(\)](#)
- [gm\\_directed\\_send\\_with\\_callback\(\)](#)
- [gm\\_put\(\)](#)
- [gm\\_free\\_send\\_token\(\)](#)
- [gm\\_free\\_send\\_tokens\(\)](#)
- [gm\\_num\\_send\\_tokens\(\)](#)
- [gm\\_send\\_token\\_available\(\)](#)



### Receiving Messages

- [gm\\_blocking\\_receive\(\)](#)
- [gm\\_blocking\\_receive\\_no\\_spin\(\)](#)
- [gm\\_get\(\)](#)
- [gm\\_memorize\\_message\(\)](#)
- [gm\\_next\\_event\\_peek\(\)](#)
- [gm\\_num\\_receive\\_tokens\(\)](#)
- [gm\\_receive\\_pending\(\)](#)

### Endian Conversion

- [gm\\_hton\\_u8\(\)](#)
- [gm\\_hton\\_u16\(\)](#)
- [gm\\_hton\\_u32\(\)](#)
- [gm\\_hton\\_u64\(\)](#)
- [gm\\_hton\\_s8\(\)](#)
- [gm\\_hton\\_s16\(\)](#)
- [gm\\_hton\\_s32\(\)](#)
- [gm\\_hton\\_s64\(\)](#)
- [gm\\_htopci\\_u32\(\)](#)

### Alarms

- [gm\\_cancel\\_alarm\(\)](#)
- [gm\\_flush\\_alarm\(\)](#)

- [gm\\_initialize\\_alarm\(\)](#)
- [gm\\_set\\_alarm\(\)](#)

#### High Availability Extensions

- [gm\\_drop\\_sends\(\)](#)
- [gm\\_resume\\_sending\(\)](#)

#### Utility Functions

- [gm\\_get\\_host\\_name\(\)](#)
- [gm\\_get\\_node\\_id\(\)](#)
- [gm\\_get\\_node\\_type\(\)](#)
- [gm\\_get\\_unique\\_board\\_id\(\)](#)
- [gm\\_get\\_mapper\\_unique\\_id\(\)](#)
- [gm\\_getpid\(\)](#)
- [\\_gm\\_get\\_route\(\)](#)
- [gm\\_get\\_port\\_id\(\)](#)
- [gm\\_host\\_name\\_to\\_node\\_id\\_ex\(\)](#)
- [gm\\_max\\_node\\_id\(\)](#)
- [gm\\_max\\_node\\_id\\_in\\_use\(\)](#)
- [gm\\_mtu\(\)](#)
- [gm\\_node\\_id\\_to\\_host\\_name\\_ex\(\)](#)
- [gm\\_node\\_id\\_to\\_unique\\_id\(\)](#)
- [gm\\_num\\_ports\(\)](#)

- [gm\\_set\\_enable\\_nack\\_down\(\)](#)
- [gm\\_unique\\_id\(\)](#)
- [gm\\_unique\\_id\\_to\\_node\\_id\(\)](#)

### Miscellaneous Routines

The following miscellaneous library functions are provided. Several are simply cover functions for standard Unix library functions, but are provided to simplify the creation of portable GM programs, or to provide the ANSI functionality on non-ANSI systems, such as Windows NT.

- [gm\\_bcopy\(\)](#)
- [gm\\_bzero\(\)](#)
- [gm\\_calloc\(\)](#)
- [gm\\_free\(\)](#)
- [gm\\_hex\\_dump\(\)](#)
- [gm\\_isprint\(\)](#)
- [gm\\_log2\\_roundup\(\)](#)
- [gm\\_malloc\(\)](#)
- [gm\\_memcmp\(\)](#)
- [gm\\_memset\(\)](#)
- [gm\\_on\\_exit\(\)](#)
- [gm\\_rand\(\)](#)
- [gm\\_rand\\_mod\(\)](#)
- [gm\\_srand\(\)](#)
- [gm\\_sleep\(\)](#)

- [gm\\_strdup\(\)](#)
- [gm\\_strerror\(\)](#)
- [gm\\_strlen\(\)](#)
- [gm\\_strncpy\(\)](#)
- [gm\\_strcmp\(\)](#)
- [gm\\_strncmp\(\)](#)
- [gm\\_strcasecmp\(\)](#)
- [gm\\_ticks\(\)](#)

### Utility Modules

These GM internal modules may be useful to GM developers.

- CRC Functions
  - [gm\\_crc\(\)](#)
  - [gm\\_crc\\_str\(\)](#)
- Hash Table Functions
  - [gm\\_create\\_hash\(\)](#)
  - [gm\\_destroy\\_hash\(\)](#)
  - [gm\\_hash\\_rekey\(\)](#)
  - [gm\\_hash\\_remove\(\)](#)
  - [gm\\_hash\\_find\(\)](#)
  - [gm\\_hash\\_insert\(\)](#)
  - [gm\\_hash\\_compare\\_strings\(\)](#)
  - [gm\\_hash\\_hash\\_string\(\)](#)
  - [gm\\_hash\\_compare longs\(\)](#)
  - [gm\\_hash\\_hash\\_long\(\)](#)
  - [gm\\_hash\\_compare\\_ints\(\)](#)
  - [gm\\_hash\\_hash\\_int\(\)](#)
  - [gm\\_hash\\_compare\\_ptrs\(\)](#)
  - [gm\\_hash\\_hash\\_ptr\(\)](#)
- Lookaside Functions
  - [gm\\_create\\_lookaside\(\)](#)
  - [gm\\_destroy\\_lookaside\(\)](#)

- [gm\\_lookaside\\_alloc\(\)](#)
- [gm\\_lookaside\\_zalloc\(\)](#)
- [gm\\_lookaside\\_free\(\)](#)
- [gm\\_create\\_mark\\_set\(\)](#)
- [gm\\_destroy\\_mark\\_set\(\)](#)

- Mark Functions

- [gm\\_mark\(\)](#)
- [gm\\_mark\\_is\\_valid\(\)](#)
- [gm\\_unmark\(\)](#)
- [gm\\_mark\\_set\\_unmark\\_all\(\)](#)

- Zone Functions

- [gm\\_zone\\_create\\_zone\(\)](#)
- [gm\\_zone\\_destroy\\_zone\(\)](#)
- [gm\\_zone\\_free\(\)](#)
- [gm\\_zone\\_malloc\(\)](#)
- [gm\\_zone\\_calloc\(\)](#)
- [gm\\_zone\\_addr\\_in\\_zone\(\)](#)

- Mutex Functions

- [gm\\_create\\_mutex\(\)](#)
- [gm\\_destroy\\_mutex\(\)](#)
- [gm\\_mutex\\_enter\(\)](#)
- [gm\\_mutex\\_exit\(\)](#)

- Buffer Debugging

- [gm\\_dump\\_buffers\(\)](#)
- [gm\\_register\\_buffer\(\)](#)
- [gm\\_unregister\\_buffer\(\)](#)

## 9.16 XIII. Utility Modules

Some of GM's internal modules may be useful to GM developers, so their APIs are exposed. These modules include the following:

### 9.16.1 1. CRC Functions

GM provides the following functions, which compute 32-bit CRCs on the contents of memory. These functions are not guaranteed to perform any particular variant of the CRC-32, but these functions are useful for creating robust hashing functions.

