



DOCA Driver APIs 2.6.0

Reference Manual

Table of Contents

Chapter 1. Change Log.....	1
Chapter 2. Modules.....	2
2.1. Host.....	2
flexio_affinity.....	3
flexio_app_attr.....	3
flexio_cmdq_attr.....	3
flexio_cq_attr.....	3
flexio_event_handler_attr.....	3
flexio_heap_mem_info.....	3
flexio_mkey_attr.....	3
flexio_msg_stream_attr_t.....	3
flexio_outbox_attr.....	3
flexio_process_attr.....	3
flexio_qmem.....	3
flexio_qp_attr.....	3
flexio_qp_attr_opt_param_mask.....	3
flexio_wq_attr.....	3
flexio_wq_sq_attr.....	3
flexio_affinity_type.....	3
flexio_cmdq_state.....	4
flexio_cq_period_mode.....	4
flexio_cqe_comp_type.....	4
flexio_err_status.....	4
flexio_log_lvl_t.....	4
flexio_memtype.....	5
flexio_msg_dev_sync_mode.....	5
flexio_qp_op_types.....	5
flexio_qp_qpc_mtu.....	5
flexio_qp_state.....	6
flexio_qp_transport_type.....	6
flexio_status.....	6
flexio_tracer_transport.....	6
flexio_func_arg_pack_fn_t.....	7
flexio_func_t.....	7
flexio_uar_device_id.....	7

flexio_uintptr_t.....	7
flexio_app_create.....	7
flexio_app_destroy.....	8
flexio_app_get_elf.....	8
flexio_app_get_elf_size.....	9
flexio_app_get_list.....	9
flexio_app_get_name.....	9
flexio_app_list_free.....	10
flexio_buf_dev_alloc.....	10
flexio_buf_dev_free.....	11
flexio_buf_dev_memset.....	11
flexio_cmdq_create.....	12
flexio_cmdq_destroy.....	12
flexio_cmdq_is_empty.....	13
flexio_cmdq_state_running.....	13
flexio_cmdq_task_add.....	13
flexio_copy_from_host.....	14
flexio_coredump_create.....	15
flexio_cq_create.....	15
flexio_cq_destroy.....	16
flexio_cq_get_cq_num.....	16
flexio_cq_modify_moderation.....	17
flexio_cq_query_moderation.....	17
flexio_crash_data.....	18
flexio_device_mkey_create.....	18
flexio_device_mkey_destroy.....	19
flexio_err_handler_fd.....	19
flexio_err_status_get.....	20
flexio_event_handler_create.....	20
flexio_event_handler_destroy.....	21
flexio_event_handler_get_id.....	21
flexio_event_handler_get_obj_id.....	21
flexio_event_handler_get_thread.....	22
flexio_event_handler_run.....	22
flexio_func_get_register_info.....	23
flexio_func_pup_register.....	24
flexio_func_register.....	25
flexio_host2dev_memcpy.....	25

flexio_log_dev_destroy.....	26
flexio_log_dev_flush.....	26
flexio_log_dev_init.....	27
flexio_log_lvl_set.....	28
flexio_mkey_get_id.....	28
flexio_msg_stream_create.....	28
flexio_msg_stream_destroy.....	29
flexio_msg_stream_flush.....	30
flexio_msg_stream_get_id.....	30
flexio_msg_stream_level_set.....	31
flexio_outbox_create.....	31
flexio_outbox_destroy.....	32
flexio_outbox_get_id.....	32
flexio_outbox_get_uar.....	32
flexio_process_call.....	33
flexio_process_create.....	33
flexio_process_destroy.....	34
flexio_process_error_handler_set.....	34
flexio_process_get_dumem_id.....	35
flexio_process_get_pd.....	35
flexio_process_get_uar.....	35
flexio_process_mem_info_get.....	36
flexio_process_udbg_token_get.....	36
flexio_qp_create.....	37
flexio_qp_destroy.....	37
flexio_qp_get_qp_num.....	38
flexio_qp_modify.....	38
flexio_qp_state_get.....	38
flexio_rmp_create.....	39
flexio_rmp_destroy.....	39
flexio_rmp_get_wq_num.....	40
flexio_rq_create.....	40
flexio_rq_destroy.....	41
flexio_rq_get_tir.....	41
flexio_rq_get_wq_num.....	41
flexio_rq_set_err_state.....	42
flexio_sq_create.....	42
flexio_sq_destroy.....	43

flexio_sq_get_wq_num.....	43
flexio_uar_create.....	43
flexio_uar_destroy.....	44
flexio_uar_extend.....	44
flexio_uar_get_extended_id.....	45
flexio_uar_get_id.....	45
flexio_window_create.....	45
flexio_window_destroy.....	46
flexio_window_get_id.....	46
FLEXIO_MAX_NAME_LEN.....	46
2.2. Dev.....	46
flexio_dev_process_tracer_ctx.....	47
flexio_tracer_msg.....	47
spinlock_s.....	47
cq_ce_mode.....	47
flexio_dev_status_t.....	47
flexio_dev_arg_unpack_func_t.....	47
flexio_dev_async_rpc_handler_t.....	47
flexio_dev_event_handler_t.....	48
flexio_dev_rpc_handler_t.....	48
flexio_uar_device_id.....	48
__packed__.....	48
g_dev_p_tracer_ctx.....	48
flexio_dev_get_pcc_table_base.....	48
flexio_dev_get_thread_ctx.....	49
flexio_dev_get_thread_id.....	49
flexio_dev_get_thread_local_storage.....	50
flexio_dev_msg.....	50
flexio_dev_nic_counters_config.....	50
flexio_dev_nic_counters_sample.....	51
flexio_dev_outbox_config.....	51
flexio_dev_outbox_config_fast.....	52
flexio_dev_outbox_config_uar_extension.....	52
flexio_dev_process_finish.....	53
flexio_dev_puts.....	53
flexio_dev_thread_finish.....	53
flexio_dev_thread_reschedule.....	54
flexio_dev_thread_retrigger.....	54

flexio_dev_trace_0.....	54
flexio_dev_trace_1.....	54
flexio_dev_trace_2.....	55
flexio_dev_trace_3.....	55
flexio_dev_trace_4.....	55
flexio_dev_trace_5.....	56
flexio_dev_trace_6.....	56
flexio_dev_tracer_flush.....	57
flexio_dev_tracer_notify_host.....	57
flexio_dev_window_config.....	57
flexio_dev_window_copy_from_host.....	58
flexio_dev_window_copy_to_host.....	59
flexio_dev_window_mkey_config.....	59
flexio_dev_window_ptr_acquire.....	60
flexio_dev_yield.....	60
flexio_dev_msg_broadcast.....	61
flexio_dev_msg_dflt.....	61
flexio_dev_print.....	61
spin_init.....	61
spin_lock.....	62
spin_trylock.....	62
spin_unlock.....	62
2.3. DevErr.....	62
flexio_dev_error_t.....	62
__attribute__.....	63
flexio_dev_get_and_rst_errno.....	63
flexio_dev_get_errno.....	63
flexio_dev_rst_errno.....	64
2.4. DevQueueAccess.....	64
flexio_ctrl_seg_t.....	64
flexio_dev_cc_db_next_act_t.....	64
flexio_dev_cc_ring_db.....	65
flexio_dev_cq_arm.....	65
flexio_dev_cqe_get_byte_cnt.....	66
flexio_dev_cqe_get_csum_ok.....	66
flexio_dev_cqe_get_err_synd.....	67
flexio_dev_cqe_get_opcode.....	67
flexio_dev_cqe_get_owner.....	67

flexio_dev_cqe_get_qpn.....	68
flexio_dev_cqe_get_user_index.....	68
flexio_dev_cqe_get_wqe_counter.....	68
flexio_dev_db_ctx_arm.....	69
flexio_dev_db_ctx_force_trigger.....	69
flexio_dev_dbr_cq_set_ci.....	70
flexio_dev_dbr_rq_inc_pi.....	70
flexio_dev_eq_update_ci.....	71
flexio_dev_eqe_get_cqn.....	71
flexio_dev_eqe_get_owner.....	72
flexio_dev_msix_send.....	72
flexio_dev_qp_sq_ring_db.....	73
flexio_dev_rwqe_get_addr.....	73
flexio_dev_swqe_seg_atomic_set.....	74
flexio_dev_swqe_seg_ctrl_set.....	74
flexio_dev_swqe_seg_eth_set.....	75
flexio_dev_swqe_seg_inline_data_set.....	76
flexio_dev_swqe_seg_mem_ptr_data_set.....	76
flexio_dev_swqe_seg_rdma_set.....	77
flexio_dev_swqe_seg_shared_receive_set.....	77
flexio_dev_swqe_seg_transpose_set.....	78
2.5. Change Log.....	78
Chapter 3. Data Structures.....	80
flexio_affinity.....	81
id.....	81
type.....	81
flexio_app_attr.....	81
app_bsize.....	81
app_name.....	81
app_ptr.....	81
app_sig_sec_name.....	81
flexio_cmdq_attr.....	81
batch_size.....	81
state.....	82
workers.....	82
flexio_cq_attr.....	82
always_armed.....	82
cc.....	82

cq_dbr_daddr.....	82
cq_max_count.....	82
cq_period.....	82
cq_period_mode.....	82
cq_ring_qmem.....	82
cqe_comp_type.....	83
element_type.....	83
emulated_eqn.....	83
log_cq_depth.....	83
no_arm.....	83
overrun_ignore.....	83
thread.....	83
uar_base_addr.....	83
uar_id.....	83
flexio_dev_cqe64.....	83
byte_cnt.....	84
csum_ok.....	84
err_syndrome.....	84
op_own.....	84
qpn.....	84
rsvd0.....	84
rsvd29.....	84
rsvd36.....	84
rsvd48.....	84
signature.....	84
srqn_uidx.....	84
wqe_counter.....	85
flexio_dev_eqe.....	85
cq_n.....	85
event_data.....	85
owner.....	85
rsvd00.....	85
rsvd00.....	85
rsvd02.....	85
rsvd3c.....	85
rsvd4.....	85
signature.....	86
sub_type.....	86

type.....	86
flexio_dev_process_tracer_ctx.....	86
tracer_ctx.....	86
flexio_dev_sqe_seg.....	86
atomic.....	86
ctrl.....	86
eth.....	87
inline_data.....	87
inline_send_data.....	87
mem_ptr_send_data.....	87
rdma.....	87
shared_receive.....	87
transpose.....	87
flexio_dev_wqe_atomic_seg.....	88
compare_data.....	88
swap_or_add_data.....	88
flexio_dev_wqe_ctrl_seg.....	88
general_id.....	88
idx_opcode.....	88
qpn_ds.....	88
signature_fm_ce_se.....	88
flexio_dev_wqe_eth_seg.....	89
cs_swp_flags.....	89
inline_hdr_bsz.....	89
inline_hdrs.....	89
mss.....	89
rsvd0.....	89
rsvd2.....	89
flexio_dev_wqe_inline_data_seg.....	89
inline_data.....	89
flexio_dev_wqe_inline_send_data_seg.....	90
byte_count.....	90
data_and_padding.....	90
flexio_dev_wqe_mem_ptr_send_data_seg.....	90
addr.....	90
byte_count.....	90
lkey.....	90
flexio_dev_wqe_rcv_data_seg.....	91

addr.....	91
byte_count.....	91
lkey.....	91
flexio_dev_wqe_rdma_seg.....	91
raddr.....	91
rkey.....	91
rsvd0.....	91
flexio_dev_wqe_shared_receive_seg.....	92
flexio_dev_wqe_transpose_seg.....	92
element_size.....	92
num_of_cols.....	92
num_of_rows.....	92
rsvd0.....	92
rsvd1.....	92
rsvd2.....	92
rsvd4.....	93
flexio_event_handler_attr.....	93
affinity.....	93
arg.....	93
continuable.....	93
host_stub_func.....	93
thread_local_storage_daddr.....	93
flexio_heap_mem_info.....	93
allocated.....	94
base_addr.....	94
requested.....	94
size.....	94
flexio_mkey_attr.....	94
access.....	94
daddr.....	94
len.....	94
pd.....	94
flexio_msg_stream_attr_t.....	95
data_bsize.....	95
level.....	95
mgmt_affinity.....	95
stream_name.....	95
sync_mode.....	95

tracer_mode.....	95
tracer_msg_formats.....	95
uar.....	96
flexio_outbox_attr.....	96
en_pcc.....	96
uar.....	96
flexio_process_attr.....	96
en_pcc.....	96
pd.....	96
flexio_qmem.....	96
daddr.....	96
humem_offset.....	97
memtype.....	97
umem_id.....	97
flexio_qp_attr.....	97
dest_mac.....	97
fl.....	97
gid_table_index.....	97
grh.....	97
isolate_vl_tc.....	97
log_rq_depth.....	97
log_rra_max.....	98
log_sq_depth.....	98
log_sra_max.....	98
min_rnr_nak_timer.....	98
next_rcv_psn.....	98
next_send_psn.....	98
next_state.....	98
no_sq.....	98
ops_flag.....	98
path_mtu.....	98
pd.....	98
qp_access_mask.....	99
qp_wq_buff_qmem.....	99
qp_wq_dbr_qmem.....	99
remote_qp_num.....	99
retry_count.....	99
rgid_or_rip.....	99

rlid.....	99
rmpqn.....	99
rq_cqn.....	99
rq_type.....	100
sq_cqn.....	100
transport_type.....	100
uar_id.....	100
udp_sport.....	100
user_index.....	100
vhca_port_num.....	100
flexio_qp_attr_opt_param_mask.....	100
min_rnr_nak_timer.....	100
qp_access_mask.....	101
flexio_tracer_msg.....	101
arg0.....	101
arg1.....	101
arg2.....	101
arg3.....	101
arg4.....	101
arg5.....	101
format_id.....	101
flexio_wq_attr.....	102
log_wq_depth.....	102
log_wq_stride.....	102
pd.....	102
rq.....	102
sq.....	102
uar_id.....	102
user_index.....	102
wq_dbr_qmem.....	102
wq_ring_qmem.....	102
flexio_wq_sq_attr.....	103
allow_multi_pkt_send_wqe.....	103
spinlock_s.....	103
locked.....	103
Chapter 4. Data Fields.....	104

Chapter 1. Change Log

This chapter lists changes in API that were introduced to the library.

