



Application Note

# Mellanox IB-Verbs API (VAPI)

Mellanox Software Programmer's

Interface for InfiniBand Verbs

---

Copyright 2001. Mellanox Technologies, Inc. All Rights Reserved.

Mellanox IB-Verbs API (VAPI)

Document Number: 2088AN

Mellanox Technologies, Inc.  
2900 Stender Way  
Santa Clara, CA 95054  
U.S.A.  
[www.Mellanox.com](http://www.Mellanox.com)

Tel: (408) 970-3400  
Fax: (408) 970-3403

Mellanox Technologies Ltd.  
POB 586 Yokneam 20692  
Israel

Tel: +972-4-909-7200  
Fax: +972-4959-3245

# Contents

<b>List of Tables</b>	<b>3</b>
<b>Chapter 1 Introduction</b>	<b>5</b>
<b>Chapter 2 HCA Verbs</b>	<b>7</b>
2.1 Open HCA	7
2.2 Get HCA Handle	8
2.3 Query HCA Capabilities	9
2.4 Query HCA Port Properties	12
2.5 Query HCA Gid Tbl	14
2.6 Query HCA Pkey Tbl	15
2.7 Modify HCA Attributes	16
2.8 Close HCA	18
2.9 Allocate Protection Domain	19
2.10 Deallocate Protection Domain	20
2.11 Allocate Reliable Datagram Domain	21
2.12 Deallocate Reliable Datagram Domain	22
<b>Chapter 3 Address Management Verbs</b>	<b>23</b>
3.1 Create Address Handle	23
3.2 Modify Address Handle	25
3.3 Query Address Handle	26
3.4 Destroy Address Handle	27
<b>Chapter 4 Queue Pair Verbs</b>	<b>29</b>
4.1 Create Queue Pair	29
4.2 Modify Queue Pair	32
4.3 Query Queue Pair	36
4.4 Destroy Queue Pair	37
4.5 Get Special QP	38
<b>Chapter 5 Completion Queue Verbs</b>	<b>41</b>
5.1 Create Completion Queue	41
5.2 Query Completion Queue	43
5.3 Resize Completion Queue	44
5.4 Destroy Completion Queue	45
<b>Chapter 6 EE Context Verbs</b>	<b>47</b>
6.1 Create EE Context	47
6.2 Modify EE Context Attributes	48
6.3 Query EE Context	49
6.4 Destroy EE Context	50
<b>Chapter 7 Memory Management Verbs</b>	<b>51</b>
7.1 Register Memory Region	51
7.2 Query Memory Region	53
7.3 Deregister Memory Region	54
7.4 Reregister Memory Region	55
7.5 Register Shared Memory Region	56
7.6 Allocate Memory Window	57
7.7 Query Memory Window	58
7.8 Bind Memory Window	59

---

7.9	Deallocate Memory Window	60
<b>Chapter 8</b>	<b>Multicast Verbs</b>	<b>61</b>
8.1	Attach QP to Multicast Group	61
8.2	Detach QP from Multicast Group	62
<b>Chapter 9</b>	<b>Work Request Verbs</b>	<b>63</b>
9.1	Post Send Request	63
9.2	Post Receive Request	66
9.3	Poll for Completion	67
9.4	Request Completion Notification	69
9.5	Request Completion Notification for a Completion Queue	70
9.6	Clear Completion Notification for a Completion Queue	71
<b>Chapter 10</b>	<b>Event Handling Verbs</b>	<b>73</b>
10.1	Set Completion Event Handler	73
10.2	Set Asynchronous Event Handler	74

## *List of Tables*

Table 1: VAPI Calls List	5
Table 2: VAPI_hca_vendor_t	9
Table 3: VAPI_hca_cap_t	10
Table 4: VAPI_hca_port_t	12
Table 5: VAPI_hca_attr_t	16
Table 6: HCA Attributes Flags	17
Table 7: HCA Attributes Mask Macros	17
Table 8: VAPI_av_t	24
Table 9: VAPI_qp_init_attr_t	30
Table 10: VAPI_qp_cap_t	30
Table 11: VAPI_qp_prop_t	31
Table 12: VAPI_qp_attr_t	33
Table 13: QP Attributes Mask Macros	34
Table 14: VAPI_mrw_t	52
Table 15: VAPI_sr_desc_t	64
Table 16: VAPI_sg_lst_entry_t	65
Table 17: VAPI_rr_desc_t	66
Table 18: VAPI_wc_desc_t	67
Table 19: VAPI_remote_node_addr_t	68
Table 20: VAPI_event_record_t	75

---



# 1 Introduction

This document provides a C-Language definition of the VERBS abstraction defined in the IBTA specification Volume 1, chapter 11.

This interface will be supported by all Mellanox HCA devices and is designed to easily absorb HW changes and support different OS without exposing those changes to the application running above the VERBS.