- [gm\\_crc\(\)](#)
- [gm\\_crc\\_str\(\)](#)

### 9.16.2 2. Hash Table

GM implements a generic hash table with a flexible interface. This module can automatically manage storage of fixed-size keys and/or data, or can allow the client to manage storage for keys and/or data. It allows the client to specify arbitrary hashing and comparison functions.

For example,

```
hash = gm_create_hash (gm_hash_compare_strings, gm_hash_hash_string,
 0, 0, 0, 0);
```

creates a hash table that uses null-terminated character string keys residing in client-managed storage, and returns pointers to data in client-managed storage. In this case, all pointers to hash keys and data passed by GM to the client will be the same as the pointers passed by the client to GM.

As another example,

```
hash = gm_create_hash (gm_hash_compare_ints, gm_hash_hash_int,
 sizeof (int), sizeof (struct my_big_struct),
 100, 0);
```

creates a hash table that uses `ints` as keys and returns pointers to copies of the inserted structures. All storage for the keys and data is automatically managed by the hash table. In this case, all pointers to hash keys and data passed by GM to the client will point to GM-managed buffers. This function also preallocates enough storage for 100 hash entries, guaranteeing that at least 100 key/data pairs can be inserted in the table if the hash table creation succeeds.

The automatic storage management option of GM not only is convenient, but also is extremely space efficient for keys and data no larger than a pointer, because when keys and data are no larger than a pointer, GM automatically stores them in the space reserved for the pointer to the key or data, rather than allocating a separate buffer.

Note that all keys and data buffers are referred to by pointers, not by value. This allows keys and data buffers of arbitrary size to be used. As a special (but common) case, however, one may wish to use pointers as keys directly, rather than use what they point to. In this special case, use the following initialization, and pass the keys (pointers) directly to the API, rather than the usual references to the keys.

```
hash = gm_create_hash (gm_hash_compare_ptrs, gm_hash_hash_ptr,
 0, DATA_LEN, MIN_CNT, FLAGS);
```

While it is possible to specify a `KEY_LEN` of `'sizeof (void *)'` during initialization and treat pointer keys just like any other keys, the API above is more efficient, more convenient, and completely architecture independent.

Some day the GM hash table API may be extended, but the current API is as follows:

- [gm\\_create\\_hash\(\)](#)
- [gm\\_destroy\\_hash\(\)](#)
- [gm\\_hash\\_rekey\(\)](#)
- [gm\\_hash\\_remove\(\)](#)
- [gm\\_hash\\_find\(\)](#)
- [gm\\_hash\\_insert\(\)](#)

The parameters are as follows:

- `GM_CLIENT_COMPARE` The function used to compare keys and may be any of [gm\\_hash\\_compare\\_ints\(\)](#), [gm\\_hash\\_compare\\_longs\(\)](#), [gm\\_hash\\_compare\\_ptrs\(\)](#), [gm\\_hash\\_compare\\_strings\(\)](#), or may be a client-defined function.

- `GM_CLIENT_HASH` The function to be used to hash keys and may be any of [gm\\_hash\\_hash\\_int\(\)](#), [gm\\_hash\\_hash\\_long\(\)](#), [gm\\_hash\\_hash\\_ptr\(\)](#), [gm\\_hash\\_hash\\_string\(\)](#), or may be a client-defined function.
- `KEY_LEN` specifies the length of the keys to be used for the hash table, or '0' if the keys should not be copied into GM-managed buffers.
- `DATA_LEN` specifies the length of the data to be stored in the hash table, or '0' if the data should not be copied into GM-managed buffers.
- `MIN_CNT` specifies the number of entries for which storage should be preallocated.
- `FLAGS` should be '0' because no flags are currently defined.

### 9.16.3 3. Lookaside List

GM implements a lookaside list, which may be used to manage small fixed-length blocks more efficiently than [gm\\_malloc\(\)](#) and [gm\\_free\(\)](#). Lookaside lists can also be used to ensure that at least a minimum number of blocks are available for allocation at all times.

GM lookaside lists have the following API:

- [gm\\_create\\_lookaside\(\)](#)
- [gm\\_destroy\\_lookaside\(\)](#)
- [gm\\_lookaside\\_alloc\(\)](#)
- [gm\\_lookaside\\_zalloc\(\)](#)
- [gm\\_lookaside\\_free\(\)](#)

### 9.16.4 4. Marks

The GM "mark" API was introduced in GM-1.4. It allows the creation and destruction of mark sets, which allow mark addition, mark removal, and test for mark in mark set



operations to be performed in constant time. Marks may be members of only one mark set at a time. Marks have the very unusual property that they need not be initialized before use.

All operations on marks are extremely efficient. Mark initialization requires zero time. Removing a mark from a mark set and testing for mark inclusion in a mark set take constant time. Addition of a mark to a mark set takes  $O(\text{constant})$  time, assuming the marks set was created with support for a sufficient number of marks; otherwise, it requires  $O(\text{constant})$  average time. Finally, creation and destruction of a mark set take time comperable to the time required for a single call to `malloc()` and `free()`, respectively.

Because marks need not be initialized before use, they can actually be used to determine if other objects have been initialized. This is done by putting a mark in the object, and adding the mark to a "mark set of marks in initialized objects" once the object has been initialized. This is similar to one common use of "magic numbers" for debugging purposes, except that it is immune to the possibility that the uninitialized magic number contained the magic number before initialization, so such marks can be used for non-debugging purposes. Therefore, marks can be used in ways that magic numbers cannot.

Marks have a nice set of properties that each mark in a mark set has a unique value and if this value is corrupted, then the mark is implicitly removed from the mark set. This makes marks useful for detecting memory corruption, and are less prone to false negatives than are magic numbers, which proliferate copies of a single value.

Finally, marks are location-dependent. This means that if a mark is copied, the copy will not be a member of the mark set.

The following APIs are provided:

- `gm_create_mark_set()`
- `gm_destroy_mark_set()`
- `gm_mark()`
- `gm_mark_is_valid()`
- `gm_unmark()`
- `gm_mark_set_unmark_all()`

### 9.16.5 5. Zones

These GM API routines manage an externally specified *zone* of memory. A zone is a chunk of memory starting and ending on page boundaries. The size and state of each area are encoded in a pair of bit-arrays. All allocated (or freed) areas are maximally aligned.