1.3.0

- ▶ Field Groups, GPU Groups, and field watches created with a handle returned from `dcgmConnect()` are now cleaned up upon disconnect. `dcgmConnect_v2()` can be used to get the old behavior of objects persisting after disconnect.
- ▶ `dcgmConnect_v2()` was added as a method for specifying additional connection options when connecting to the host engine.
- ▶ `dcgmUnwatchFields()` was added as a method of unwatching fields that were previously watched with `dcgmWatchFields()`
- ▶ `dcgmActionValidate_v2()` was added to be able to pass more parameters to the DCGM GPU Diagnostic.
- ▶ `dcgmDiagResponse_t` was increased from v2 to v3. See `dcgmDiagResponse_v3` for details

1.2.3

- ▶ No API changes in this version.

1.1.1

- ▶ `dcgmGetAllSupportedDevices()` was added as a method to get DCGM-supported GPU IDs. `dcgmGetAllDevices()` can still be used to get all GPU IDs in the system.

1.0.0

- ▶ Initial Release.

Chapter 2. Modules

Here is a list of all modules:

- ▶ [Host](#)
- ▶ [Dev](#)
- ▶ [DevErr](#)
- ▶ [DevQueueAccess](#)
- ▶ [DevQueueTypes](#)

2.1. Host

Flex IO SDK host API for DPA programs. Mostly used for DPA resource management and invocation of DPA programs.

```
struct flexio_affinity
struct flexio_app_attr
struct flexio_cmdq_attr
struct flexio_cq_attr
struct flexio_event_handler_attr
struct flexio_heap_mem_info
struct flexio_mkey_attr
struct flexio_msg_stream_attr_t
struct flexio_outbox_attr
struct flexio_process_attr
struct flexio_qmem
struct flexio_qp_attr
struct flexio_qp_attr_opt_param_mask
struct flexio_wq_attr
struct flexio_wq_sq_attr
enum flexio_affinity_type
```

Flex IO thread affinity types.

Values

FLEXIO_AFFINITY_NONE = 0
FLEXIO_AFFINITY_STRICT
FLEXIO_AFFINITY_GROUP

enum flexio_cmdq_state

Flex IO command queue states.

Values

FLEXIO_CMDQ_STATE_PENDING = 0
FLEXIO_CMDQ_STATE_RUNNING = 1

enum flexio_cq_period_mode

Flex IO CQ CQE compression period modes.

Values

FLEXIO_CQ_PERIOD_MODE_EVENT = 0x0
FLEXIO_CQ_PERIOD_MODE_CQE = 0x1

enum flexio_cqe_comp_type

Flex IO CQ CQE compression modes.

Values

FLEXIO_CQE_COMP_NONE = 0x0
FLEXIO_CQE_COMP_BASIC = 0x1
FLEXIO_CQE_COMP_ENH = 0x2

enum flexio_err_status

Flex IO DEV error status.

Values

FLEXIO_DEV_NO_ERROR = 0
FLEXIO_DEV_FATAL_ERROR = 1
FLEXIO_DEV_USER_ERROR = 2

enum flexio_log_lvl_t

Flex IO SDK host logging levels

Values

FLEXIO_LOG_LVL_ERR = 0
FLEXIO_LOG_LVL_WARN = 1
FLEXIO_LOG_LVL_INFO = 2
FLEXIO_LOG_LVL_DBG = 3

enum flexio_memtype

Flex IO memory types.

Values

FLEXIO_MEMTYPE_DPA = 0
FLEXIO_MEMTYPE_HOST = 1

enum flexio_msg_dev_sync_mode

Flex IO device messaging synchronization modes.

Values

FLEXIO_LOG_DEV_SYNC_MODE_SYNC = 0
FLEXIO_LOG_DEV_SYNC_MODE_ASYNC = 1
FLEXIO_LOG_DEV_SYNC_MODE_BATCH = 2
FLEXIO_LOG_DEV_SYNC_MODE_TRACER = 3

enum flexio_qp_op_types

Flex IO QP operation types.

Values

FLEXIO_QP_WR_RDMA_WRITE = 0x4
FLEXIO_QP_WR_RDMA_READ = 0x8
FLEXIO_QP_WR_ATOMIC_CMP_AND_SWAP = 0x10

enum flexio_qp_qpc_mtu

Flex IO QP possible MTU values.

Values

FLEXIO_QP_QPC_MTU_BYTES_256 = 0x1
FLEXIO_QP_QPC_MTU_BYTES_512 = 0x2
FLEXIO_QP_QPC_MTU_BYTES_1K = 0x3
FLEXIO_QP_QPC_MTU_BYTES_2K = 0x4

FLEXIO_QP_QPC_MTU_BYTES_4K = 0x5

enum flexio_qp_state

Flex IO QP states.

Values

FLEXIO_QP_STATE_RST = 0x0

FLEXIO_QP_STATE_INIT = 0x1

FLEXIO_QP_STATE_RTR = 0x2

FLEXIO_QP_STATE_RTS = 0x3

FLEXIO_QP_STATE_ERR = 0x6

enum flexio_qp_transport_type

Flex IO QP states.

Values

FLEXIO_QPC_ST_RC = 0x0

FLEXIO_QPC_ST_UC = 0x1

FLEXIO_QPC_ST_UD = 0x2

FLEXIO_QPC_ST_XRC = 0x3

FLEXIO_QPC_ST_IBL2 = 0x4

FLEXIO_QPC_ST_DCI = 0x5

FLEXIO_QPC_ST_QP0 = 0x7

FLEXIO_QPC_ST_QP1 = 0x8

FLEXIO_QPC_ST_RAW_DATAGRAM = 0x9

FLEXIO_QPC_ST_REG_UMR = 0xc

FLEXIO_QPC_ST_DC_CNAK = 0x10

enum flexio_status

Flex IO API function return codes.

Values

FLEXIO_STATUS_SUCCESS = 0

FLEXIO_STATUS_FAILED = 1

FLEXIO_STATUS_TIMEOUT = 2

FLEXIO_STATUS_FATAL_ERR = 3

enum flexio_tracer_transport

Flex IO device messaging tracer transport modes.

Values

FLEXIO_TRACER_TRANSPORT_QP = 0

FLEXIO_TRACER_TRANSPORT_WINDOW = 1

typedef void (flexio_func_arg_pack_fn_t)

Callback function to pack the arguments for a function.

This function is called internally from the FlexIO runtime upon user making a call (e.g., flexio_process_call). It packs the arguments for a user function into the argument buffer provided in `argbuf`. The argument list can be arbitrarily long and is represented by `ap`. The correct usage of this function requires the caller to initialize the list using `va_start`.

typedef void (flexio_func_t)

Flex IO application function prototype.

typedef uint32_t flexio_uar_device_id

Flex IO UAR extension ID prototype.

typedef uint64_t flexio_uintptr_t

Flex IO address type.

flexio_status flexio_app_create (flexio_app_attr *fattr, flexio_app **app)

Create a container for a FlexIO App.

Parameters

fattr

- A pointer to the application attributes struct.

app

Returns

flexio status value.

Description

This function creates a named app with a given ELF buffer. It is called from within the constructor generated by the compiler.

flexio_status flexio_app_destroy (flexio_app *app)

Destroy a flexio app.

Parameters

app

- App that was created before.

Returns

flexio status value.

Description

This function destroys the state associated with the app and all registered functions. This function will free the internal elf buffer. It is called from within the destructor generated by the compiler.

flexio_status flexio_app_get_elf (flexio_app *app, uint64_t *bin_buff, size_t bin_size)

Retrieve ELF binary associated with application.

Parameters

app

- App that created before.

bin_buff

- Pointer to buffer to copy ELF binary.

bin_size

- Size of buffer pointed by bin_buff. If parameter is smaller than ELF binary size function will fail.

Returns

flexio status value.

Description

This function registers the function name, stub address with the runtime. Compiler calls this from within the constructor.

size_t flexio_app_get_elf_size (flexio_app *app)

Gets a Flex IO application size.

Parameters

app

- A pointer to a Flex IO application.

Returns

the application's size (bytes) or NULL on error.

flexio_status flexio_app_get_list (flexio_appapp_list, uint32_t *num_apps)

Get a list of FlexIO Apps that are available.

Parameters

app_list

- A list of apps that are available.

num_apps

Returns

flexio status value.

Description

This function returns a list of Flex IO apps that are loaded.

const char *flexio_app_get_name (flexio_app *app)

Gets a Flex IO application name.

Parameters

app

- A pointer to a Flex IO application.

Returns

the application's name or NULL on error.

flexio_status flexio_app_list_free (flexio_app **apps_list)

Free the list of flexio apps.

Returns

flexio status value.

Description

This function frees the list of apps obtained from `flexio_app_get_list`.

flexio_status flexio_buf_dev_alloc (flexio_process *process, size_t buff_bsize, flexio_uintptr_t *dest_daddr_p)

Allocates a buffer on Flex IO heap memory.

Parameters

process

- A pointer to the Flex IO process context.

buff_bsize

- The size of the buffer to allocate.

dest_daddr_p

- A pointer to the Flex IO address, where the buffer was allocated.

Returns

flexio status value.

Description

This function allocates a buffer with the requested size on the Flex IO heap memory. On success - sets `dest_daddr_p` to the start address of the allocated buffer. On Failure - sets `dest_daddr_p` to 0x0.

flexio_status flexio_buf_dev_free (flexio_process *process, flexio_uintptr_t daddr)

Deallocates Flex IO heap memory buffer.

Parameters

process

- A pointer to the Flex IO process context.

daddr

- A pointer to an address of allocated memory on the Flex IO heap. Zero value is valid argument.

Returns

flexio status value.

Description

This function frees Flex IO heap memory buffer by address.

flexio_status flexio_buf_dev_memset (flexio_process *process, int value, size_t buff_bsize, flexio_uintptr_t dest_daddr)

Sets DPA heap memory buffer to a given value.

Parameters

process

- A pointer to the Flex IO process context.

value

- A value to set the DPA heap memory buffer to.

buff_bsize

- The size of the Flex IO heap memory buffer.

dest_daddr

- Flex IO heap memory buffer address to set.

Returns

flexio status value.

```
flexio_status flexio_cmdq_create (flexio_process
*process, flexio_cmdq_attr *fattr, flexio_cmdq
**cmdq)
```

Create asynchronous rpc command queue.

Parameters

process

- A pointer to the process context.

fattr

- A pointer to the command queue attributes struct.

cmdq

- A pointer to the created command queue context pointer.

Returns

flexio status value.

Description

This function creates the asynchronous rpc command queue infrastructure allowing background tasks execution.

```
flexio_status flexio_cmdq_destroy (flexio_cmdq
*cmdq)
```

Destroy the command queue infrastructure.

Parameters

cmdq

- A pointer to the command queue context.

Returns

flexio status value.

Description

This function destroy the command queue infrastructure and release all its resources.

flexio_cmdq_is_empty (flexio_cmdq *cmdq)

Check if command queue is empty.

Parameters

cmdq

- A pointer to the command queue context.

Returns

boolean.

Description

This function checks if the command queue is empty and all jobs up to this point were performed.

flexio_status flexio_cmdq_state_running (flexio_cmdq *cmdq)

Move command queue to running state.

Parameters

cmdq

- A pointer to the command queue context.

Returns

flexio status value.

Description

This function moves the command queue to running state in the case the queue was create in pending state. Otherwise has no affect.

flexio_status flexio_cmdq_task_add (flexio_cmdq *cmdq, flexio_func_t *host_func, uint64_t arg)

Add a task to the asynchronous rpc command queue.

Parameters

cmdq

- A pointer to the command queue context.

host_func

- host stub function for DPA function to execute.

arg

- user argument to function.

Returns

flexio status value.

Description

This function adds a task to the asynchronous rpc command queue to be executed by DPA in background. allowing background jobs execution.

```
flexio_status flexio_copy_from_host  
(flexio_process *process, void *src_haddr, size_t  
buff_bsize, flexio_uintptr_t *dest_daddr_p)
```

Copy from host memory to Flex IO heap memory buffer.

Parameters**process**

- A pointer to the Flex IO process context.

src_haddr

- An address of the buffer on the host memory.

buff_bsize

- The size of the buffer to copy.

dest_daddr_p

- A pointer to the Flex IO address, where the buffer was copied to.

Returns

flexio status value.

Description

This function copies data from a buffer on the host memory to the Flex IO memory. The function allocates memory on the device heap which dest_address points to. It is the caller responsibility to deallocate this memory when it is no longer used.

flexio_status flexio_coredump_create (flexio_process *process, const char *outfile)

Create a DPA core dump of the process.

Parameters

process

- A pointer to a flexio_process

outfile

- pathname to write ELF formatted core dump data too. If NULL - filename will be generated in form flexio_dev.NNN.core, where NNN is the process id. If outfile is not NULL - suffix .NNN.core will be added. If outfile starts from slash (/pathname) - it will be passed with suffix described above to fopen() otherwise outfile will be created in the current directory or (if failed) in /tmp directory

Returns

flexio status value.

Description

This function creates a core dump image of a process and all it's threads, and is intended to be used after a fatal error or abnormal termination to allow the user to debug DPA application code.

There must be sufficient free memory to allocate 2-3 times the maximum core file size for intermediate processing before the elf file is written.

Memory windows that may be referenced by DPA code are **not** dumped by this code and must be handled separately if the data is desired.

flexio_status flexio_cq_create (flexio_process *process, ibv_context *ibv_ctx, const flexio_cq_attr *fattr, flexio_cq **cq)

Creates a Flex IO CQ.

Parameters

process

- A pointer to the Flex IO process.

ibv_ctx

- A pointer to an IBV device context (might be different than process'). If NULL - process' will be used.

fattr

- A pointer to the CQ attributes struct.

cq

- A pointer to the created CQ context pointer.

Returns

flexio status value.

Description

This function creates a Flex IO CQ.

flexio_status flexio_cq_destroy (flexio_cq *cq)

Destroys a Flex IO CQ.

Parameters**cq**

- A pointer to a CQ context.

Returns

flexio status value.

Description

This function destroys a Flex IO CQ.

uint32_t flexio_cq_get_cq_num (flexio_cq *cq)

Gets the Flex IO CQ number.

Parameters**cq**

- A pointer to a Flex IO CQ.

Returns

the CQ number or UINT32_MAX on error.

flexio_status flexio_cq_modify_moderation
 (flexio_cq *cq, uint16_t max_count, uint16_t period,
 uint16_t mode)

Modifies a Flex IO CQ moderation configuration.

Parameters

cq

- A pointer to a CQ context.

max_count

- CQ moderation max count value.

period

- CQ moderation period value.

mode

- CQ moderation mode value.

Returns

flexio status value.

flexio_status flexio_cq_query_moderation
 (flexio_cq *cq, uint16_t *max_count, uint16_t
 *period, uint16_t *mode)

Queries a Flex IO CQ moderation configuration.

Parameters

cq

- A pointer to a CQ context.

max_count

- A pointer to the CQ moderation max count value.

period

- A pointer to the CQ moderation period value.

mode

- A pointer to the CQ moderation mode value.

Returns

flexio status value.

flexio_status flexio_crash_data (flexio_process *process, const char *outfile)

Provide crash info in textual form.

Parameters

process

- A pointer to a flexio_process

outfile

- pathname to write ELF formatted core dump data too. If NULL - filename will be generated in form flexio_dev.NNN.crash, where NNN is the process id. If outfile is not NULL - suffix .NNN.crash will be added. If outfile starts from slash (/pathname) - it will be passed with suffix described above to fopen() otherwise outfile will be created in the current directory or (if failed) in /tmp directory

Returns

flexio status value.

Description

This function displays useful crash info in textual form. Info will be printed on console and duplicated to outfile

flexio_status flexio_device_mkey_create (flexio_process *process, flexio_mkey_attr *fattr, flexio_mkey **mkey)

Creates an Mkey to the process device UMEM.

Parameters

process

- A pointer to the Flex IO process context.

fattr

- A pointer to a Flex IO MKey attribute struct.

mkey

- A pointer to a pointer to the created MKey struct.

Returns

flexio status value.

Description

This function creates an MKey over the provided PD for the provided process device UMEM. The mkey_id will point to the field in the containing flexio_mkey object.

flexio_status flexio_device_mkey_destroy (flexio_mkey *mkey)

destroys an MKey object containing the given ID

Parameters

mkey

- A pointer to the Flex IO MKey to destroy. NULL is a valid value.

Returns

flexio status value.

Description

This function destroys an Mkey object containing the given ID.

flexio_err_handler_fd (flexio_process *process)

Get file descriptor for error handler.

Parameters

process

- A pointer to the Flex IO process.

Returns

- file descriptor.

Description

User should get fd in order to monitor for nonrecoverable errors

User can poll all created processes, using select/poll/epoll functions family.

flexio_err_status flexio_err_status_get (flexio_process *process)

Check if unrecoverable error occurred.

Parameters

process

- A pointer to the Flex IO process. NULL is a valid value.

Returns

- nonzero value if error happen.

Description

It is suggested to check error status after every negotiation with DPA and periodically later.

flexio_status flexio_event_handler_create (flexio_process *process, flexio_event_handler_attr *fattr, flexio_event_handler **event_handler_ptr)

Creates a Flex IO event handler.

Parameters

process

- A pointer to the Flex IO process.

fattr

- A pointer to the event handler attributes struct.

event_handler_ptr

- A pointer to the created event handler context pointer.

Returns

flexio status value.

Description

This function creates a Flex IO event handler for an existing Flex IO process.

flexio_status flexio_event_handler_destroy (flexio_event_handler *event_handler)

Destroys a Flex IO event handler.

Parameters

event_handler

- A pointer to an event handler context.

Returns

flexio status value.

Description

This function destroys a Flex IO event handler.

uint32_t flexio_event_handler_get_id (flexio_event_handler *event_handler)

Gets the ID from a Flex IO event handler's thread metadata.

Parameters

event_handler

- A pointer to a Flex IO event handler.

Returns

the event handler's thread ID or UINT32_MAX on error.

uint32_t flexio_event_handler_get_obj_id (flexio_event_handler *event_handler)

Gets the object ID of a Flex IO event handler.

Parameters

event_handler

- A pointer to a Flex IO event handler.

Returns

the event handler's thread object ID or UINT32_MAX on error.

```
flexio_thread *flexio_event_handler_get_thread  
(flexio_event_handler *event_handler)
```

Gets a Flex IO thread object from a Flex IO event handler.

Parameters

event_handler

- A pointer to a Flex IO event handler.

Returns

the event handler's thread or NULL on error.

```
flexio_status flexio_event_handler_run  
(flexio_event_handler *event_handler, uint64_t  
user_arg)
```

Run a Flex IO event handler.

Parameters

event_handler

- A pointer to an event handler context.

user_arg

- A 64 bit argument for the event handler's thread.

Returns

flexio status value.

Description

This function makes a Flex IO event handler start running.

```
flexio_status flexio_func_get_register_info
(flexio_app *app, flexio_func_t
*host_stub_func_addr, uint32_t *pup, char
*dev_func_name, char *dev_unpack_func_name,
size_t func_name_size, size_t *argbuf_size,
flexio_func_arg_pack_fn_t **host_pack_func,
flexio_uintptr_t *dev_func_addr, flexio_uintptr_t
*dev_unpack_func_addr)
```

Obtain info for previously registered function.

Parameters

app

- FlexIO app.

host_stub_func_addr

- Known host stub func addr.

pup

- Whether function has been registered with pack/unpack support (0: No, 1:Yes).

dev_func_name

- Name of device function.

dev_unpack_func_name

- Name of unpack routine on device, NA if pup == 0.

func_name_size

- Size of function name len allocated.

argbuf_size

- Size of argument buffer, NA if pup == 0.

host_pack_func

- Function pointer to host packing routine, NA if pup == 0.

dev_func_addr

- address of device function.

dev_unpack_func_addr

- address of device unpack function.

Returns

flexio status value.

Description

This function is used to obtain info about a previously registered function. It is used to compose higher-level libraries on top of DPACC / FlexIO interface. It is not intended to be used directly by the user.

The caller must ensure that the string pointers have been allocated and are at least ``FLEXIO_MAX_NAME_LEN + 1`` long to ensure that the call doesn't fail to copy full function name.

```
flexio_status flexio_func_pup_register (flexio_app
*app, const char *dev_func_name, const
char *dev_unpack_func_name, flexio_func_t
*host_stub_func_addr, size_t argbuf_size,
flexio_func_arg_pack_fn_t *host_pack_func)
```

Register a function name at application start.

Parameters

app

- App that created before.

dev_func_name

- The device function name (entry point). Length of name should be up to `FLEXIO_MAX_NAME_LEN` bytes.

dev_unpack_func_name

- The device wrapper function that unpacks the argument buffer. Length of name should be up to `FLEXIO_MAX_NAME_LEN` bytes.

host_stub_func_addr

- The host stub function that is used by the application to reference the device function.

argbuf_size

- Size of the argument buffer required by this function.

host_pack_func

- Host callback function that packs the arguments.

Returns

flexio status value.

Description

This function registers the function name, stub address with the runtime. It is called from within the constructor generated by the compiler.

```
flexio_status flexio_func_register (flexio_app
*app, const char *dev_func_name, flexio_func_t
**out_func)
```

Register a function to be used later.

Parameters

app

- previously created flexio app.

dev_func_name

- name of flexio function on device that will be called. Length of name should be up to FLEXIO_MAX_NAME_LEN bytes.

out_func

- opaque handle to use with [flexio_process_call\(\)](#), [flexio_event_handler_create\(\)](#), ...

Returns

flexio status value.

Description

This function is intended to be called directly by user in the situation where they don't desire pack/unpack support that is typically done by the compiler interface.

It is the user's responsibility to ensure that a function was annotated for event handler with `__dpa_global__`. The runtime will not provide any type checking. A mismatched call will result in undefined behavior.

```
flexio_status flexio_host2dev_memcpy
(flexio_process *process, void *src_haddr, size_t
buff_bsize, flexio_uintptr_t dest_daddr)
```

Copy from host memory to a pre-allocated Flex IO heap memory buffer.

Parameters

process

- A pointer to the Flex IO process context.

src_haddr

- An address of the buffer on the host memory.

buff_bsize

- The size of the buffer to copy.

dest_daddr

- Flex IO heap memory buffer address to copy to.

Returns

flexio status value.

Description

This function copies data from a buffer on the host memory to a buffer on the Flex IO heap memory.

flexio_status flexio_log_dev_destroy (flexio_process *process)

Destroys a flexio device messaging default stream environment.

Parameters**process**

- A pointer to the Flex IO process.

Returns

flexio status value.

Description

This function destroys and releases all resources, allocated for process messaging needs, which were allocated by [flexio_log_dev_init\(\)](#) in purpose of serving the default stream.

flexio_status flexio_log_dev_flush (flexio_process *process)

Flush the default msg stream's buffer in case of asynchronous messaging mode.

Parameters**process**

- A pointer to the Flex IO process.

Returns

flexio status value.

Description

All data from the default msg stream buffer will be flushed to the file defined in [flexio_log_dev_init\(\)](#).

In case of synchronous device messaging this functions does nothing. This function allocates resources to support messaging from Flex IO to HOST.

```
flexio_status flexio_log_dev_init (flexio_process
*process, flexio_msg_stream_attr_t *stream_fattr,
FILE *out, pthread_t *ppthread)
```

Create environment to support messages output from DPA.

Parameters

process

- A pointer to the Flex IO process.

stream_fattr

- A pointer to the messaging attributes struct.

out

- file to save data from Flex IO. Use stdout if you want receive data on HOST's console

ppthread

- A pointer to receive pthread ID of created thread. May be NULL if user doesn't need it.

Returns

flexio status value.

Description

This function allocates resources to support messages output from Flex IO to HOST. It can only allocate and create the default stream.

Device messaging works in the following modes: synchronous or asynchronous. Under synchronous mode, a dedicated thread starts to receive data and outputs it immediately. When asynchronous mode is in operation, all message stream buffers will be flushed by [flexio_log_dev_flush\(\)](#). Buffer can be overrun.

This function doesn't have a "destroy" procedure. All messaging infrastructure will be closed and the resources will be released using the [flexio_process_destroy\(\)](#) function.

enum flexio_log_lvl flexio_log_lvl_set (enum flexio_log_lvl lvl)

Sets host SDK logging level.

Parameters

lvl

- logging level to set. All entries with this or higher priority level will be printed.

Returns

the previous host logging level.

Description

This function sets the host logging level. Changing the logging level may change the visibility of some logging entries in the SDK code.

uint32_t flexio_mkey_get_id (flexio_mkey *mkey)

Gets the Flex IO MKey ID.

Parameters

mkey

- A pointer to a Flex IO MKey.

Returns

the Flex IO mkey ID or UINT32_MAX on error.

flexio_status flexio_msg_stream_create (flexio_process *process, flexio_msg_stream_attr_t *stream_fattr, FILE *out, pthread_t *ppthread, flexio_msg_stream **stream)

Create a Flex IO msg stream that can contain output messages sent from the DPA.

Parameters

process

- A pointer to the Flex IO process.

stream_fattr

- A pointer to the flexio_msg_stream attributes struct.

out

- file to save data from Flex IO. Use stdout if you want receive data on HOST's console

pthead

- A pointer to receive pthread ID of created thread. May be NULL if user doesn't need it.

stream

- A pointer to the created stream context pointer.

Returns

flexio status value.

Description

This function can create a flexio_msg_stream that could have device messages directed to it. Directing messages from the device to the host, could be done to any and all open streams, including the default stream.

The function creates the same resources created in flexio_log_dev_init for any new stream. It can also create the default stream. It creates it with the FLEXIO_MSG_DEV_INFO stream level, and that could be modified using flexio_msg_stream_level_set.

flexio_status flexio_msg_stream_destroy (flexio_msg_stream *stream)

Destroys a Flex IO msg stream.

Parameters**stream**

- A pointer to the stream context.

Returns

flexio status value.

Description

This function destroys any Flex IO msg stream.

flexio_status flexio_msg_stream_flush (flexio_msg_stream *stream)

Flush a msg stream's buffer in case of asynchronous messaging mode.

Parameters

stream

- A pointer to the Flex IO msg stream.

Returns

flexio status value.

Description

All data from the msg stream buffer will be flushed to the file defined in [flexio_msg_stream_create\(\)](#).

In case of synchronous device messaging this functions does nothing. This function allocates resources to support messaging from Flex IO to HOST.

flexio_msg_stream_get_id (flexio_msg_stream *stream)

Gets the Flex IO device message stream's ID (aka file descriptor).

Parameters

stream

- A pointer to a Flex IO message stream.

Returns

the stream_id or -1 in case of error.

Description

Using this function on a destroyed stream will result in unpredictable behavior.

flexio_status flexio_msg_stream_level_set (flexio_msg_stream *stream, flexio_msg_dev_level level)

Change the provided device message stream's level.

Parameters

stream

- A pointer to a Flex IO message stream.

level

- The new desired level, ranges between FLEXIO_MSG_DEV_NO_PRINT FLEXIO_MSG_DEV_DEBUG. FLEXIO_MSG_DEV_ALWAYS_PRINT cannot be used here.

Returns

flexio status value.

Description

The default stream's level cannot be altered. Note that modifying the stream's level while messages are being sent may result in missing or unwanted messages.

flexio_status flexio_outbox_create (flexio_process *process, flexio_outbox_attr *fattr, flexio_outbox **outbox)

Creates a Flex IO outbox.

Parameters

process

- A pointer to the Flex IO process.

fattr

- A pointer to the outbox attributes struct.

outbox

- A pointer to the created outbox context pointer.

Returns

flexio status value.

Description

This function Creates a Flex IO outbox for the given process.

flexio_status flexio_outbox_destroy (flexio_outbox *outbox)

Destroys a Flex IO outbox.

Parameters

outbox

- A pointer to a outbox context.

Returns

flexio status value.

Description

This function destroys a Flex IO outbox.

uint32_t flexio_outbox_get_id (flexio_outbox *outbox)

Gets the Flex IO outbox ID.

Parameters

outbox

- A pointer to a Flex IO outbox.

Returns

the Flex IO outbox ID or UINT32_MAX on error.

flexio_uar *flexio_outbox_get_uar (flexio_outbox *outbox)

Gets a Flex IO UAR object from a Flex IO outbox.

Parameters

outbox

- A pointer to a Flex IO outbox.

Returns

the Flex IO outbox UAR object or NULL on error.

```
flexio_status flexio_process_call (flexio_process
*process, flexio_func_t *host_func, uint64_t
*func_ret, ...)
```

Calls a Flex IO process.

Parameters

process

- A pointer to the Flex IO process to run.

host_func

- The host stub function that is used by the application to reference the device function.

func_ret

- A pointer to the ELF function return value.

Returns

flexio status value.

```
flexio_status flexio_process_create (ibv_context
*ibv_ctx, flexio_app *app, const flexio_process_attr
*process_attr, flexio_process **process_ptr)
```

Create a new Flex IO process.

Parameters

ibv_ctx

- A pointer to a device context.

app

- Device side application handle.

process_attr

- Optional, process attributes for create. Can be NULL.

process_ptr

- A pointer to the created process pointer.

Returns

flexio status value.

Description

This function creates a new Flex IO process with requested image.

flexio_status flexio_process_destroy (flexio_process *process)

Destroys a Flex IO process.

Parameters

process

- A pointer to a process. NULL is a valid value.

Returns

flexio status value.

Description

This function destroys a Flex IO process.

flexio_status flexio_process_error_handler_set (flexio_process *process, flexio_func_t *error_handler)

Set the Flexio process error handler.