The following APIs are provided:

- [gm\\_zone\\_create\\_zone\(\)](#)
- [gm\\_zone\\_destroy\\_zone\(\)](#)
- [gm\\_zone\\_free\(\)](#)
- [gm\\_zone\\_malloc\(\)](#)
- [gm\\_zone\\_calloc\(\)](#)
- [gm\\_zone\\_addr\\_in\\_zone\(\)](#)

### 9.16.6 6. Mutexes

GM mutex routines have the following API:

- [gm\\_create\\_mutex\(\)](#)
- [gm\\_destroy\\_mutex\(\)](#)
- [gm\\_mutex\\_enter\(\)](#)
- [gm\\_mutex\\_exit\(\)](#)

### 9.16.7 7. Buffer Debugging

There are three API functions provided for buffer debugging:

- [gm\\_dump\\_buffers\(\)](#)

- [gm\\_register\\_buffer\(\)](#)
- [gm\\_unregister\\_buffer\(\)](#)

---

# Index

- `_GM_NEW_PUT_NOTIFICATION_EVENT`
  - `gm.h`, [56](#)
  - `_gm_get_route`
  - `gm.h`, [74](#)
  - Deprecated GM API functions, [13](#)
  - `gm.h`, [43](#)
    - `_GM_NEW_PUT_NOTIFICATION_EVENT`, [56](#)
    - `_gm_get_route`, [74](#)
    - `gm_abort`, [56](#)
    - `GM_ABORTED`, [54](#)
    - `GM_ACCESS_DENIED`, [54](#)
    - `GM_ALARM_EVENT`, [56](#)
    - `gm_alloc_pages`, [64](#)
    - `gm_alloc_send_token`, [57](#)
    - `gm_allow_remote_memory_access`, [57](#)
    - `GM_API_VERSION`, [51](#)
    - `GM_API_VERSION_1_0`, [50](#)
    - `GM_API_VERSION_1_1`, [50](#)
    - `GM_API_VERSION_1_2`, [50](#)
    - `GM_API_VERSION_1_3`, [50](#)
    - `GM_API_VERSION_1_4`, [51](#)
    - `GM_API_VERSION_1_5`, [51](#)
    - `GM_API_VERSION_1_6`, [51](#)
    - `GM_API_VERSION_2_0`, [51](#)
    - `GM_API_VERSION_2_0_6`, [51](#)
    - `gm_bcopy`, [57](#)
    - `gm_blocking_receive`, [58](#)
    - `gm_blocking_receive_no_spin`, [58](#)
    - `GM_BUSY`, [53](#)
    - `gm_bzero`, [58](#)
    - `gm_calloc`, [58](#)
    - `gm_cancel_alarm`, [59](#)
    - `GM_CLONE_BUSY`, [54](#)
    - `gm_close`, [59](#)
    - `GM_CPU_alpha`, [51](#)
    - `gm_crc`, [82](#)
    - `gm_crc_str`, [83](#)
    - `gm_create_hash`, [78](#)
    - `gm_create_lookaside`, [75](#)
    - `gm_create_mark_set`, [97](#)
    - `gm_create_mutex`, [86](#)
    - `GM_DATA_CORRUPTED`, [55](#)
    - `gm_datagram_send`, [59](#)
    - `gm_datagram_send_4`, [59](#)
    - `gm_deregister_memory`, [59](#)
    - `gm_destroy_hash`, [78](#)
    - `gm_destroy_lookaside`, [76](#)
    - `gm_destroy_mark_set`, [98](#)
    - `gm_destroy_mutex`, [86](#)
    - `GM_DEV_NOT_FOUND`, [55](#)
    - `gm_directcopy_get`, [91](#)
    - `gm_directed_send_with_callback`, [59](#)
    - `gm_dma_calloc`, [60](#)
    - `gm_dma_free`, [60](#)
    - `gm_dma_malloc`, [60](#)
    - `gm_drop_sends`, [90](#)
    - `gm_dump_buffers`, [74](#)
    - `gm_eprintf`, [95](#)
    - `gm_exit`, [93](#)
    - `GM_FAILURE`, [53](#)
    - `GM_FAST_HIGH_PEER_RECV_EVENT`, [56](#)
    - `GM_FAST_HIGH_RECV_EVENT`, [56](#)
    - `GM_FAST_PEER_RECV_EVENT`, [56](#)
    - `GM_FAST_RECV_EVENT`, [56](#)
-

- gm\_finalize, 85
- GM\_FIRMWARE\_NOT\_RUNNING, 55
- gm\_flush\_alarm, 61
- gm\_free, 61
- gm\_free\_pages, 65
- gm\_free\_send\_token, 61
- gm\_free\_send\_tokens, 61
- gm\_get\_host\_name, 61
- gm\_get\_mapper\_unique\_id, 62
- gm\_get\_node\_id, 62
- gm\_get\_node\_type, 62
- gm\_get\_port\_id, 69
- gm\_get\_unique\_board\_id, 62
- gm\_getpid, 91
- gm\_global\_id\_to\_node\_id, 100
- GM\_HARDWARE\_FAULT, 55
- gm\_hash\_compare\_ints, 80
- gm\_hash\_compare longs, 79
- gm\_hash\_compare\_ptrs, 81
- gm\_hash\_compare\_strings, 79
- gm\_hash\_find, 78
- gm\_hash\_hash\_int, 81
- gm\_hash\_hash\_long, 80
- gm\_hash\_hash\_ptr, 82
- gm\_hash\_hash\_string, 79
- gm\_hash\_insert, 78
- gm\_hash\_rekey, 78
- gm\_hash\_remove, 78
- gm\_hex\_dump, 62
- GM\_HIGH\_PEER\_RECV\_EVENT, 56
- GM\_HIGH\_PRIORITY, 55
- GM\_HIGH\_RECV\_EVENT, 56
- gm\_host\_name\_to\_node\_id, 62
- gm\_host\_name\_to\_node\_id\_ex, 100
- GM\_INCOMPATIBLE\_LIB\_AND\_DRIVER, 54
- gm\_init, 84
- gm\_initialize\_alarm, 63
- GM\_INPUT\_BUFFER\_TOO\_SMALL, 53
- GM\_INTERNAL\_ERROR, 54
- GM\_INTERRUPTED, 54
- GM\_INVALID\_COMMAND, 54
- GM\_INVALID\_PARAMETER, 54
- GM\_INVALID\_PORT\_NUMBER, 55
- gm\_isprint, 63
- gm\_log2\_roundup, 85
- gm\_log2\_roundup\_table, 100
- gm\_lookaside\_alloc, 76
- gm\_lookaside\_free, 77
- gm\_lookaside\_zalloc, 77
- GM\_LOW\_PRIORITY, 55
- gm\_malloc, 63
- gm\_mark, 96
- gm\_mark\_is\_valid, 97
- GM\_MAX\_DMA\_GRANULARITY, 52
- GM\_MAX\_HOST\_NAME\_LEN, 51
- gm\_max\_length\_for\_size, 65
- gm\_max\_node\_id, 66
- gm\_max\_node\_id\_in\_use, 95
- GM\_MAX\_PORT\_NAME\_LEN, 51
- gm\_memcmp, 66
- gm\_memorize\_message, 66
- GM\_MEMORY\_FAULT, 54
- gm\_memset, 95
- gm\_min\_message\_size, 67
- gm\_min\_size\_for\_length, 67
- GM\_MINOR\_OVERFLOW, 55
- gm\_mtu, 68
- gm\_mutex\_enter, 86
- gm\_mutex\_exit, 87
- GM\_NEW\_FAST\_RECV\_EVENT, 56
- GM\_NEW\_RECV\_EVENT, 56
- GM\_NEW\_SENDS\_FAILED\_EVENT, 56
- gm\_next\_event\_peek, 69
- GM\_NO\_DRIVER\_SUPPORT, 55
- GM\_NO\_MATCH, 55
- GM\_NO\_RECV\_EVENT, 55
- GM\_NO\_SUCH\_DEVICE, 54
- GM\_NO\_SUCH\_NODE\_ID, 51

- GM\_NODE\_ID\_NOT\_YET\_SET, 54
- gm\_node\_id\_to\_global\_id, 100
- gm\_node\_id\_to\_host\_name, 68
- gm\_node\_id\_to\_host\_name\_ex, 100
- gm\_node\_id\_to\_unique\_id, 68
- GM\_NOT\_SUPPORTED\_IN\_KERNEL, 55
- GM\_NOT\_SUPPORTED\_ON\_ARCH, 55
- GM\_NUM\_ELEM, 52
- gm\_num\_ports, 68
- GM\_NUM\_PRIORITIES, 55
- gm\_num\_receive\_tokens, 69
- GM\_NUM\_RECV\_EVENT\_TYPES, 56
- gm\_num\_send\_tokens, 68
- gm\_on\_exit, 99
- gm\_open, 69
- GM\_OUT\_OF\_MEMORY, 54
- GM\_OUTPUT\_BUFFER\_TOO\_SMALL, 53
- gm\_page\_alloc, 64
- gm\_page\_free, 64
- GM\_PAGE\_TABLE\_FULL, 55
- GM\_PEER\_RECV\_EVENT, 56
- GM\_PERMISSION\_DENIED, 54
- gm\_perror, 92
- gm\_pid\_t, 53
- GM\_POWER\_OF\_TWO, 52
- gm\_printf, 93
- gm\_priority, 55
- gm\_provide\_receive\_buffer\_with\_tag, 69
- GM\_PTE\_REF\_CNT\_OVERFLOW, 55
- gm\_put, 99
- gm\_rand, 83
- gm\_rand\_mod, 84
- GM\_RDMA\_GRANULARITY, 52
- gm\_receive, 69
- gm\_receive\_pending, 69
- GM\_RECV\_EVENT, 56
- gm\_rcv\_event\_type, 55
- gm\_register\_buffer, 74
- gm\_register\_memory, 69
- gm\_remote\_ptr\_t, 53
- gm\_resume\_sending, 89
- gm\_send\_completion\_callback\_t, 53
- GM\_SEND\_DROPPED, 54
- GM\_SEND\_ORPHANED, 55
- GM\_SEND\_PORT\_CLOSED, 54
- GM\_SEND\_REJECTED, 54
- GM\_SEND\_TARGET\_NODE\_UNREACHABLE, 54
- GM\_SEND\_TARGET\_PORT\_CLOSED, 54
- GM\_SEND\_TIMED\_OUT, 54
- gm\_send\_to\_peer\_with\_callback, 70
- gm\_send\_token\_available, 70
- gm\_send\_with\_callback, 70
- GM\_SENDS\_FAILED\_EVENT, 55
- gm\_set\_acceptable\_sizes, 70
- gm\_set\_alarm, 71
- gm\_set\_enable\_nack\_down, 94
- gm\_sleep, 93
- gm\_srand, 84
- gm\_status, 53
- gm\_status\_t, 53
- GM\_STILL\_SHUTTING\_DOWN, 54
- gm\_strcmp, 72
- gm\_strdup, 96
- gm\_strerror, 94
- gm\_strlen, 71
- gm\_strncasecmp, 73
- gm\_strncmp, 72
- gm\_strncpy, 71
- GM\_STRUCT\_CONTAINING, 52
- GM\_SUCCESS, 53
- gm\_ticks, 73
- GM\_TIMED\_OUT, 55
- GM\_TRY\_AGAIN, 53
- GM\_UC\_ERROR, 55
- GM\_UNATTACHED, 54
- gm\_unique\_id, 74

- gm\_unique\_id\_to\_node\_id, 74
- gm\_unknown, 74
- gm\_unmark, 98
- gm\_unmark\_all, 98
- gm\_unregister\_buffer, 75
- GM\_UNSUPPORTED\_DEVICE, 54
- GM\_UNTRANSLATED\_SYSTEM\_ERROR, 54
- GM\_USER\_ERROR, 55
- GM\_YP\_NO\_MATCH, 55
- gm\_zone\_addr\_in\_zone, 89
- gm\_zone\_calloc, 89
- gm\_zone\_create\_zone, 87
- gm\_zone\_destroy\_zone, 88
- gm\_zone\_free, 88
- gm\_zone\_malloc, 88
- gm\_abort
  - gm.h, 56
  - gm\_abort.c, 101
- gm\_abort.c, 101
  - gm\_abort, 101
- GM\_ABORTED
  - gm.h, 54
- GM\_ACCESS\_DENIED
  - gm.h, 54
- GM\_ALARM\_EVENT
  - gm.h, 56
- gm\_alloc\_pages
  - gm.h, 64
  - gm\_alloc\_pages.c, 102
- gm\_alloc\_pages.c, 102
  - gm\_alloc\_pages, 102
  - gm\_free\_pages, 102
- gm\_alloc\_send\_token
  - gm.h, 57
  - gm\_alloc\_send\_token.c, 104
- gm\_alloc\_send\_token.c, 104
  - gm\_alloc\_send\_token, 104
- gm\_allow\_remote\_memory\_access
  - gm.h, 57
  - gm\_allow\_remote\_memory\_access.c, 105
- gm\_allow\_remote\_memory\_access.c, 105
  - gm\_allow\_remote\_memory\_access, 105
- gm\_allow\_remote\_memory\_access, 105
  - gm\_allow\_remote\_memory\_access, 105
- GM\_API\_VERSION
  - gm.h, 51
- GM\_API\_VERSION\_1\_0
  - gm.h, 50
- GM\_API\_VERSION\_1\_1
  - gm.h, 50
- GM\_API\_VERSION\_1\_2
  - gm.h, 50
- GM\_API\_VERSION\_1\_3
  - gm.h, 50
- GM\_API\_VERSION\_1\_4
  - gm.h, 51
- GM\_API\_VERSION\_1\_5
  - gm.h, 51
- GM\_API\_VERSION\_1\_6
  - gm.h, 51
- GM\_API\_VERSION\_2\_0
  - gm.h, 51
- GM\_API\_VERSION\_2\_0\_6
  - gm.h, 51
- GM\_AREA\_FOR\_PTR
  - gm\_zone.c, 247
- gm\_bcopy
  - gm.h, 57
  - gm\_bcopy.c, 106
- gm\_bcopy.c, 106
  - gm\_bcopy, 106
- gm\_blocking\_receive
  - gm.h, 58
  - gm\_blocking\_receive.c, 107
- gm\_blocking\_receive.c, 107
  - gm\_blocking\_receive, 107
- gm\_blocking\_receive\_no\_spin
  - gm.h, 58
  - gm\_blocking\_receive\_no\_spin.c, 109
- gm\_blocking\_receive\_no\_spin.c, 109
  - gm\_blocking\_receive\_no\_spin, 109
- gm\_blocking\_receive\_no\_spin, 109
  - gm\_blocking\_receive\_no\_spin, 109
- gm\_buf\_status\_name
  - gm\_debug\_buffers.c, 122
- GM\_BUSY
  - gm.h, 53
- gm\_bzero