Parameters

process

- A pointer to a process

error_handler

- The host stub function that is used as a reference to the error handler function.

Returns

flexio status value.

Description

This function sets the Flex IO process error handler. The error handler must be set after the process is created, and before the first thread is created. The function registered for error handler should be annotated with `__dpa_global__`.

uint32_t flexio_process_get_dumem_id (flexio_process *process)

Gets the Flex IO process DUMEM ID.

Parameters

process

- A pointer to a Flex IO process.

Returns

the Flex IO process DUMEM ID or UINT32_MAX on error.

ibv_pd *flexio_process_get_pd (flexio_process *process)

Gets a Flex IO IBV PD object from a Flex IO process.

Parameters

process

- A pointer to a Flex IO process.

Returns

the process's PD object or NULL on error.

flexio_uar *flexio_process_get_uar (flexio_process *process)

Gets a Flex IO UAR object from a Flex IO process.

Parameters

process

- A pointer to a Flex IO process.

Returns

the Flex IO process UAR object or NULL on error.

```
flexio_status flexio_process_mem_info_get (const  
flexio_process *process, flexio_heap_mem_info  
*info)
```

Get process memory info.

Parameters

process

- A pointer to the Flex IO process context.

info

- A pointer to [flexio_heap_mem_info](#) struct to fill info.

Returns

flexio status value.

Description

This function returns the process heap memory base address and its available size.

```
uint64_t flexio_process_udbg_token_get  
(flexio_process *process)
```

Get token for Flex IO process debug access.

Parameters

process

- A pointer to the Flex IO process context.

Returns

the requested token. Zero value means - User Debug access for the process is not allowed.

Description

This function returns the token, needed for user debug syscalls access.


```
flexio_status flexio_qp_create (flexio_process
*process, ibv_context *ibv_ctx, flexio_qp_attr
*qp_fattr, flexio_qp **qp_ptr)
```

Creates a Flex IO QP.

Parameters

process

- A pointer to the Flex IO process.

ibv_ctx

- A pointer to an IBV device context (might be different than process'). If NULL - process' will be used.

qp_fattr

- A pointer to the QP attributes struct.

qp_ptr

- A pointer to the created QP context pointer.

Returns

flexio status value.

Description

This function creates a Flex IO QP.

```
flexio_status flexio_qp_destroy (flexio_qp *qp)
```

Destroys a Flex IO QP.

Parameters

qp

- A pointer to the QP context.

Returns

flexio status value.

Description

This function destroys a Flex IO QP.

uint32_t flexio_qp_get_qp_num (flexio_qp *qp)

Gets the Flex IO QP number.

Parameters

qp

- A pointer to a Flex IO QP.

Returns

the QP number or UINT32_MAX on error.

flexio_status flexio_qp_modify (flexio_qp *qp, flexio_qp_attr *fattr, flexio_qp_attr_opt_param_mask *mask)

Modify Flex IO QP.

Parameters

qp

- A pointer to the QP context.

fattr

- A pointer to the QP attributes struct that will also define the QP connection.

mask

- A pointer to the optional QP attributes mask.

Returns

flexio status value.

Description

This function modifies Flex IO QP and transition it between states. At the end of the procedure Flex IO QP would have moved from it's current state to to next state, given in the fattr, if the move is a legal transition in the QP's state machine.

flexio_qp_state flexio_qp_state_get (flexio_qp *qp)

retrieve the device QP state.

Parameters

qp

- A pointer to a Flex IO QP.

Returns

enum flexio_qp_state.

Description

This function return the device QP state it is currently in.

```
flexio_status flexio_rmp_create (flexio_process
*process, ibv_context *ibv_ctx, const flexio_wq_attr
*fattr, flexio_rmp **flexio_rmp_ptr)
```

Creates a Flex IO RMP.

Parameters

process

- A pointer to the Flex IO process.

ibv_ctx

- A pointer to an IBV device context (might be different than process'). If NULL - process' will be used.

fattr

- A pointer to the WQ attributes struct.

flexio_rmp_ptr

- A pointer to the created RMP context pointer.

Returns

flexio status value.

Description

This function creates a Flex IO RMP.

```
flexio_status flexio_rmp_destroy (flexio_rmp
*flexio_rmp)
```

Destroys a Flex IO RMP.

Parameters

flexio_rmp

- A pointer to an RMP context.

Returns

flexio status value.

Description

This function destroys a Flex IO RMP.

uint32_t flexio_rmp_get_wq_num (flexio_rmp *rmp)

Gets the Flex IO RMP number.

Parameters

rmp

- A pointer to a Flex IO RMP.

Returns

the RQ number or UINT32_MAX on error.

flexio_status flexio_rq_create (flexio_process *process, ibv_context *ibv_ctx, uint32_t cq_num, const flexio_wq_attr *fattr, flexio_rq **flexio_rq_ptr)

Creates a Flex IO RQ.

Parameters

process

- A pointer to the Flex IO process.

ibv_ctx

- A pointer to an IBV device context (might be different than process'). If NULL - process' will be used.

cq_num

- A CQ number.

fattr

- A pointer to the RQ WQ attributes struct.

flexio_rq_ptr

- A pointer to the created RQ context pointer.

Returns

flexio status value.

Description

This function creates a Flex IO RQ.

flexio_status flexio_rq_destroy (flexio_rq *flexio_rq)

Destroys a Flex IO RQ.

Parameters

flexio_rq

- A pointer to an RQ context.

Returns

flexio status value.

Description

This function destroys a Flex IO RQ.

mlx5dv_devx_obj *flexio_rq_get_tir (flexio_rq *rq)

Gets the Flex IO RQ TIR object.

Parameters

rq

- A pointer to a Flex IO RQ.

Returns

the RQ TIR object or NULL on error.

uint32_t flexio_rq_get_wq_num (flexio_rq *rq)

Gets the Flex IO RQ number.

Parameters

rq

- A pointer to a Flex IO RQ.

Returns

the RQ number or UINT32_MAX on error.

flexio_status flexio_rq_set_err_state (flexio_rq *rq)

Sets a Flex IO RQ to error state.

Parameters

rq

- A pointer to the RQ context to move to error state.

Returns

flexio status value.

Description

This function sets a Flex IO RQ to error state.

flexio_status flexio_sq_create (flexio_process *process, ibv_context *ibv_ctx, uint32_t cq_num, const flexio_wq_attr *fattr, flexio_sq **flexio_sq_ptr)

Creates a Flex IO SQ.

Parameters

process

- A pointer to the Flex IO process.

ibv_ctx

- A pointer to an IBV device context (might be different than process'). If NULL - process' will be used.

cq_num

- A CQ number (can be Flex IO or host CQ).

fattr

- A pointer to the SQ attributes struct.

flexio_sq_ptr

Returns

flexio status value.

Description

This function creates a Flex IO SQ.

flexio_status flexio_sq_destroy (flexio_sq *flexio_sq)

Destroys a Flex IO SQ.

Returns

flexio status value.

Description

This function destroys a Flex IO SQ.

uint32_t flexio_sq_get_wq_num (flexio_sq *sq)

Gets the Flex IO SQ number.

Parameters

sq

- A pointer to a Flex IO SQ.

Returns

the SQ number or UINT32_MAX on error.

flexio_status flexio_uar_create (flexio_process *process, flexio_uar **flexio_uar)

Creates a Flex IO UAR object.

Parameters

process

- A pointer to the Flex IO process context.

flexio_uar

- A pointer to a pointer to the created Flex IO UAR struct.

Returns

flexio status value.

Description

This function creates a Flex IO UAR object.

flexio_status flexio_uar_destroy (flexio_uar *uar)

destroys a Flex IO UAR object

Parameters

uar

- A pointer to the Flex IO UAR to destroy.

Returns

flexio status value.

Description

This function destroys a Flex IO UAR object.

flexio_status flexio_uar_extend (flexio_uar *in_uar, ibv_context *to_extend, flexio_uar **extended)

Extend UAR to an ibv context.

Parameters

in_uar

- A pointer to the Flex IO uar.

to_extend

- A pointer to an IBV device context to be extended to.

extended

- A pointer to the UAR context pointer.

Returns

flexio status value.

Description

This function extend the UAR to an ibv context to allow handling its queues.

flexio_uar_device_id flexio_uar_get_extended_id (flexio_uar *uar)

Gets the Flex IO extended UAR ID.

Parameters

uar

- A pointer to a Flex IO extended UAR.

Returns

the Flex IO UAR extended ID or UINT32_MAX on error.

uint32_t flexio_uar_get_id (flexio_uar *uar)

Gets the Flex IO UAR ID.

Parameters

uar

- A pointer to a Flex IO UAR.

Returns

the Flex IO UAR ID or UINT32_MAX on error.

flexio_status flexio_window_create (flexio_process *process, ibv_pd *pd, flexio_window **window)

Creates a Flex IO window.

Parameters

process

- A pointer to the Flex IO process.

pd

- A pointer to a protection domain struct to the memory the window should access.

window

- A pointer to the created window context pointer.

Returns

flexio status value.

Description

This function Creates a Flex IO window for the given process.

flexio_status flexio_window_destroy (flexio_window *window)

Destroys a Flex IO window.

Parameters

window

- A pointer to a window context.

Returns

flexio status value.

Description

This function destroys a Flex IO window.

uint32_t flexio_window_get_id (flexio_window *window)

Gets the Flex IO window ID.

Parameters

window

- A pointer to a Flex IO window.

Returns

the Flex IO window ID or UINT32_MAX on error.

#define FLEXIO_MAX_NAME_LEN (256)

Maximum length of application and device function names

2.2. Dev

Flex IO SDK device API for DPA programs. Includes services for DPA programs.

Flex IO SDK message stream device API for DPA programs. Includes message stream services for DPA programs.

`struct flexio_dev_process_tracer_ctx`

`struct flexio_tracer_msg`

`struct spinlock_s`

`enum cq_ce_mode`

Flex IO dev CQ CQE creation modes.

Values

MLX5_CTRL_SEG_CE_CQE_ON_CQE_ERROR = 0x0

MLX5_CTRL_SEG_CE_CQE_ON_FIRST_CQE_ERROR = 0x1

MLX5_CTRL_SEG_CE_CQE_ALWAYS = 0x2

MLX5_CTRL_SEG_CE_CQE_AND_EQE = 0x3

`enum flexio_dev_status_t`

Return status of Flex IO dev API functions.

Values

FLEXIO_DEV_STATUS_SUCCESS = 0

FLEXIO_DEV_STATUS_FAILED = 1

`typedef uint64_t (flexio_dev_arg_unpack_func_t)`

Unpack the arguments and call the user function.

This callback function is used at runtime to unpack the arguments from the call on Host and then call the function on DPA. This function is called internally from flexio dev.

argbuf - Argument buffer that was written by Host. func - Function pointer to user function.

return uint64_t - result of the RPC function.

`typedef void (flexio_dev_async_rpc_handler_t)`

Asynchronous RPC handler callback function type.

Defines an RPC handler callback function.

arg - argument of the RPC function.

return void.

typedef void (flexio_dev_event_handler_t)

Event handler callback function type.

Defines an event handler callback function. On handler function end, need to call [flexio_dev_process_finish\(\)](#) instead of a regular return statement, in order to properly release resources back to the OS.

thread_arg - an argument for the executing thread.

return void.

typedef uint64_t (flexio_dev_rpc_handler_t)

RPC handler callback function type.

Defines an RPC handler for most useful callback function.

arg - argument of the RPC function.

return uint64_t - result of the RPC function.

typedef uint32_t flexio_uar_device_id

Flex IO UAR extension ID prototype.

struct flexio_tracer_msg ::__packed__

Describes Flex IO trace message. This struct is used to communicate the tracer raw data from device to host.

flexio_dev_process_tracer_ctx *::g_dev_p_tracer_ctx

Global process tracer context struct instance.

uint64_t flexio_dev_get_pcc_table_base (uint16_t gvmi)

get programable congestion control table base address

Parameters

gvmi

- PCC table GVMI.

Returns

PCC table base address for the given GVMI.

Description

This function gets the programable congestion control table base address.

flexio_dev_get_thread_ctx (flexio_dev_thread_ctx **dtctx)

Request thread context.

Parameters

dtctx

- A pointer to a pointer of flexio_dev_thread_ctx structure.

Returns

0 on success negative value on failure.

Description

This function requests the thread context. Should be called for every start of thread.

uint32_t flexio_dev_get_thread_id (flexio_dev_thread_ctx *dtctx)

Get thread ID from thread context.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

Returns

thread ID value.

Description

This function queries a thread context for its thread ID (from thread metadata).

flexio_uintptr_t flexio_dev_get_thread_local_storage (flexio_dev_thread_ctx *dtctx)

Get thread local storage address from thread context.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

Returns

thread local storage value.

Description

This function queries a thread context for its thread local storage (from thread metadata).

flexio_dev_msg (int stream_id, flexio_msg_dev_level level, const char *format, ...)

Creates message entry and outputs from the device to the host side. Same as a regular printf but with protection from simultaneous print from different threads.

Returns

- same as from regular printf.

Description

[in] stream_id - the relevant msg stream, created and passed from the host. [in] level - messaging level. [in] format, ... - same as for regular printf.

flexio_dev_nic_counters_config (uint32_t *values, uint32_t *counter_ids, uint32_t num_coutners)

Prepare a list of counters to read.

Parameters

values

- buffer to store counters values (32b) read by [flexio_dev_nic_counters_sample\(\)](#).

counter_ids

- An array of counter ids.

num_counters**Description**

The list is stored in kernel memory. A single counters config per process is supported. Note that arrays memory must be defined in global or heap memory only.

flexio_dev_nic_counters_sample (void)

Sample counters according to the prior configuration call.

Description

Sample counter_ids, num_counters and values buffer provided in the last successful call to flexio_dev_config_nic_counters(). This call ensures fastest sampling on a pre-checked counter ids and buffers.

flexio_dev_status_t flexio_dev_outbox_config (flexio_dev_thread_ctx *dtctx, uint16_t outbox_config_id)

Config thread outbox object.

Parameters**dtctx**

- A pointer to flexio_dev_thread_ctx structure.

outbox_config_id

- The outbox object config id.

Returns

flexio_dev_status_t.

Description

This function updates the thread outbox object of the given thread context.

flexio_dev_outbox_config_fast (flexio_dev_thread_ctx *dtctx, uint16_t outbox_config_id)

Config thread outbox object without any checks.

Parameters

dtctx

- A pointer to flexio_dev_thread_ctx structure.

outbox_config_id

- The outbox object config id.

Description

This function updates the thread outbox object of the given thread context, but it doesn't check for correctness or redundancy (same ID as current configured).

flexio_dev_status_t flexio_dev_outbox_config_uar_extension (flexio_dev_thread_ctx *dtctx, flexio_uar_device_id device_id)

set extension ID for outbox

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

device_id

- The device ID.

Returns

flexio_dev_status_t.

Description

This function sets the GVMI for the outbox to operate on.

flexio_dev_process_finish (void)

Exit flexio process (no errors).

Description

This function releases resources back to OS and returns '0x40' in dpa_process_status. All threads for the current process will stop executing and no new threads will be able to trigger for this process. Threads state will NOT be changes to 'finished' (will remain as is).

flexio_dev_puts (flexio_dev_thread_ctx *dtctx, char *str)

Put a string to messaging queue.

Parameters

dtctx

- A pointer to a pointer of flexio_dev_thread_ctx structure.

str

- A pointer to string.

Returns

length of messaged string.

Description

This function puts a string to host's default stream messaging queue. This queue has been serviced by host application. Would have no effect, if the host application didn't configure device messaging stream environment. In order to initialize/configure device messaging environment - On HOST side - after flexio_process_create, a stream should be created, therefore flexio_msg_stream_create should be called, and the default stream should be created. On DEV side - before using flexio_dev_puts, the thread context is needed, therefore flexio_dev_get_thread_ctx should be called before.

flexio_dev_thread_finish (void)

Exit from a thread, mark it as finished.

Description

This function releases resources back to OS. The thread will be marked as finished so next DUAR will not trigger it.

flexio_dev_thread_reschedule (void)

Exit from a thread, leave process active.

Description

This function releases resources back to OS. For the next DUAR the thread will restart from the beginning.

flexio_dev_thread_retrigger (void)

Exit from a thread, and retrigger it.

Description

This function asks the OS to retrigger the thread. The thread will not wait for the next DUAR to be triggered but will be triggered immediately.

flexio_dev_trace_0 (uint8_t tracer_id, flexio_msg_dev_level level, int format_id)

Creates trace message entry with no arguments.

Description

Using the trace mechanism for fast logging. Call the appropriate function according to number of needed arguments.

[in] tracer_id - the relevant msg stream id. [in] level - messaging level. [in] format_id -the template format id to print message accordingly.

flexio_dev_trace_1 (uint8_t tracer_id, flexio_msg_dev_level level, int format_id, uint64_t arg0)

Creates trace message entry with 1 arguments.

Description

Using the trace mechanism for fast logging. Call the appropriate function according to number of needed arguments.

[in] tracer_id - the relevant msg stream id. [in] level - messaging level. [in] format_id -the template format id to print message accordingly. [in] arg0 - argument #0 to format into the template.

`flexio_dev_trace_2 (uint8_t tracer_id, flexio_msg_dev_level level, int format_id, uint64_t arg0, uint64_t arg1)`

Creates trace message entry with 2 arguments.

Description

Using the trace mechanism for fast logging. Call the appropriate function according to number of needed arguments.

[in] tracer_id - the relevant msg stream id. [in] level - messaging level. [in] format_id -the template format id to print message accordingly. [in] arg0 - argument #0 to format into the template. [in] arg1 - argument #1 to format into the template.

`flexio_dev_trace_3 (uint8_t tracer_id, flexio_msg_dev_level level, int format_id, uint64_t arg0, uint64_t arg1, uint64_t arg2)`

Creates trace message entry with 3 arguments.

Description

Using the trace mechanism for fast logging. Call the appropriate function according to number of needed arguments.

[in] tracer_id - the relevant msg stream id. [in] level - messaging level. [in] format_id -the template format id to print message accordingly. [in] arg0 - argument #0 to format into the template. [in] arg1 - argument #1 to format into the template. [in] arg2 - argument #2 to format into the template.

`flexio_dev_trace_4 (uint8_t tracer_id, flexio_msg_dev_level level, int format_id, uint64_t arg0, uint64_t arg1, uint64_t arg2, uint64_t arg3)`

Creates trace message entry with 4 arguments.

Description

Using the trace mechanism for fast logging. Call the appropriate function according to number of needed arguments.

[in] tracer_id - the relevant msg stream id. [in] level - messaging level. [in] format_id -the template format id to print message accordingly. [in] arg0 - argument #0 to format into

the template. [in] arg1 - argument #1 to format into the template. [in] arg2 - argument #2 to format into the template. [in] arg3 - argument #3 to format into the template.

`flexio_dev_trace_5 (uint8_t tracer_id, flexio_msg_dev_level level, int format_id, uint64_t arg0, uint64_t arg1, uint64_t arg2, uint64_t arg3, uint64_t arg4)`

Creates trace message entry with 5 arguments.

Description

Using the trace mechanism for fast logging. Call the appropriate function according to number of needed arguments.

[in] tracer_id - the relevant msg stream id. [in] level - messaging level. [in] format_id -the template format id to print message accordingly. [in] arg0 - argument #0 to format into the template. [in] arg1 - argument #1 to format into the template. [in] arg2 - argument #2 to format into the template. [in] arg3 - argument #3 to format into the template. [in] arg4 - argument #4 to format into the template.

`flexio_dev_trace_6 (uint8_t tracer_id, flexio_msg_dev_level level, int format_id, uint64_t arg0, uint64_t arg1, uint64_t arg2, uint64_t arg3, uint64_t arg4, uint64_t arg5)`

Creates trace message entry with 6 arguments.

Description

Using the trace mechanism for fast logging. Call the appropriate function according to number of needed arguments.

[in] tracer_id - the relevant msg stream id. [in] level - messaging level. [in] format_id -the template format id to print message accordingly. [in] arg0 - argument #0 to format into the template. [in] arg1 - argument #1 to format into the template. [in] arg2 - argument #2 to format into the template. [in] arg3 - argument #3 to format into the template. [in] arg4 - argument #4 to format into the template. [in] arg5 - argument #5 to format into the template.

flexio_dev_tracer_flush (uint8_t tracer_id)

Flush not full buffer.

Description

As soon as a buffer is fully occupied it is internal sent to host, however user can ask partially occupied buffer to be sent to host. Its intended use is at end of run to flush whatever messages left. Flush is also performed by the host stream destroy call.

NOTE: this call is not thread safe, user responsibility to avoid calling it while any device trace APIs are in use. Frequent call to this API might cause performance issues.

[in] tracer_id - ID of tracer to flush.

flexio_dev_tracer_notify_host (uint8_t tracer_id, uint32_t num_msg)

Send a tracer buffer to host side.

Description

Send current used buffer to the host. Main usage is to send a full buffer to not risk writing to the buffer from other threads while sending.

NOTE: this call is not thread safe, user responsibility to avoid calling it while any device trace APIs are in use.

[in] tracer_id - ID of tracer to send a notification for.

flexio_dev_status_t flexio_dev_window_config (flexio_dev_thread_ctx *dtctx, uint16_t window_config_id, uint32_t mkey)

Config thread window object.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

window_config_id

- The window object id.

mkey

- mkey object.

Returns

flexio_dev_status_t.

Description

This function updates the thread window object of the given thread context.

flexio_dev_status_t

flexio_dev_window_copy_from_host

(flexio_dev_thread_ctx *dtctx, void *daddr, uint64_t haddr, uint32_t size)

Copy a buffer from host memory to device memory.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

daddr

- A pointer to the device memory buffer.

haddr

- A pointer to the host memory allocated buffer.

size

- Number of bytes to copy.

Returns

flexio_dev_status_t.

Description

This function copies specified number of bytes from host memory to device memory. UNSUPPORTED at this time.

flexio_dev_status_t
flexio_dev_window_copy_to_host
 (flexio_dev_thread_ctx *dtctx, uint64_t haddr,
 const void *daddr, uint32_t size)

Copy a buffer from device memory to host memory.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

haddr

- A pointer to the host memory allocated buffer.

daddr

- A pointer to the device memory buffer.

size

- Number of bytes to copy.

Returns

flexio_dev_status_t.

Description

This function copies specified number of bytes from device memory to host memory.

flexio_dev_status_t
flexio_dev_window_mkey_config
 (flexio_dev_thread_ctx *dtctx, uint32_t mkey)

Config thread window mkey object.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

mkey

- mkey object.

Returns

flexio_dev_status_t.

Description

This function updates the thread window mkey object of the given thread context.

```
flexio_dev_status_t flexio_dev_window_ptr_acquire
(flexio_dev_thread_ctx *dtctx, uint64_t haddr,
flexio_uintptr_t *daddr)
```

Generate device address from host allocated memory.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

haddr

- Host allocated address.

daddr

- A pointer to write the device generated matching address.

Returns

flexio_dev_status_t.

Description

This function generates a memory address to be used by device to access host side memory, according to already create window object. from a host allocated address.

```
flexio_dev_yield (flexio_dev_thread_ctx *dtctx)
```

exit point for continuable event handler routine

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

Description

This function is used to mark the exit point on continuable event handler where user wishes to continue execution on next event. In order to use this API the event handler must be created with continuable flag enabled, otherwise call will have no effect.


```
#define flexio_dev_msg_broadcast
flexio_dev_msg(FLEXIO_MSG_DEV_BROADCAST_STREAM,
lvl, __VA_ARGS__)
```

Create message entry and outputs from the device to all of the host's open streams. Same as a regular printf but with protection from simultaneous print from different threads.

[in] level - messaging level. [in] ... - format and the parameters. Same as for regular printf.

```
#define flexio_dev_msg_dflt
flexio_dev_msg(FLEXIO_MSG_DEV_DEFAULT_STREAM_ID,
lvl, __VA_ARGS__)
```

Create message entry and outputs from the device to host's default stream. Same as a regular printf but with protection from simultaneous print from different threads.

[in] level - messaging level. [in] ... - format and the parameters. Same as for regular printf.

```
#define flexio_dev_print
flexio_dev_msg(FLEXIO_MSG_DEV_DEFAULT_STREAM_ID,
FLEXIO_MSG_DEV_INFO, __VA_ARGS__)
```

Create message entry and outputs from the device to host's default stream, with FLEXIO_MSG_DEV_INFO message level. Same as a regular printf but with protection from simultaneous print from different threads.

[in] ... - format and the parameters. Same as for regular printf.

```
#define spin_init __atomic_store_n(&((lock)-
>locked), 0, __ATOMIC_SEQ_CST)
```

Initialize a spinlock mechanism.

Initialize a spinlock mechanism, must be called before use.

```
#define spin_lock do { \ while
(__atomic_exchange_n(&((lock)->locked), 1,
__ATOMIC_SEQ_CST)) {;} \ } while (0)
```

Lock a spinlock mechanism.

Lock a spinlock mechanism.

```
#define spin_trylock
__atomic_exchange_n(&((lock)->locked), 1,
__ATOMIC_SEQ_CST)
```

Atomic try to catch lock.

makes attempt to take lock. Returns immediately.

```
#define spin_unlock __atomic_store_n(&((lock)-
>locked), 0, __ATOMIC_SEQ_CST)
```

Unlock a spinlock mechanism.

Unlock a spinlock mechanism.

2.3. DevErr

Flex IO SDK device API for DPA programs error handling.

```
enum flexio_dev_error_t
```

Flex IO dev errors.

Values

FLEXIO_DEV_ERROR_ILLEGAL_ERR = 0x42

`__attribute__((__noreturn__))`

Exit the process and return a user (fatal) error code.

Description

Error codes returned to the host in the `dpa_process_status` field of the `DPA_PROCESS` object are defined as follows: 0: OK 1-63: RTOS or Firmware errors 64-127: Flexio-SDK errors 129-255: User defined

`uint64_t flexio_dev_get_and_rst_errno (flexio_dev_thread_ctx *dtctx)`

Get and Reset thread error flag (errno) of recoverable (non fatal) error.

Parameters

dtctx

- A pointer to a `flexio_dev_thread_ctx` structure.

Returns

- void.

`uint64_t flexio_dev_get_errno (flexio_dev_thread_ctx *dtctx)`

Get thread error flag (errno) of recoverable (non fatal) error.

Parameters

dtctx

- A pointer to a `flexio_dev_thread_ctx` structure.

Returns

thread error code.

Description

This function queries an `errno` field from thread context.

flexio_dev_rst_errno (flexio_dev_thread_ctx *dtctx)

Reset thread error flag (errno) of recoverable (non fatal) error.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

2.4. DevQueueAccess

Flex IO SDK device API for DPA programs queue access. Provides an API for handling networking queues (WQs/CQs).

enum flexio_ctrl_seg_t

Flex IO dev WQE control segment types.

Values

FLEXIO_CTRL_SEG_SEND_EN = 0

FLEXIO_CTRL_SEG_SEND_RC = 1

FLEXIO_CTRL_SEG_LDMA = 2

FLEXIO_CTRL_SEG_RDMA_WRITE = 3

FLEXIO_CTRL_SEG_RDMA_READ = 4

FLEXIO_CTRL_SEG_ATOMIC_COMPARE_AND_SWAP = 5

FLEXIO_CTRL_SEG_LSO = 6

FLEXIO_CTRL_SEG_NOP = 7

FLEXIO_CTRL_SEG_RDMA_WRITE_IMM = 8

FLEXIO_CTRL_SEG_TRANSPOSE = 9

enum flexio_dev_cc_db_next_act_t

Flex IO dev congestion control next action types.

Values

CC_DB_NEXT_ACT_SINGLE = 0x0

CC_DB_NEXT_ACT_MULTIPLE = 0x1

CC_DB_NEXT_ACT_FW = 0x2

```
flexio_dev_status_t flexio_dev_cc_ring_db
(flexio_dev_thread_ctx *dtctx, uint16_t
ccq_id, uint32_t rate, uint32_t rtt_req,
flexio_dev_cc_db_next_act_t next_act)
```

Rings CC doorbell.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

ccq_id

- CC queue ID to update.

rate

- Rate to set.

rtt_req

- RTT measure request to set.

next_act

- Next action to set.

Returns

flexio_dev_status_t.

Description

This function rings CC doorbell for the requested CC queue, which sets the requested rate, RTT request and next action.

```
flexio_dev_status_t flexio_dev_cq_arm
(flexio_dev_thread_ctx *dtctx, uint32_t ci, uint32_t
qnum)
```

Arm CQ function.

Parameters

dtctx

- A pointer to a pointer of flexio_dev_thread_ctx structure.

ci

- Current CQ consumer index.

qnum

- Number of the CQ to arm.

Returns

flexio_dev_status_t.

Description

Moves a CQ to 'armed' state. This means that next CQE created for this CQ will result in an EQE on the relevant EQ.

uint32_t flexio_dev_cqe_get_byte_cnt (flexio_dev_cqe64 *cqe)

Get byte count field from CQE function.

Parameters

cqe

- CQE to parse.