- gm.h, 58
- gm\_bzero.c, 111
- gm\_bzero.c, 111
  - gm\_bzero, 111
- gm\_calloc
  - gm.h, 58
  - gm\_calloc.c, 112
- gm\_calloc.c, 112
  - gm\_calloc, 112
- gm\_cancel\_alarm
  - gm.h, 59
  - gm\_set\_alarm.c, 229
- GM\_CLONE\_BUSY
  - gm.h, 54
- gm\_close
  - gm.h, 59
  - gm\_close.c, 113
- gm\_close.c, 113
  - gm\_close, 113
- GM\_CPU\_alpha
  - gm.h, 51
- gm\_crc
  - gm.h, 82
  - gm\_crc.c, 115
- gm\_crc.c, 115
  - gm\_crc, 115
  - gm\_crc\_str, 115
- gm\_crc\_str
  - gm.h, 83
  - gm\_crc.c, 115
- gm\_create\_hash
  - gm.h, 78
  - gm\_hash.c, 152
- gm\_create\_lookaside
  - gm.h, 75
  - gm\_lookaside.c, 170
- gm\_create\_mark\_set
  - gm.h, 97
  - gm\_mark.c, 176
- gm\_create\_mutex
  - gm.h, 86
  - gm\_mutex.c, 189
- GM\_DATA\_CORRUPTED
  - gm.h, 55
- gm\_datagram\_send
  - gm.h, 59
- gm\_datagram\_send.c, 117
- gm\_datagram\_send.c, 117
  - gm\_datagram\_send, 117
- gm\_datagram\_send\_4
  - gm.h, 59
  - gm\_datagram\_send\_4.c, 119
- gm\_datagram\_send\_4.c, 119
  - gm\_datagram\_send\_4, 119
- gm\_debug\_buffers.c, 121
  - gm\_buf\_status\_name, 122
  - gm\_dump\_buffers, 122
  - gm\_register\_buffer, 122
  - gm\_unregister\_buffer, 121
- gm\_deregister.c, 124
  - gm\_deregister\_memory, 124
- gm\_deregister\_memory
  - gm.h, 59
  - gm\_deregister.c, 124
- gm\_destroy\_hash
  - gm.h, 78
  - gm\_hash.c, 153
- gm\_destroy\_lookaside
  - gm.h, 76
  - gm\_lookaside.c, 170
- gm\_destroy\_mark\_set
  - gm.h, 98
  - gm\_mark.c, 176
- gm\_destroy\_mutex
  - gm.h, 86
  - gm\_mutex.c, 189
- GM\_DEV\_NOT\_FOUND
  - gm.h, 55
- gm\_directcopy.c, 126
  - gm\_directcopy\_get, 126
- gm\_directcopy\_get
  - gm.h, 91
  - gm\_directcopy.c, 126
- gm\_directed\_send.c, 128
- gm\_directed\_send\_with\_callback
  - gm.h, 59
- gm\_dma\_calloc
  - gm.h, 60
  - gm\_dma\_calloc.c, 129
- gm\_dma\_calloc.c, 129
  - gm\_dma\_calloc, 129
- gm\_dma\_free



- gm.h, 60
- gm\_dma\_malloc.c, 130
- gm\_dma\_malloc
  - gm.h, 60
  - gm\_dma\_malloc.c, 131
- gm\_dma\_malloc.c, 130
  - gm\_dma\_free, 130
  - gm\_dma\_malloc, 131
- gm\_drop\_sends
  - gm.h, 90
  - gm\_drop\_sends.c, 132
- gm\_drop\_sends.c, 132
  - gm\_drop\_sends, 132
- gm\_dump\_buffers
  - gm.h, 74
  - gm\_debug\_buffers.c, 122
- gm\_eprintf
  - gm.h, 95
  - gm\_eprintf.c, 134
- gm\_eprintf.c, 134
  - gm\_eprintf, 134
- gm\_exit
  - gm.h, 93
  - gm\_exit.c, 135
- gm\_exit.c, 135
  - gm\_exit, 135
- GM\_FAILURE
  - gm.h, 53
- GM\_FAST\_HIGH\_PEER\_RECV\_-  
EVENT
  - gm.h, 56
- GM\_FAST\_HIGH\_RECV\_EVENT
  - gm.h, 56
- GM\_FAST\_PEER\_RECV\_EVENT
  - gm.h, 56
- GM\_FAST\_RECV\_EVENT
  - gm.h, 56
- gm\_finalize
  - gm.h, 85
  - gm\_init.c, 164
- GM\_FIRMWARE\_NOT\_RUNNING
  - gm.h, 55
- gm\_flush\_alarm
  - gm.h, 61
  - gm\_flush\_alarm.c, 136
- gm\_flush\_alarm.c, 136
  - gm\_flush\_alarm, 136
- gm\_free
  - gm.h, 61
  - gm\_free.c, 137
- gm\_free.c, 137
  - gm\_free, 137
- gm\_free\_mdebug, 15
- gm\_free\_page, 16
- gm\_free\_pages
  - gm.h, 65
  - gm\_alloc\_pages.c, 102
- gm\_free\_send\_token
  - gm.h, 61
  - gm\_free\_send\_token.c, 138
- gm\_free\_send\_token.c, 138
  - gm\_free\_send\_token, 138
- gm\_free\_send\_tokens
  - gm.h, 61
  - gm\_free\_send\_tokens.c, 139
- gm\_free\_send\_tokens.c, 139
  - gm\_free\_send\_tokens, 139
- gm\_get
  - gm\_get.c, 140
- gm\_get.c, 140
  - gm\_get, 140
- gm\_get\_host\_name
  - gm.h, 61
  - gm\_get\_host\_name.c, 142
- gm\_get\_host\_name.c, 142
  - gm\_get\_host\_name, 142
- gm\_get\_mapper\_unique\_id
  - gm.h, 62
  - gm\_get\_mapper\_unique\_id.c, 143
- gm\_get\_mapper\_unique\_id.c, 143
  - gm\_get\_mapper\_unique\_id, 143
- gm\_get\_node\_id
  - gm.h, 62
  - gm\_get\_node\_id.c, 144
- gm\_get\_node\_id.c, 144
  - gm\_get\_node\_id, 144
- gm\_get\_node\_type
  - gm.h, 62
  - gm\_get\_node\_type.c, 145
- gm\_get\_node\_type.c, 145
  - gm\_get\_node\_type, 145
- gm\_get\_port\_id