Returns

uint32_t - Byte count field value of the CQE.

Description

Parse a CQE for its byte count field.

uint8_t flexio_dev_cqe_get_csum_ok (flexio_dev_cqe64 *cqe)

Get csum OK field from CQE function.

Parameters

cqe

- CQE to parse.

Returns

uint8_t - csum_ok field value of the CQE.

Description

Parse a CQE for its csum OK field.

uint32_t flexio_dev_cqe_get_err_synd (flexio_dev_cqe64 *cqe)

Get error syndrome field from CQE function.

Parameters

cqe

- CQE to parse.

Returns

uint32_t - Error syndrome field value of the CQE.

Description

Parse a CQE for its error syndrome field.

uint8_t flexio_dev_cqe_get_opcode (flexio_dev_cqe64 *cqe)

Get the opcode field from CQE function.

Parameters

cqe

- CQE to parse.

Returns

uint8_t - Opcode field value of the CQE.

uint8_t flexio_dev_cqe_get_owner (flexio_dev_cqe64 *cqe)

Get owner field from CQE function.

Parameters

cqe

- CQE to parse.

Returns

uint8_t - Owner field value of the CQE.

Description

Parse a CQE for its owner field.

```
uint32_t flexio_dev_cqe_get_qpn  
(flexio_dev_cqe64 *cqe)
```

Get QP number field from CQE function.

Parameters

cqe

- CQE to parse.

Returns

uint32_t - QP number field value of the CQE.

Description

Parse a CQE for its QP number field.

```
uint32_t flexio_dev_cqe_get_user_index  
(flexio_dev_cqe64 *cqe)
```

Get the user index field from CQE function.

Parameters

cqe

- CQE to parse.

Returns

uint32_t - User index field value of the CQE.

```
uint16_t flexio_dev_cqe_get_wqe_counter  
(flexio_dev_cqe64 *cqe)
```

Get WQE counter field from CQE function.

Parameters

cqe

- CQE to parse.

Returns

uint16_t - WQE counter field value of the CQE.

Description

Parse a CQE for its WQE counter field.

```
flexio_dev_status_t flexio_dev_db_ctx_arm
(flexio_dev_thread_ctx *dtctx, uint32_t qnum,
uint32_t emu_ctx_id)
```

arm the emulation context

Parameters

dtctx

- A pointer to a pointer of flexio_dev_thread_ctx structure.

qnum

emu_ctx_id

- Emulation context ID, provided by a call on the host to flexio_emu_db_to_cq_ctx_get_id.

```
flexio_dev_status_t
flexio_dev_db_ctx_force_trigger
(flexio_dev_thread_ctx *dtctx, uint32_t cqn,
uint32_t emu_ctx_id)
```

force trigger of emulation context

Parameters

dtctx

cqn

- CQ number provided by host.

emu_ctx_id

- Emulation context ID, provided by a call on the host to flexio_emu_db_to_cq_ctx_get_id.

flexio_dev_status_t flexio_dev_dbr_cq_set_ci (uint32_t *cq_dbr, uint32_t ci)

Set consumer index value for a CQ function.

Parameters

cq_dbr

- A pointer to the CQ's doorbell record address.

ci

- The consumer index value to update.

Returns

flexio_dev_status_t.

Description

Writes an updated consumer index number to a CQ's doorbell record

flexio_dev_status_t flexio_dev_dbr_rq_inc_pi (uint32_t *rq_dbr)

Increment producer index of an RQ by 1 function.

Parameters

rq_dbr

- A pointer to the CQ's doorbell record address.

Returns

flexio_dev_status_t.

Description

Mark a WQE for reuse by incrementing the relevant RQ producer index by 1

`flexio_dev_status_t flexio_dev_eq_update_ci
(flexio_dev_thread_ctx *dtctx, uint32_t ci, uint32_t
qnum)`

Update an EQ consumer index function.

Parameters

dtctx

- A pointer to a pointer of flexio_dev_thread_ctx structure.

ci

- Current EQ consumer index.

qnum

- Number of the EQ to update.

Returns

flexio_dev_status_t.

Description

Updates the consumer index of an EQ after handling an EQE.

`uint32_t flexio_dev_eqe_get_cqn (flexio_dev_eqe
*eqe)`

Get CQ number field from EQE function.

Parameters

eqe

- EQE to parse.

Returns

uint32_t - CQ number field value of the EQE.

Description

Parse an EQE for its CQ number field.

uint8_t flexio_dev_eqe_get_owner (flexio_dev_eqe *eqe)

Get owner field from EQE function.

Parameters

eqe

- EQE to parse.

Returns

uint32_t - owner field value of the EQE.

Description

Parse an EQE for its owner field.

flexio_dev_status_t flexio_dev_msix_send (flexio_dev_thread_ctx *dtctx, uint32_t cqn)

Send msix on the cq linked to the msix eq.

Parameters

dtctx

- A pointer to a flexio_dev_thread_ctx structure.

cqn

- CQ number to trigger db on. Trigger is done via currently configured outbox, this can be changed with outbox config API according to CQ.

Returns

flexio_dev_status_t.

Description

This function trigger msix on the given cq.

```
flexio_dev_status_t flexio_dev_qp_sq_ring_db  
(flexio_dev_thread_ctx *dtctx, uint16_t pi, uint32_t  
qnum)
```

QP/SQ ring doorbell function.

Parameters

dtctx

- A pointer to a pointer of flexio_dev_thread_ctx structure.

pi

- Current queue producer index.

qnum

- Number of the queue to update.

Returns

flexio_dev_status_t.

Description

Rings the doorbell of a QP or SQ in order to alert the HW of pending work.

```
void *flexio_dev_rwqe_get_addr  
(flexio_dev_wqe_rcv_data_seg *rwqe)
```

Get address field from receive WQE function.

Parameters

rwqe

- WQE to parse.

Returns

void* - Address field value of the receive WQE.

Description

Parse a receive WQE for its address field.

`flexio_dev_status_t`
`flexio_dev_swqe_seg_atomic_set`
 (`flexio_dev_sqe_seg *swqe, uint64_t`
`swap_or_add_data, uint64_t compare_data`)

Fill out an Atomic send queue wqe segment function.

Parameters

swqe

- Send WQE segment to fill.

swap_or_add_data

- The data that will be swapped in or the data that will be added.

compare_data

- The data that will be compared with. Unused in fetch & add operation.

Returns

`flexio_dev_status_t`.

Description

Fill the fields of a send WQE segment (2 DWORDs) with Atomic segment information. This segment can service a compare & swap or fetch & add operation.

`flexio_dev_status_t flexio_dev_swqe_seg_ctrl_set`
 (`flexio_dev_sqe_seg *swqe, uint32_t sq_pi,`
`uint32_t sq_number, uint32_t ce, flexio_ctrl_seg_t`
`ctrl_seg_type`)

Fill out a control send queue wqe segment function.

Parameters

swqe

- Send WQE segment to fill.

sq_pi

- Producer index of the send WQE.

sq_number

- SQ number that holds the WQE.

ce

- wanted CQ policy for CQEs. Value is taken from `cq_ce_mode` enum.

ctrl_seg_type

- Type of control segment.

Returns

flexio_dev_status_t.

Description

Fill the fields of a send WQE segment (4 DWORDs) with control segment information. This should always be the 1st segment of the WQE.

```
flexio_dev_status_t flexio_dev_swqe_seg_eth_set
(flexio_dev_sqe_seg *swqe, uint16_t cs_swp_flags,
uint16_t mss, uint16_t inline_hdr_bsz, uint8_t
inline_hdrs)
```

Fill out an ethernet send queue wqe segment function.

Parameters**swqe**

- Send WQE segment to fill.

cs_swp_flags

- Flags for checksum and swap, see PRM section 8.9.4.2, Send WQE Construction Summary.

mss

- Maximum Segment Size - For LSO WQEs - the number of bytes in the TCP payload to be transmitted in each packet. Must be 0 on non LSO WQEs.

inline_hdr_bsz

- Length of inlined packet headers in bytes. This includes the headers in the inline_data segment as well.

inline_hdrs

- First 2 bytes of the inlined packet headers.

Returns

flexio_dev_status_t.

Description

Fill the fields of a send WQE segment (4 DWORDs) with Ethernet segment information.

```
flexio_dev_status_t
flexio_dev_swqe_seg_inline_data_set
(flexio_dev_sqe_seg *swqe, uint32_t data_sz,
uint32_t *data)
```

Fill out an inline data send queue wqe segment function.

Parameters

swqe

- Send WQE segment to fill.

data_sz

- Size of the data.

data

- Inline data array (3 DWORDs).

Returns

flexio_dev_status_t.

Description

Fill the fields of a send WQE segment (4 DWORDs) with inline data segment information.

```
flexio_dev_status_t
flexio_dev_swqe_seg_mem_ptr_data_set
(flexio_dev_sqe_seg *swqe, uint32_t data_sz,
uint32_t lkey, uint64_t data_addr)
```

Fill out a memory pointer data send queue wqe segment function.

Parameters

swqe

- Send WQE segment to fill.

data_sz

- Size of the data.

lkey

- Local memory access key for the data operation.

data_addr

- Address of the data for the data operation.

Returns

flexio_dev_status_t.

Description

Fill the fields of a send WQE segment (4 DWORDs) with memory pointer data segment information.

flexio_dev_status_t flexio_dev_swqe_seg_rdma_set
(flexio_dev_sqe_seg *swqe, uint32_t rkey, uint64_t raddr)

Fill out an RDMA send queue wqe segment function.

Parameters

swqe

- Send WQE segment to fill.

rkey

- Remote memory access key for the RDMA operation.

raddr

Returns

flexio_dev_status_t.

Description

Fill the fields of a send WQE segment (4 DWORDs) with RDMA segment information.

flexio_dev_status_t
flexio_dev_swqe_seg_shared_receive_set
(flexio_dev_sqe_seg *swqe, uint16_t
next_wqe_index, uint8_t signature)

Fill out a Shared receive queue wqe segment function.

Parameters

swqe

- Send WQE segment to fill.

next_wqe_index

- The next wqe index.

signature

- The signature.

Returns

flexio_dev_status_t.

Description

Fill the fields of a linked list shared receive WQE segment.

flexio_dev_status_t**flexio_dev_swqe_seg_transpose_set**

(flexio_dev_sqe_seg *swqe, uint8_t element_size, uint8_t num_of_cols, uint8_t num_of_rows)

Fill out a Transpose send wqe segment function.

Parameters**swqe**

- Send WQE segment to fill.

element_size

- The Matrix element_size.

num_of_cols

- Number of columns in the matrix.

num_of_rows

- Number of rows in the matrix.

Returns

flexio_dev_status_t.

Description

Fill the fields of a send WQE segment (4 DWORDs) with Transpose segment information.

2.5. Change Log

This chapter list changes in API that were introduced to the library.

1.3.0

- ▶ Field Groups, GPU Groups, and field watches created with a handle returned from `dcgmConnect()` are now cleaned up upon disconnect. `dcgmConnect_v2()` can be used to get the old behavior of objects persisting after disconnect.
- ▶ `dcgmConnect_v2()` was added as a method for specifying additional connection options when connecting to the host engine.
- ▶ `dcgmUnwatchFields()` was added as a method of unwatching fields that were previously watched with `dcgmWatchFields()`
- ▶ `dcgmActionValidate_v2()` was added to be able to pass more parameters to the DCGM GPU Diagnostic.
- ▶ `dcgmDiagResponse_t` was increased from v2 to v3. See `dcgmDiagResponse_v3` for details

1.2.3

- ▶ No API changes in this version.

1.1.1

- ▶ `dcgmGetAllSupportedDevices()` was added as a method to get DCGM-supported GPU Ids. `dcgmGetAllDevices()` can still be used to get all GPU Ids in the system.

1.0.0

- ▶ Initial Release.

Chapter 3. Data Structures

Here are the data structures with brief descriptions:

flexio_affinity

flexio_app_attr

flexio_cmdq_attr

flexio_cq_attr

flexio_dev_cqe64

flexio_dev_eqe

flexio_dev_process_tracer_ctx

flexio_dev_sqe_seg

flexio_dev_wqe_atomic_seg

flexio_dev_wqe_ctrl_seg

flexio_dev_wqe_eth_seg

flexio_dev_wqe_inline_data_seg

flexio_dev_wqe_inline_send_data_seg

flexio_dev_wqe_mem_ptr_send_data_seg

flexio_dev_wqe_rcv_data_seg

flexio_dev_wqe_rdma_seg

flexio_dev_wqe_shared_receive_seg

flexio_dev_wqe_transpose_seg

flexio_event_handler_attr

flexio_heap_mem_info

flexio_mkey_attr

flexio_msg_stream_attr_t

flexio_outbox_attr

flexio_process_attr

flexio_qmem

flexio_qp_attr

flexio_qp_attr_opt_param_mask

flexio_tracer_msg

flexio_wq_attr

flexio_wq_sq_attr

spinlock_s

3.1. flexio_affinity Struct Reference

Describes Flex IO thread affinity information.

`uint32_t flexio_affinity::id`

ID of the chosen resource (EU / DPA EU group). Reserved if affinity type none is set.

`enum flexio_affinity_type flexio_affinity::type`

Affinity type to use for a Flex IO thread (none, strict or group).

3.2. flexio_app_attr Struct Reference

Describes process attributes for creating a Flex IO application.

`size_t flexio_app_attr::app_bsize`

DPA application size (bytes).

`const char *flexio_app_attr::app_name`

DPA application name.

`void *flexio_app_attr::app_ptr`

Pointer in the ELF file for the DPA application.

`char *flexio_app_attr::app_sig_sec_name`

Application signature section name.

3.3. flexio_cmdq_attr Struct Reference

Describes process attributes for creating a Flex IO command queue (async RPC).

`int flexio_cmdq_attr::batch_size`

Number of tasks to be executed to completion by invoked thread.

enum flexio_cmdq_state flexio_cmdq_attr::state

Command queue initial state.

int flexio_cmdq_attr::workers

Number of available workers, each worker can handle up to batch_size number of tasks in a single invocation.

3.4. flexio_cq_attr Struct Reference

Describes attributes for creating a Flex IO CQ.

uint8_t flexio_cq_attr::always_armed

Indication to always arm for the created CQ

bool flexio_cq_attr::cc

Indication to enable collapsed CQE for the created CQ.

flexio_uintptr_t flexio_cq_attr::cq_dbr_daddr

DBR memory address for the created CQ.

uint16_t flexio_cq_attr::cq_max_count

CQE compression max count (number of CQEs before creating an event).

uint16_t flexio_cq_attr::cq_period

CQE compression period (number of usecs before creating an event).

enum flexio_cq_period_mode

flexio_cq_attr::cq_period_mode

CQE compression period mode (by CQE or by event).

struct flexio_qmem flexio_cq_attr::cq_ring_qmem

Ring memory info for the created CQ.