- gm.h, 69
- gm\_get\_port\_id.c, 146
- gm\_get\_port\_id.c, 146
  - gm\_get\_port\_id, 146
- gm\_get\_unique\_board\_id
  - gm.h, 62
  - gm\_get\_unique\_board\_id.c, 147
- gm\_get\_unique\_board\_id.c, 147
  - gm\_get\_unique\_board\_id, 147
- gm\_getpid
  - gm.h, 91
  - gm\_getpid.c, 148
- gm\_getpid.c, 148
  - gm\_getpid, 148
- gm\_global\_id\_to\_node\_id
  - gm.h, 100
- gm\_handle\_sent\_tokens.c, 149
- GM\_HARDWARE\_FAULT
  - gm.h, 55
- gm\_hash, 17
- gm\_hash.c, 150
  - gm\_create\_hash, 152
  - gm\_destroy\_hash, 153
  - gm\_hash\_compare\_ints, 157
  - gm\_hash\_compare longs, 156
  - gm\_hash\_compare\_ptrs, 158
  - gm\_hash\_compare\_strings, 155
  - gm\_hash\_entry\_t, 152
  - gm\_hash\_find, 154
  - gm\_hash\_hash\_int, 158
  - gm\_hash\_hash\_long, 157
  - gm\_hash\_hash\_ptr, 159
  - gm\_hash\_hash\_string, 156
  - gm\_hash\_insert, 155
  - gm\_hash\_rekey, 154
  - gm\_hash\_remove, 154
  - gm\_hash\_segment\_t, 152
  - gm\_hash\_t, 152
- gm\_hash\_compare\_ints
  - gm.h, 80
  - gm\_hash.c, 157
- gm\_hash\_compare longs
  - gm.h, 79
  - gm\_hash.c, 156
- gm\_hash\_compare\_ptrs
  - gm.h, 81
- gm\_hash.c, 158
- gm\_hash\_compare\_strings
  - gm.h, 79
  - gm\_hash.c, 155
- gm\_hash\_entry, 18
- gm\_hash\_entry\_t
  - gm\_hash.c, 152
- gm\_hash\_find
  - gm.h, 78
  - gm\_hash.c, 154
- gm\_hash\_hash\_int
  - gm.h, 81
  - gm\_hash.c, 158
- gm\_hash\_hash\_long
  - gm.h, 80
  - gm\_hash.c, 157
- gm\_hash\_hash\_ptr
  - gm.h, 82
  - gm\_hash.c, 159
- gm\_hash\_hash\_string
  - gm.h, 79
  - gm\_hash.c, 156
- gm\_hash\_insert
  - gm.h, 78
  - gm\_hash.c, 155
- gm\_hash\_rekey
  - gm.h, 78
  - gm\_hash.c, 154
- gm\_hash\_remove
  - gm.h, 78
  - gm\_hash.c, 154
- gm\_hash\_segment, 19
- gm\_hash\_segment\_t
  - gm\_hash.c, 152
- gm\_hash\_t
  - gm\_hash.c, 152
- gm\_hex\_dump
  - gm.h, 62
  - gm\_hex\_dump.c, 160
- gm\_hex\_dump.c, 160
  - gm\_hex\_dump, 160
- GM\_HIGH\_PEER\_RECV\_EVENT
  - gm.h, 56
- GM\_HIGH\_PRIORITY
  - gm.h, 55
- GM\_HIGH\_RECV\_EVENT

- gm.h, 56
- gm\_host\_name\_to\_node\_id
  - gm.h, 62
  - gm\_host\_name\_to\_node\_id.c, 162
- gm\_host\_name\_to\_node\_id.c, 161
  - gm\_host\_name\_to\_node\_id, 162
  - gm\_host\_name\_to\_node\_id\_ex, 161
- gm\_host\_name\_to\_node\_id\_ex
  - gm.h, 100
  - gm\_host\_name\_to\_node\_id.c, 161
- GM\_INCOMPATIBLE\_LIB\_AND\_DRIVER
  - gm.h, 54
- gm\_init
  - gm.h, 84
  - gm\_init.c, 163
- gm\_init.c, 163
  - gm\_finalize, 164
  - gm\_init, 163
- gm\_initialize\_alarm
  - gm.h, 63
  - gm\_set\_alarm.c, 229
- GM\_INPUT\_BUFFER\_TOO\_SMALL
  - gm.h, 53
- GM\_INTERNAL\_ERROR
  - gm.h, 54
- GM\_INTERRUPTED
  - gm.h, 54
- GM\_INVALID\_COMMAND
  - gm.h, 54
- GM\_INVALID\_PARAMETER
  - gm.h, 54
- GM\_INVALID\_PORT\_NUMBER
  - gm.h, 55
- gm\_isprint
  - gm.h, 63
  - gm\_isprint.c, 165
- gm\_isprint.c, 165
  - gm\_isprint, 165
- gm\_log2.c, 166
  - gm\_log2\_roundup, 166
  - gm\_log2\_roundup\_table, 167
- gm\_log2\_roundup
  - gm.h, 85
  - gm\_log2.c, 166
- gm\_log2\_roundup\_table
  - gm.h, 100
  - gm\_log2.c, 167
- gm\_lookaside, 20
  - segment\_list, 20
- gm\_lookaside.c, 168
  - gm\_create\_lookaside, 170
  - gm\_destroy\_lookaside, 170
  - gm\_lookaside\_alloc, 169
  - gm\_lookaside\_free, 170
  - gm\_lookaside\_zalloc, 169
- gm\_lookaside::gm\_lookaside\_segment\_list, 21
- gm\_lookaside\_alloc
  - gm.h, 76
  - gm\_lookaside.c, 169
- gm\_lookaside\_free
  - gm.h, 77
  - gm\_lookaside.c, 170
- gm\_lookaside\_segment, 22
- gm\_lookaside\_zalloc
  - gm.h, 77
  - gm\_lookaside.c, 169
- GM\_LOW\_PRIORITY
  - gm.h, 55
- gm\_malloc
  - gm.h, 63
  - gm\_malloc.c, 172
- gm\_malloc.c, 172
  - gm\_malloc, 172
- gm\_mark
  - gm.h, 96
  - gm\_mark.c, 175
- gm\_mark.c, 173
  - gm\_create\_mark\_set, 176
  - gm\_destroy\_mark\_set, 176
  - gm\_mark, 175
  - gm\_mark\_is\_valid, 175
  - gm\_mark\_reference\_t, 174
  - gm\_unmark, 175
  - gm\_unmark\_all, 177
- gm\_mark\_is\_valid
  - gm.h, 97
  - gm\_mark.c, 175
- gm\_mark\_reference, 23
- gm\_mark\_reference\_t