`enum flexio_cqe_comp_type`
`flexio_cq_attr::cqe_comp_type`

CQE compression type to use for the CQ (none, basic or enhanced).

`uint8_t flexio_cq_attr::element_type`

Type of the element attached to the created CQ (thread, EQ, none, emulated EQ).

`uint32_t flexio_cq_attr::emulated_eqn`

Emulated EQ number to attach to the created CQ

`uint8_t flexio_cq_attr::log_cq_depth`

Log number of entries for the created CQ.

`bool flexio_cq_attr::no_arm`

Indication to not arm the CQ on creation.

`uint8_t flexio_cq_attr::overrun_ignore`

Indication to ignore overrun for the created CQ.

`flexio_thread *flexio_cq_attr::thread`

Thread object to attach to the created CQ (only valid for element type thread).

`void *flexio_cq_attr::uar_base_addr`

CQ UAR base address, relevant for devx UAR only, otherwise must be NULL.

`uint32_t flexio_cq_attr::uar_id`

CQ UAR ID (devx UAR ID for host queues, otherwise `flexio_uar`).

3.5. `flexio_dev_cqe64` Struct Reference

Describes Flex IO dev CQE.

`__be32 flexio_dev_cqe64::byte_cnt`

0Bh - Byte count.

`uint8_t flexio_dev_cqe64::csum_ok`

07h 24..26 - checksum ok bits.

`__be32 flexio_dev_cqe64::err_syndrome`

0Dh Error syndrome

`uint8_t flexio_dev_cqe64::op_own`

0Fh 0 - Ownership bit.

`__be32 flexio_dev_cqe64::qpn`

0Eh - QPN.

`__be32 flexio_dev_cqe64::rsvd0`

00h..06h - Reserved.

`uint8_t flexio_dev_cqe64::rsvd29`

07h 0..23 - Reserved.

`__be32 flexio_dev_cqe64::rsvd36`

09h..0Ah - Reserved.

`__be32 flexio_dev_cqe64::rsvd48`

0Ch.Reserved.

`uint8_t flexio_dev_cqe64::signature`

0Fh 8..15 - Signature.

`__be32 flexio_dev_cqe64::srqn_uidx`

08h - SRQ number or user index.

`__be16 flexio_dev_cqe64::wqe_counter`

0Fh 16..31 - WQE counter.

3.6. flexio_dev_eqe Struct Reference

Describes Flex IO dev EQE.

`__be32 flexio_dev_eqe::cqcn`

18h 24 lsb - CQN.

`flexio_dev_eqe::@10 flexio_dev_eqe::event_data`

20h - Event data.

`uint8_t flexio_dev_eqe::owner`

3Fh - Owner.

`__be32 flexio_dev_eqe::rsvd00`

00h..17h - Reserved.

`uint8_t flexio_dev_eqe::rsvd00`

00h - Reserved.

`uint8_t flexio_dev_eqe::rsvd02`

02h - Reserved.

`__be16 flexio_dev_eqe::rsvd3c`

3Ch - Reserved.

`uint8_t flexio_dev_eqe::rsvd4`

04h..1fh - Reserved.

`uint8_t flexio_dev_eqe::signature`

3Eh - Signature.

`uint8_t flexio_dev_eqe::sub_type`

03h - Sub type.

`uint8_t flexio_dev_eqe::type`

01h - EQE type.

3.7. flexio_dev_process_tracer_ctx Struct Reference

Describes Flex IO process trace context. This struct is used for managing different tracers for a DPA process.

`flexio_tracer_streams_data`

`*flexio_dev_process_tracer_ctx::tracer_ctx`

< Process trace contexts.

3.8. flexio_dev_sqe_seg Union Reference

Describes Flex IO dev send WQE segments. Only one segment can be set at a given time.

`struct flexio_dev_wqe_atomic_seg`

`flexio_dev_sqe_seg::atomic`

Atomic segment.

`struct flexio_dev_wqe_ctrl_seg`

`flexio_dev_sqe_seg::ctrl`

Control segment.

```
struct flexio_dev_wqe_eth_seg  
flexio_dev_sqe_seg::eth
```

Ethernet segment.

```
struct flexio_dev_wqe_inline_data_seg  
flexio_dev_sqe_seg::inline_data
```

Inline data segment.

```
struct flexio_dev_wqe_inline_send_data_seg  
flexio_dev_sqe_seg::inline_send_data
```

Inline send data segment.

```
struct flexio_dev_wqe_mem_ptr_send_data_seg  
flexio_dev_sqe_seg::mem_ptr_send_data
```

Memory pointer send data segment.

```
struct flexio_dev_wqe_rdma_seg  
flexio_dev_sqe_seg::rdma
```

RDMA segment.

```
struct flexio_dev_wqe_shared_receive_seg  
flexio_dev_sqe_seg::shared_receive
```

Shared receive.

```
struct flexio_dev_wqe_transpose_seg  
flexio_dev_sqe_seg::transpose
```

Transpose segment.

3.9. flexio_dev_wqe_atomic_seg Struct Reference

Describes Flex IO dev WQE ATOMIC segment.

`__be64 flexio_dev_wqe_atomic_seg::compare_data`

02h..03h - Compare operation data.

`__be64`

`flexio_dev_wqe_atomic_seg::swap_or_add_data`

00h..01h - Swap or Add operation data.

3.10. flexio_dev_wqe_ctrl_seg Struct Reference

Describes Flex IO dev WQE control segment.

`__be32 flexio_dev_wqe_ctrl_seg::general_id`

03h - Control general ID.

`__be32 flexio_dev_wqe_ctrl_seg::idx_opcode`

00h - WQE index and opcode.

`__be32 flexio_dev_wqe_ctrl_seg::qpn_ds`

01h - QPN and number of data segments.

`__be32`

`flexio_dev_wqe_ctrl_seg::signature_fm_ce_se`

02h - Signature, fence mode, completion mode and solicited event.

3.11. flexio_dev_wqe_eth_seg Struct Reference

Describes Flex IO dev WQE ethernet segment.

`__be16 flexio_dev_wqe_eth_seg::cs_swp_flags`

01h 16..31 - CS and SWP flags.

`__be16 flexio_dev_wqe_eth_seg::inline_hdr_bsz`

03h 16..31 - Inline headers size (bytes).

`uint8_t flexio_dev_wqe_eth_seg::inline_hdrs`

03h 0..15 - Inline headers (first two bytes).

`__be16 flexio_dev_wqe_eth_seg::mss`

01h 0..15 - Max segment size.

`__be32 flexio_dev_wqe_eth_seg::rsvd0`

00h - Reserved.

`__be32 flexio_dev_wqe_eth_seg::rsvd2`

02h - Reserved.

3.12. flexio_dev_wqe_inline_data_seg Struct Reference

Describes Flex IO dev WQE inline data segment.

`uint8_t flexio_dev_wqe_inline_data_seg::inline_data`

00h..03h - Inline data.

3.13. flexio_dev_wqe_inline_send_data_seg Struct Reference

Describes Flex IO dev WQE inline send data segment.

`__be32`

`flexio_dev_wqe_inline_send_data_seg::byte_count`

00h - Byte count.

`__be32`

`flexio_dev_wqe_inline_send_data_seg::data_and_padding`

01h..03h - Data and padding array.

3.14. flexio_dev_wqe_mem_ptr_send_data_seg Struct Reference

Describes Flex IO dev WQE memory pointer send data segment.

`__be64`

`flexio_dev_wqe_mem_ptr_send_data_seg::addr`

02h..03h - Address.

`__be32`

`flexio_dev_wqe_mem_ptr_send_data_seg::byte_count`

00h - Byte count.

`__be32`

`flexio_dev_wqe_mem_ptr_send_data_seg::lkey`

01h - Local key.

3.15. flexio_dev_wqe_rcv_data_seg Struct Reference

Describes Flex IO dev WQE receive data segment.

`__be64 flexio_dev_wqe_rcv_data_seg::addr`

02h..03h - Address.

`__be32 flexio_dev_wqe_rcv_data_seg::byte_count`

00h - Byte count.

`__be32 flexio_dev_wqe_rcv_data_seg::lkey`

01h - Local key.

3.16. flexio_dev_wqe_rdma_seg Struct Reference

Describes Flex IO dev WQE RDMA segment.

`__be64 flexio_dev_wqe_rdma_seg::raddr`

00h..01h - Remote address.

`__be32 flexio_dev_wqe_rdma_seg::rkey`

02h - Remote key.

`__be32 flexio_dev_wqe_rdma_seg::rsvd0`

03h - Reserved.

3.17. flexio_dev_wqe_shared_receive_seg Struct Reference

Describes Flex IO dev shared receive WQE.

3.18. flexio_dev_wqe_transpose_seg Struct Reference

Describes Flex IO dev WQE transpose segment.

uint8_t

flexio_dev_wqe_transpose_seg::element_size

00h 0..7 - Matrix element size.

uint8_t

flexio_dev_wqe_transpose_seg::num_of_cols

01h 16..22 - Number of columns in matrix (7b).

uint8_t

flexio_dev_wqe_transpose_seg::num_of_rows

01h 0..6 - Number of rows in matrix (7b).

uint8_t flexio_dev_wqe_transpose_seg::rsvd0

00h 8..31 - Reserved.

uint8_t flexio_dev_wqe_transpose_seg::rsvd1

01h - Reserved.

uint8_t flexio_dev_wqe_transpose_seg::rsvd2

01h - Reserved.

`uint8_t flexio_dev_wqe_transpose_seg::rsvd4`

02h..03h - Reserved.

3.19. `flexio_event_handler_attr` Struct Reference

Describes attributes for creating a Flex IO event handler.

`struct flexio_affinity`
`flexio_event_handler_attr::affinity`

Thread's affinity information.

`uint64_t flexio_event_handler_attr::arg`

Thread argument.

`int flexio_event_handler_attr::continuable`

Thread continuable flag.

`flexio_func_t`
`*flexio_event_handler_attr::host_stub_func`

Stub for the entry function of the thread.

`flexio_uintptr_t`
`flexio_event_handler_attr::thread_local_storage_daddr`

Address of the local storage buffer of the thread.

3.20. `flexio_heap_mem_info` Struct Reference

Describes process heap memory information

`size_t flexio_heap_mem_info::allocated`

Process heap memory allocated in bytes.

`uint64_t flexio_heap_mem_info::base_addr`

Process heap memory base address.

`size_t flexio_heap_mem_info::requested`

Process heap memory requested in bytes.

`size_t flexio_heap_mem_info::size`

Process heap memory size in bytes.

3.21. `flexio_mkey_attr` Struct Reference

Describes process attributes for creating a Flex IO MKey.

`int flexio_mkey_attr::access`

`access` contains the access mask for the MKey (Expected values: `IBV_ACCESS_REMOTE_WRITE`, `IBV_ACCESS_LOCAL_WRITE`).

`flexio_uintptr_t flexio_mkey_attr::daddr`

DPA address the MKey is created for.

`size_t flexio_mkey_attr::len`

Length of the address space the MKey is created for.

`ibv_pd *flexio_mkey_attr::pd`

IBV protection domain information for the created MKey.

3.22. flexio_msg_stream_attr_t Struct Reference

Describes DPA msg thread attributes for messaging from the Device to the Host side.

`size_t flexio_msg_stream_attr_t::data_bsize`

Size of buffer, used for data transfer from Flex IO to HOST MUST be power of two and be at least 2Kb.

`flexio_msg_dev_level`
`flexio_msg_stream_attr_t::level`

Log level of the stream.

`struct flexio_affinity`
`flexio_msg_stream_attr_t::mgmt_affinity`

EU affinity for stream management operations creation, modification and destruction
 Passing a nullified struct will set affinity type to 'NONE'.

`char *flexio_msg_stream_attr_t::stream_name`

The name of the stream.

`flexio_msg_dev_sync_mode`
`flexio_msg_stream_attr_t::sync_mode`

Select sync mode scheme.

`enum flexio_tracer_transport`
`flexio_msg_stream_attr_t::tracer_mode`

Tracer transport mode.

`**flexio_msg_stream_attr_t::tracer_msg_formats`

Tracer print format templates array, last entry must be NULL. Device message format ID is used as index to this array.

`flexio_uar *flexio_msg_stream_attr_t::uar`

Deprecated field. Value will be ignored. `flexio_process` UAR be used instead.

3.23. `flexio_outbox_attr` Struct Reference

Describes attributes for creating a Flex IO outbox.

`uint32_t flexio_outbox_attr::en_pcc`

Create outbox with support for CC operations.

`flexio_uar *flexio_outbox_attr::uar`

Deprecated field. Value will be ignored. `flexio_process` UAR will be used instead.

3.24. `flexio_process_attr` Struct Reference

Describes attributes for creating a Flex IO process.

`int flexio_process_attr::en_pcc`

Enable PCC configuration for the created process.

`ibv_pd *flexio_process_attr::pd`

IBV protection domain information for the created process. Passing NULL will result in an internal PD being created and used for the process.

3.25. `flexio_qmem` Struct Reference

Describes queue memory, which may be either host memory or DPA memory

`flexio_uintptr_t flexio_qmem::daddr`

DPA address of the queue memory (only valid for memtype `FLEXIO_MEMTYPE_DPA`).

`uint64_t flexio_qmem::humem_offset`

Address offset in the umem of the queue memory (only valid for memtype FLEXIO_MEMTYPE_HOST).

`enum flexio_memtype flexio_qmem::memtype`

Type of memory to use (FLEXIO_MEMTYPE_DPA or FLEXIO_MEMTYPE_HOST).

`uint32_t flexio_qmem::umem_id`

UMEM ID of the queue memory.

3.26. `flexio_qp_attr` Struct Reference

Describes attributes for creating a Flex IO QP.

`uint8_t *flexio_qp_attr::dest_mac`

Destination MAC address to set for the modified QP

`uint8_t flexio_qp_attr::fl`

Indication to enable force loopback for the modified QP.

`uint8_t flexio_qp_attr::gid_table_index`

GID table index to set for the modified QP

`uint8_t flexio_qp_attr::grh`

GRH to set for the modified QP.

`uint8_t flexio_qp_attr::isolate_vl_tc`

When set, the QP will transmit on an isolated VL/TC if available.

`int flexio_qp_attr::log_rq_depth`

Log number of entries of the QP's RQ.

`uint8_t flexio_qp_attr::log_rra_max`

Log of the number of allowed outstanding RDMA read/atomic operations

`int flexio_qp_attr::log_sq_depth`

Log number of entries of the QP's SQ.