- gm\_mark.c, 174
- gm\_mark\_set, 24
- GM\_MAX\_DMA\_GRANULARITY
  - gm.h, 52
- GM\_MAX\_HOST\_NAME\_LEN
  - gm.h, 51
- gm\_max\_length\_for\_size
  - gm.h, 65
  - gm\_max\_length\_for\_size.c, 178
- gm\_max\_length\_for\_size.c, 178
  - gm\_max\_length\_for\_size, 178
- gm\_max\_node\_id
  - gm.h, 66
  - gm\_max\_node\_id.c, 179
- gm\_max\_node\_id.c, 179
  - gm\_max\_node\_id, 179
- gm\_max\_node\_id\_in\_use
  - gm.h, 95
  - gm\_max\_node\_id\_in\_use.c, 180
- gm\_max\_node\_id\_in\_use.c, 180
  - gm\_max\_node\_id\_in\_use, 180
- GM\_MAX\_PORT\_NAME\_LEN
  - gm.h, 51
- gm\_mdebug, 25
- gm\_memcmp
  - gm.h, 66
  - gm\_memcmp.c, 182
- gm\_memcmp.c, 182
  - gm\_memcmp, 182
- gm\_memorize\_message
  - gm.h, 66
  - gm\_memorize\_message.c, 183
- gm\_memorize\_message.c, 183
  - gm\_memorize\_message, 183
- GM\_MEMORY\_FAULT
  - gm.h, 54
- gm\_memset
  - gm.h, 95
  - gm\_memset.c, 185
- gm\_memset.c, 185
  - gm\_memset, 185
- gm\_min\_message\_size
  - gm.h, 67
  - gm\_min\_message\_size.c, 186
- gm\_min\_message\_size.c, 186
  - gm\_min\_message\_size, 186
- gm\_min\_size\_for\_length
  - gm.h, 67
  - gm\_min\_size\_for\_length.c, 187
- gm\_min\_size\_for\_length.c, 187
  - gm\_min\_size\_for\_length, 187
- GM\_MINOR\_OVERFLOW
  - gm.h, 55
- gm\_mtu
  - gm.h, 68
  - gm\_mtu.c, 188
- gm\_mtu.c, 188
  - gm\_mtu, 188
- gm\_mutex.c, 189
  - gm\_create\_mutex, 189
  - gm\_destroy\_mutex, 189
  - gm\_mutex\_enter, 190
  - gm\_mutex\_exit, 190
- gm\_mutex\_enter
  - gm.h, 86
  - gm\_mutex.c, 190
- gm\_mutex\_exit
  - gm.h, 87
  - gm\_mutex.c, 190
- GM\_NEW\_FAST\_RECV\_EVENT
  - gm.h, 56
- GM\_NEW\_RECV\_EVENT
  - gm.h, 56
- GM\_NEW SENDS\_FAILED\_EVENT
  - gm.h, 56
- gm\_next\_event\_peek
  - gm.h, 69
  - gm\_next\_event\_peek.c, 191
- gm\_next\_event\_peek.c, 191
  - gm\_next\_event\_peek, 191
- GM\_NO\_DRIVER\_SUPPORT
  - gm.h, 55
- GM\_NO\_MATCH
  - gm.h, 55
- GM\_NO\_RECV\_EVENT
  - gm.h, 55
- GM\_NO\_SUCH\_DEVICE
  - gm.h, 54
- GM\_NO\_SUCH\_NODE\_ID
  - gm.h, 51
- GM\_NODE\_ID\_NOT\_YET\_SET
  - gm.h, 54

- gm\_node\_id\_to\_global\_id
  - gm.h, 100
- gm\_node\_id\_to\_host\_name
  - gm.h, 68
  - gm\_node\_id\_to\_host\_name.c, 192
- gm\_node\_id\_to\_host\_name.c, 192
  - gm\_node\_id\_to\_host\_name, 192
  - gm\_node\_id\_to\_host\_name\_ex, 193
- gm\_node\_id\_to\_host\_name\_ex
  - gm.h, 100
  - gm\_node\_id\_to\_host\_name.c, 193
- gm\_node\_id\_to\_unique\_id
  - gm.h, 68
  - gm\_node\_id\_to\_unique\_id.c, 194
- gm\_node\_id\_to\_unique\_id.c, 194
  - gm\_node\_id\_to\_unique\_id, 194
- GM\_NOT\_SUPPORTED\_IN\_KERNEL
  - gm.h, 55
- GM\_NOT\_SUPPORTED\_ON\_ARCH
  - gm.h, 55
- GM\_NUM\_ELEM
  - gm.h, 52
- gm\_num\_ports
  - gm.h, 68
  - gm\_num\_ports.c, 195
- gm\_num\_ports.c, 195
  - gm\_num\_ports, 195
- GM\_NUM\_PRIORITIES
  - gm.h, 55
- gm\_num\_receive\_tokens
  - gm.h, 69
  - gm\_num\_receive\_tokens.c, 196
- gm\_num\_receive\_tokens.c, 196
  - gm\_num\_receive\_tokens, 196
- GM\_NUM\_RECV\_EVENT\_TYPES
  - gm.h, 56
- gm\_num\_send\_tokens
  - gm.h, 68
  - gm\_num\_send\_tokens.c, 197
- gm\_num\_send\_tokens.c, 197
  - gm\_num\_send\_tokens, 197
- gm\_on\_exit
  - gm.h, 99
  - gm\_on\_exit.c, 198
- gm\_on\_exit.c, 198
  - gm\_on\_exit, 198
  - gm\_on\_exit\_record\_t, 198
- gm\_on\_exit\_record, 26
- gm\_on\_exit\_record\_t
  - gm\_on\_exit.c, 198
- gm\_open
  - gm.h, 69
  - gm\_open.c, 200
- gm\_open.c, 200
  - gm\_open, 200
- GM\_OUT\_OF\_MEMORY
  - gm.h, 54
- GM\_OUTPUT\_BUFFER\_TOO\_SMALL
  - gm.h, 53
- gm\_page\_alloc
  - gm.h, 64
  - gm\_page\_alloc.c, 202
- gm\_page\_alloc.c, 202
  - gm\_page\_alloc, 202
  - gm\_page\_free, 203
- gm\_page\_allocation\_record, 27
- gm\_page\_free
  - gm.h, 64
  - gm\_page\_alloc.c, 203
- GM\_PAGE\_TABLE\_FULL
  - gm.h, 55
- GM\_PEER\_RECV\_EVENT
  - gm.h, 56
- GM\_PERMISSION\_DENIED
  - gm.h, 54
- gm\_perror
  - gm.h, 92
  - gm\_perror.c, 204
- gm\_perror.c, 204
  - gm\_perror, 204
- gm\_pid\_t
  - gm.h, 53
- GM\_POWER\_OF\_TWO
  - gm.h, 52
- gm\_printf
  - gm.h, 93
  - gm\_printf.c, 205
- gm\_printf.c, 205
  - gm\_printf, 205

- gm\_priority
  - gm.h, 55
- gm\_provide\_receive\_buffer.c, 206
  - gm\_provide\_receive\_buffer\_-with\_tag, 206
- gm\_provide\_receive\_buffer\_with\_tag
  - gm.h, 69
  - gm\_provide\_receive\_buffer.c, 206
- GM\_PTE\_REF\_CNT\_OVERFLOW
  - gm.h, 55
- gm\_put
  - gm.h, 99
  - gm\_put.c, 208
- gm\_put.c, 208
  - gm\_put, 208
- gm\_rand
  - gm.h, 83
  - gm\_rand.c, 210
- gm\_rand.c, 210
  - gm\_rand, 210
  - gm\_srand, 210
- gm\_rand\_mod
  - gm.h, 84
  - gm\_rand\_mod.c, 212
- gm\_rand\_mod.c, 212
  - gm\_rand\_mod, 212
- GM\_RDMA\_GRANULARITY
  - gm.h, 52
- gm\_receive
  - gm.h, 69
  - gm\_receive.c, 213
- gm\_receive.c, 213
  - gm\_receive, 213
  - gm\_receive\_debug\_buffers, 214
- gm\_receive\_debug\_buffers
  - gm\_receive.c, 214
- gm\_receive\_pending
  - gm.h, 69
  - gm\_receive\_pending.c, 215
- gm\_receive\_pending.c, 215
  - gm\_receive\_pending, 215
- GM\_RECV\_EVENT
  - gm.h, 56
- gm\_recv\_event\_type
  - gm.h, 55
- gm\_register.c, 216
  - gm\_register\_memory, 217
  - gm\_register\_memory\_ex, 216
- gm\_register\_buffer
  - gm.h, 74
  - gm\_debug\_buffers.c, 122
- gm\_register\_memory
  - gm.h, 69
  - gm\_register.c, 217
- gm\_register\_memory\_ex
  - gm\_register.c, 216
- gm\_remote\_ptr\_n\_t, 28
- gm\_remote\_ptr\_t
  - gm.h, 53
- gm\_resume\_sending
  - gm.h, 89
  - gm\_resume\_sending.c, 219
- gm\_resume\_sending.c, 219
  - gm\_resume\_sending, 219
- gm\_s16\_n\_t, 29
- gm\_s32\_n\_t, 30
- gm\_s64\_n\_t, 31
- gm\_s8\_n\_t, 32
- gm\_s\_e\_context\_t, 33
- gm\_send.c, 221
  - gm\_send\_with\_callback, 221
- gm\_send\_completion\_callback\_t
  - gm.h, 53
- GM\_SEND\_DROPPED
  - gm.h, 54
- GM\_SEND\_ORPHANED
  - gm.h, 55
- GM\_SEND\_PORT\_CLOSED
  - gm.h, 54
- GM\_SEND\_REJECTED
  - gm.h, 54
- GM\_SEND\_TARGET\_NODE\_UNREACHABLE
  - gm.h, 54
- GM\_SEND\_TARGET\_PORT\_CLOSED
  - gm.h, 54
- GM\_SEND\_TIMED\_OUT
  - gm.h, 54
- gm\_send\_to\_peer.c, 223
  - gm\_send\_to\_peer\_with\_callback, 223