`uint8_t flexio_qp_attr::log_sra_max`

Log of the number of allowed outstanding RDMA read/atomic operations as requester

`uint32_t flexio_qp_attr::min_rnr_nak_timer`

Minimal RNR NACK timer to set for the modified QP.

`uint32_t flexio_qp_attr::next_rcv_psn`

Next receive PSN to set for the modified QP.

`uint32_t flexio_qp_attr::next_send_psn`

Next send PSN to set for the modified QP.

`enum flexio_qp_state flexio_qp_attr::next_state`

QP state to move the QP to (reset, init, RTS, RTR).

`int flexio_qp_attr::no_sq`

Indication to create the QP without an SQ.

`int flexio_qp_attr::ops_flag`

deprecated.

`enum flexio_qp_qpc_mtu flexio_qp_attr::path_mtu`

Path MTU to set for the modified QP.

`ibv_pd *flexio_qp_attr::pd`

IBV protection domain information for the created QP.

int flexio_qp_attr::qp_access_mask

QP's access permission (Expected values: IBV_ACCESS_REMOTE_WRITE, IBV_ACCESS_REMOTE_READ, IBV_ACCESS_REMOTE_ATOMIC, IBV_ACCESS_LOCAL_WRITE).

struct flexio_qmem

flexio_qp_attr::qp_wq_buff_qmem

Ring memory info for the created QP's WQ.

struct flexio_qmem

flexio_qp_attr::qp_wq_dbr_qmem

DBR memory info for the created QP's WQ.

uint32_t flexio_qp_attr::remote_qp_num

Remote QP number to set for the modified QP.

uint8_t flexio_qp_attr::retry_count

Retry count to set for the modified QP.

ibv_gid flexio_qp_attr::rgid_or_rip

Remote GID or remote IP to set for the modified QP.

uint16_t flexio_qp_attr::rlid

Remote LID to set for the modified QP.

uint32_t flexio_qp_attr::rmpqn

RMP queue number, relevant only if QP RQ is RMP.

uint32_t flexio_qp_attr::rq_cqn

CQ number of the QP's RQ. Not relevant for RMP

`int flexio_qp_attr::rq_type`

QP's RQ type (regular, RMP, zero-RQ)

`uint32_t flexio_qp_attr::sq_cqn`

CQ number of the QP's SQ.

`uint32_t flexio_qp_attr::transport_type`

QP's transport type (currently only FLEXIO_QPC_ST_RC is supported).

`uint32_t flexio_qp_attr::uar_id`

QP UAR ID.

`uint16_t flexio_qp_attr::udp_sport`

UDP port to set for the modified QP.

`uint32_t flexio_qp_attr::user_index`

User defined user_index for the created QP.

`uint8_t flexio_qp_attr::vhca_port_num`

VHCA port number to set for the modified QP.

3.27. flexio_qp_attr_opt_param_mask Struct Reference

Describes QP modify operation mask.

`bool`

`flexio_qp_attr_opt_param_mask::min_rnr_nak_timer`

Indication to modify the QP's min_rnr_nak_timer field.

`bool`

`flexio_qp_attr_opt_param_mask::qp_access_mask`

Indication to modify the QP's `qp_access_mask` field.

3.28. `flexio_tracer_msg` Struct Reference

Describes Flex IO trace message. This struct is used to communicate the tracer raw data from device to host.

`uint64_t flexio_tracer_msg::arg0`

Argument 0 for trace string format.

`uint64_t flexio_tracer_msg::arg1`

Argument 1 for trace string format.

`uint64_t flexio_tracer_msg::arg2`

Argument 2 for trace string format.

`uint64_t flexio_tracer_msg::arg3`

Argument 3 for trace string format.

`uint64_t flexio_tracer_msg::arg4`

Argument 4 for trace string format.

`uint64_t flexio_tracer_msg::arg5`

Argument 5 for trace string format.

`uint32_t flexio_tracer_msg::format_id`

Format ID for trace string template to use.

3.29. flexio_wq_attr Struct Reference

Describes attributes for creating a Flex IO WQ.

`uint8_t flexio_wq_attr::log_wq_depth`

Log number of entries for the created WQ.

`uint8_t flexio_wq_attr::log_wq_stride`

Log size of entry for the created WQ. If this parameter is not provided, it will be set to default value 4.

`ibv_pd *flexio_wq_attr::pd`

IBV protection domain struct to use for creating the WQ.

`flexio_wq_rq_attr flexio_wq_attr::rq`

RMP attributes (used only for RMPs). Cyclic WQ type is not supported.

`struct flexio_wq_sq_attr flexio_wq_attr::sq`

SQ attributes (used only for SQs).

`uint32_t flexio_wq_attr::uar_id`

WQ UAR ID.

`uint32_t flexio_wq_attr::user_index`

User defined user_index for the created WQ.

`struct flexio_qmem flexio_wq_attr::wq_dbr_qmem`

DBR memory address for the created WQ.

`struct flexio_qmem flexio_wq_attr::wq_ring_qmem`

Ring memory info for the created WQ.

3.30. flexio_wq_sq_attr Struct Reference

Describes attributes for creating a Flex IO SQ.

uint8_t

flexio_wq_sq_attr::allow_multi_pkt_send_wqe

Indication enable multi packet send WQE for the created SQ.

3.31. spinlock_s Struct Reference

Describes Flex IO dev spinlock.

uint32_t spinlock_s::locked

Indication for spinlock lock state.

Chapter 4. Data Fields

Here is a list of all documented struct and union fields with links to the struct/union documentation for each field:

A

access

[flexio_mkey_attr](#)

addr

[flexio_dev_wqe_mem_ptr_send_data_seg](#)

[flexio_dev_wqe_rcv_data_seg](#)

affinity

[flexio_event_handler_attr](#)

allocated

[flexio_heap_mem_info](#)

allow_multi_pkt_send_wqe

[flexio_wq_sq_attr](#)

always_armed

[flexio_cq_attr](#)

app_bsize

[flexio_app_attr](#)

app_name

[flexio_app_attr](#)

app_ptr

[flexio_app_attr](#)

app_sig_sec_name

[flexio_app_attr](#)

arg

[flexio_event_handler_attr](#)

arg0

[flexio_tracer_msg](#)

arg1

[flexio_tracer_msg](#)

arg2

[flexio_tracer_msg](#)

arg3[flexio_tracer_msg](#)**arg4**[flexio_tracer_msg](#)**arg5**[flexio_tracer_msg](#)**atomic**[flexio_dev_sqe_seg](#)**B****base_addr**[flexio_heap_mem_info](#)**batch_size**[flexio_cmdq_attr](#)**byte_cnt**[flexio_dev_cqe64](#)**byte_count**[flexio_dev_wqe_mem_ptr_send_data_seg](#)[flexio_dev_wqe_inline_send_data_seg](#)[flexio_dev_wqe_rcv_data_seg](#)**C****cc**[flexio_cq_attr](#)**compare_data**[flexio_dev_wqe_atomic_seg](#)**continuable**[flexio_event_handler_attr](#)**cq_dbr_daddr**[flexio_cq_attr](#)**cq_max_count**[flexio_cq_attr](#)**cq_period**[flexio_cq_attr](#)**cq_period_mode**[flexio_cq_attr](#)**cq_ring_qmem**[flexio_cq_attr](#)**cqe_comp_type**[flexio_cq_attr](#)**cqn**[flexio_dev_eqe](#)

cs_swp_flags[flexio_dev_wqe_eth_seg](#)**csum_ok**[flexio_dev_cqe64](#)**ctrl**[flexio_dev_sqe_seg](#)**D****daddr**[flexio_qmem](#)[flexio_mkey_attr](#)**data_and_padding**[flexio_dev_wqe_inline_send_data_seg](#)**data_bsize**[flexio_msg_stream_attr_t](#)**dest_mac**[flexio_qp_attr](#)**E****element_size**[flexio_dev_wqe_transpose_seg](#)**element_type**[flexio_cq_attr](#)**emulated_eqn**[flexio_cq_attr](#)**en_pcc**[flexio_process_attr](#)[flexio_outbox_attr](#)**err_syndrome**[flexio_dev_cqe64](#)**eth**[flexio_dev_sqe_seg](#)**event_data**[flexio_dev_eqe](#)**F****fl**[flexio_qp_attr](#)**format_id**[flexio_tracer_msg](#)

G

general_id[flexio_dev_wqe_ctrl_seg](#)**gid_table_index**[flexio_qp_attr](#)**grh**[flexio_qp_attr](#)

H

host_stub_func[flexio_event_handler_attr](#)**humem_offset**[flexio_qmem](#)

I

id[flexio_affinity](#)**idx_opcode**[flexio_dev_wqe_ctrl_seg](#)**inline_data**[flexio_dev_sqe_seg](#)[flexio_dev_wqe_inline_data_seg](#)**inline_hdr_bsz**[flexio_dev_wqe_eth_seg](#)**inline_hdrs**[flexio_dev_wqe_eth_seg](#)**inline_send_data**[flexio_dev_sqe_seg](#)**isolate_vl_tc**[flexio_qp_attr](#)

L

len[flexio_mkey_attr](#)**level**[flexio_msg_stream_attr_t](#)**lkey**[flexio_dev_wqe_rcv_data_seg](#)[flexio_dev_wqe_mem_ptr_send_data_seg](#)**locked**[spinlock_s](#)

log_cq_depth[flexio_cq_attr](#)**log_rq_depth**[flexio_qp_attr](#)**log_rra_max**[flexio_qp_attr](#)**log_sq_depth**[flexio_qp_attr](#)**log_sra_max**[flexio_qp_attr](#)**log_wq_depth**[flexio_wq_attr](#)**log_wq_stride**[flexio_wq_attr](#)**M****mem_ptr_send_data**[flexio_dev_sqe_seg](#)**memtype**[flexio_qmem](#)**mgmt_affinity**[flexio_msg_stream_attr_t](#)**min_rnr_nak_timer**[flexio_qp_attr_opt_param_mask](#)[flexio_qp_attr](#)**mss**[flexio_dev_wqe_eth_seg](#)**N****next_rcv_psn**[flexio_qp_attr](#)**next_send_psn**[flexio_qp_attr](#)**next_state**[flexio_qp_attr](#)**no_arm**[flexio_cq_attr](#)**no_sq**[flexio_qp_attr](#)**num_of_cols**[flexio_dev_wqe_transpose_seg](#)**num_of_rows**[flexio_dev_wqe_transpose_seg](#)

O

op_own[flexio_dev_cqe64](#)**ops_flag**[flexio_qp_attr](#)**overrun_ignore**[flexio_cq_attr](#)**owner**[flexio_dev_eqe](#)

P

path_mtu[flexio_qp_attr](#)**pd**[flexio_mkey_attr](#)[flexio_qp_attr](#)[flexio_process_attr](#)[flexio_wq_attr](#)

Q

qp_access_mask[flexio_qp_attr_opt_param_mask](#)[flexio_qp_attr](#)**qp_wq_buff_qmem**[flexio_qp_attr](#)**qp_wq_dbr_qmem**[flexio_qp_attr](#)**qpn**[flexio_dev_cqe64](#)**qpn_ds**[flexio_dev_wqe_ctrl_seg](#)

R

raddr[flexio_dev_wqe_rdma_seg](#)**rdma**[flexio_dev_sqe_seg](#)**remote_qp_num**[flexio_qp_attr](#)**requested**[flexio_heap_mem_info](#)

retry_count[flexio_qp_attr](#)**rgid_or_rip**[flexio_qp_attr](#)**rkey**[flexio_dev_wqe_rdma_seg](#)**rlid**[flexio_qp_attr](#)**rmpqn**[flexio_qp_attr](#)**rq**[flexio_wq_attr](#)**rq_cqn**[flexio_qp_attr](#)**rq_type**[flexio_qp_attr](#)**rsvd0**[flexio_dev_wqe_transpose_seg](#)[flexio_dev_cqe64](#)[flexio_dev_wqe_eth_seg](#)[flexio_dev_wqe_rdma_seg](#)**rsvd00**[flexio_dev_eqe](#)**rsvd02**[flexio_dev_eqe](#)**rsvd1**[flexio_dev_wqe_transpose_seg](#)**rsvd2**[flexio_dev_wqe_transpose_seg](#)[flexio_dev_wqe_eth_seg](#)**rsvd29**[flexio_dev_cqe64](#)**rsvd36**[flexio_dev_cqe64](#)**rsvd3c**[flexio_dev_eqe](#)**rsvd4**[flexio_dev_eqe](#)[flexio_dev_wqe_transpose_seg](#)**rsvd48**[flexio_dev_cqe64](#)

S

shared_receive[flexio_dev_sqe_seg](#)**signature**[flexio_dev_eqe](#)[flexio_dev_cqe64](#)**signature_fm_ce_se**[flexio_dev_wqe_ctrl_seg](#)**size**[flexio_heap_mem_info](#)**sq**[flexio_wq_attr](#)**sq_cqn**[flexio_qp_attr](#)**srqn_uidx**[flexio_dev_cqe64](#)**state**[flexio_cmdq_attr](#)**stream_name**[flexio_msg_stream_attr_t](#)**sub_type**[flexio_dev_eqe](#)**swap_or_add_data**[flexio_dev_wqe_atomic_seg](#)**sync_mode**[flexio_msg_stream_attr_t](#)

T

thread[flexio_cq_attr](#)**thread_local_storage_daddr**[flexio_event_handler_attr](#)**tracer_ctx**[flexio_dev_process_tracer_ctx](#)**tracer_mode**[flexio_msg_stream_attr_t](#)**tracer_msg_formats**[flexio_msg_stream_attr_t](#)**transport_type**[flexio_qp_attr](#)**transpose**[flexio_dev_sqe_seg](#)

type

[flexio_affinity](#)
[flexio_dev_eqe](#)

U

uar

[flexio_msg_stream_attr_t](#)
[flexio_outbox_attr](#)

uar_base_addr

[flexio_cq_attr](#)

uar_id

[flexio_cq_attr](#)
[flexio_wq_attr](#)
[flexio_qp_attr](#)

udp_sport

[flexio_qp_attr](#)

umem_id

[flexio_qmem](#)

user_index

[flexio_wq_attr](#)
[flexio_qp_attr](#)

V

vhca_port_num

[flexio_qp_attr](#)

W

workers

[flexio_cmdq_attr](#)

wq_dbr_qmem

[flexio_wq_attr](#)

wq_ring_qmem

[flexio_wq_attr](#)

wqe_counter

[flexio_dev_cqe64](#)

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: "NVIDIA") make no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assume no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of Mellanox Technologies Ltd. and/or NVIDIA Corporation in the U.S. and in other countries. The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2023 NVIDIA Corporation & affiliates. All rights reserved.