- gm\_send\_to\_peer\_with\_callback
  - gm.h, 70
  - gm\_send\_to\_peer.c, 223
- gm\_send\_token\_available
  - gm.h, 70
  - gm\_send\_token\_available.c, 225
- gm\_send\_token\_available.c, 225
  - gm\_send\_token\_available, 225
- gm\_send\_with\_callback
  - gm.h, 70
  - gm\_send.c, 221
- GM\_SENDS\_FAILED\_EVENT
  - gm.h, 55
- gm\_set\_acceptable\_sizes
  - gm.h, 70
  - gm\_set\_acceptable\_sizes.c, 226
- gm\_set\_acceptable\_sizes.c, 226
  - gm\_set\_acceptable\_sizes, 226
- gm\_set\_alarm
  - gm.h, 71
  - gm\_set\_alarm.c, 229
- gm\_set\_alarm.c, 228
  - gm\_cancel\_alarm, 229
  - gm\_initialize\_alarm, 229
  - gm\_set\_alarm, 229
- gm\_set\_enable\_nack\_down
  - gm.h, 94
  - gm\_set\_enable\_nack\_down.c, 231
- gm\_set\_enable\_nack\_down.c, 231
  - gm\_set\_enable\_nack\_down, 231
- gm\_simple\_example.h, 232
- gm\_sleep
  - gm.h, 93
  - gm\_sleep.c, 233
- gm\_sleep.c, 233
  - gm\_sleep, 233
- gm\_srand
  - gm.h, 84
  - gm\_rand.c, 210
- gm\_status
  - gm.h, 53
- gm\_status\_t
  - gm.h, 53
- GM\_STILL\_SHUTTING\_DOWN
  - gm.h, 54
- gm\_strcmp
  - gm.h, 72
  - gm\_strcmp.c, 234
- gm\_strcmp.c, 234
  - gm\_strcmp, 234
- gm\_strdup
  - gm.h, 96
  - gm\_strdup.c, 235
- gm\_strdup.c, 235
  - gm\_strdup, 235
- gm\_strerror
  - gm.h, 94
  - gm\_strerror.c, 236
- gm\_strerror.c, 236
  - gm\_strerror, 236
- gm\_strlen
  - gm.h, 71
  - gm\_strlen.c, 237
- gm\_strlen.c, 237
  - gm\_strlen, 237
- gm\_strncasecmp
  - gm.h, 73
  - gm\_strncasecmp.c, 238
- gm\_strncasecmp.c, 238
  - gm\_strncasecmp, 238
- gm\_strncmp
  - gm.h, 72
  - gm\_strncmp.c, 239
- gm\_strncmp.c, 239
  - gm\_strncmp, 239
- gm\_strncpy
  - gm.h, 71
  - gm\_strncpy.c, 240
- gm\_strncpy.c, 240
  - gm\_strncpy, 240
- GM\_STRUCT\_CONTAINING
  - gm.h, 52
- GM\_SUCCESS
  - gm.h, 53
- gm\_ticks
  - gm.h, 73
  - gm\_ticks.c, 241
- gm\_ticks.c, 241
  - gm\_ticks, 241
- GM\_TIMED\_OUT
  - gm.h, 55
- GM\_TRY\_AGAIN

- gm.h, 53
- gm\_u16\_n\_t, 34
- gm\_u32\_n\_t, 35
- gm\_u64\_n\_t, 36
- gm\_u8\_n\_t, 37
- GM\_UC\_ERROR
  - gm.h, 55
- GM\_UNATTACHED
  - gm.h, 54
- gm\_unique\_id
  - gm.h, 74
  - gm\_unique\_id.c, 242
- gm\_unique\_id.c, 242
  - gm\_unique\_id, 242
- gm\_unique\_id\_to\_node\_id
  - gm.h, 74
  - gm\_unique\_id\_to\_node\_id.c, 243
- gm\_unique\_id\_to\_node\_id.c, 243
  - gm\_unique\_id\_to\_node\_id, 243
- gm\_unknown
  - gm.h, 74
  - gm\_unknown.c, 245
- gm\_unknown.c, 245
  - gm\_unknown, 245
- gm\_unmark
  - gm.h, 98
  - gm\_mark.c, 175
- gm\_unmark\_all
  - gm.h, 98
  - gm\_mark.c, 177
- gm\_unregister\_buffer
  - gm.h, 75
  - gm\_debug\_buffers.c, 121
- GM\_UNSUPPORTED\_DEVICE
  - gm.h, 54
- GM\_UNTRANSLATED\_SYSTEM\_ERROR
  - gm.h, 54
- gm\_up\_n\_t, 38
- GM\_USER\_ERROR
  - gm.h, 55
- GM\_YP\_NO\_MATCH
  - gm.h, 55
- gm\_zone, 39
  - gm\_zone.c, 246
  - GM\_AREA\_FOR\_PTR, 247
  - gm\_zone\_addr\_in\_zone, 249
  - gm\_zone\_area\_t, 247
  - gm\_zone\_calloc, 249
  - gm\_zone\_create\_zone, 247
  - gm\_zone\_destroy\_zone, 248
  - gm\_zone\_free, 248
  - gm\_zone\_malloc, 249
  - gm\_zone\_t, 247
- gm\_zone\_addr\_in\_zone
  - gm.h, 89
  - gm\_zone.c, 249
- gm\_zone\_area, 40
- gm\_zone\_area\_t
  - gm\_zone.c, 247
- gm\_zone\_calloc
  - gm.h, 89
  - gm\_zone.c, 249
- gm\_zone\_create\_zone
  - gm.h, 87
  - gm\_zone.c, 247
- gm\_zone\_destroy\_zone
  - gm.h, 88
  - gm\_zone.c, 248
- gm\_zone\_free
  - gm.h, 88
  - gm\_zone.c, 248
- gm\_zone\_malloc
  - gm.h, 88
  - gm\_zone.c, 249
- gm\_zone\_t
  - gm\_zone.c, 247
- hash\_entry, 41
- preallocated\_record\_chunk, 42
- segment\_list
  - gm\_lookaside, 20