**Table 1 VAPI Calls List**

<i>VAPI CALL</i>	<i>Description</i>	<i>Status</i>	<i>Page</i>
<a href="#">VAPI_open_hca</a>	<a href="#">Open HCA</a>	Supported	<a href="#">p.7</a>
<a href="#">EVAPI_get_hca_hndl</a>	<a href="#">Get HCA Handle</a>	Supported	<a href="#">p.9</a>
<a href="#">VAPI_query_hca_cap</a>	<a href="#">Query HCA Capabilities</a>	Supported	<a href="#">p.10</a>
<a href="#">VAPI_query_hca_port_prop</a>	<a href="#">Query HCA Port Properties</a>	Supported	<a href="#">p.13</a>
<a href="#">VAPI_query_hca_gid_tbl</a>	<a href="#">Query HCA Gid Tbl</a>	Supported	<a href="#">p.15</a>
<a href="#">VAPI_query_hca_pkey_tbl</a>	<a href="#">Query HCA Pkey Tbl</a>	Supported	<a href="#">p.16</a>
<a href="#">VAPI_modify_hca_attr</a>	<a href="#">Modify HCA Attributes</a>	Supported	<a href="#">p.17</a>
<a href="#">VAPI_close_hca</a>	<a href="#">Close HCA</a>	Supported	<a href="#">p.20</a>
<a href="#">VAPI_alloc_pd</a>	<a href="#">Allocate Protection Domain</a>	Supported	<a href="#">p.21</a>
<a href="#">VAPI_dealloc_pd</a>	<a href="#">Deallocate Protection Domain</a>	Supported	<a href="#">p.22</a>
<a href="#">VAPI_alloc_rdd</a>	<a href="#">Allocate Reliable Datagram Domain</a>	None	<a href="#">p.23</a>
<a href="#">VAPI_dealloc_rdd</a>	<a href="#">Deallocate Reliable Datagram Domain</a>	None	<a href="#">p.24</a>
<a href="#">VAPI_create_addr_hndl</a>	<a href="#">Create Address Handle</a>	Supported	<a href="#">p.25</a>
<a href="#">VAPI_modify_addr_hndl</a>	<a href="#">Modify Address Handle</a>	Supported	<a href="#">p.27</a>
<a href="#">VAPI_query_addr_hndl</a>	<a href="#">Query Address Handle</a>	Supported	<a href="#">p.28</a>
<a href="#">VAPI_destroy_addr_hndl</a>	<a href="#">Destroy Address Handle</a>	Supported	<a href="#">p.29</a>
<a href="#">VAPI_create_qp</a>	<a href="#">Create Queue Pair</a>	Supported	<a href="#">p.31</a>
<a href="#">VAPI_modify_qp</a>	<a href="#">Modify Queue Pair</a>	Supported	<a href="#">p.34</a>
<a href="#">VAPI_query_qp</a>	<a href="#">Query Queue Pair</a>	Supported	<a href="#">p.37</a>
<a href="#">VAPI_destroy_qp</a>	<a href="#">Destroy Queue Pair</a>	Supported	<a href="#">p.38</a>
<a href="#">VAPI_get_special_qp</a>	<a href="#">Get Special QP</a>	Supported	<a href="#">p.39</a>
<a href="#">VAPI_create_cq</a>	<a href="#">Create Completion Queue</a>	Supported	<a href="#">p.41</a>
<a href="#">VAPI_query_cq</a>	<a href="#">Query Completion Queue</a>	Supported	<a href="#">p.43</a>

Table 1 VAPI Calls List (Continued)

<i>VAPI CALL</i>	<i>Description</i>	<i>Status</i>	<i>Page</i>
VAPI_resize_cq	Resize Completion Queue	None	<a href="#">p.44</a>
VAPI_destroy_cq	Destroy Completion Queue	Supported	<a href="#">p.45</a>
VAPI_create_eec	Create EE Context	None	<a href="#">p.47</a>
VAPI_modify_eec_attr	Modify EE Context Attributes	None	<a href="#">p.48</a>
VAPI_query_eec_attr	Query EE Context	None	<a href="#">p.49</a>
VAPI_destroy_eec	Destroy EE Context	None	<a href="#">p.50</a>
VAPI_register_mr	Register Memory Region	Supported	<a href="#">p.51</a>
VAPI_query_mr	Query Memory Region	Supported	<a href="#">p.53</a>
VAPI_deregister_mr	Deregister Memory Region	Supported	<a href="#">p.54</a>
VAPI_reregister_mr	Reregister Memory Region	None	<a href="#">p.55</a>
VAPI_register_smr	Register Shared Memory Region	None	<a href="#">p.56</a>
VAPI_alloc_mw	Allocate Memory Window	None	<a href="#">p.57</a>
VAPI_query_mw	Query Memory Window	None	<a href="#">p.58</a>
VAPI_bind_mw	Bind Memory Window	None	<a href="#">p.59</a>
VAPI_dealloc_mw	Deallocate Memory Window	None	<a href="#">p.60</a>
VAPI_attach_to_multicast	Attach QP to Multicast Group	Supported	<a href="#">p.61</a>
VAPI_detach_from_multicast	Detach QP from Multicast Group	Supported	<a href="#">p.62</a>
VAPI_post_sr	Post Send Request	Supported	<a href="#">p.63</a>
VAPI_post_rr	Post Receive Request	Supported	<a href="#">p.66</a>
VAPI_poll_cq	Poll for Completion	Supported	<a href="#">p.67</a>
VAPI_req_comp_notif	Request Completion Notification	Supported	<a href="#">p.69</a>
EVAPI_set_comp_eventh	Request Completion Notification for a Completion Queue	Supported	<a href="#">p.70</a>
EVAPI_clear_comp_eventh	Clear Completion Notification for a Completion Queue	Supported	<a href="#">p.71</a>
VAPI_set_comp_event_handler	Set Completion Event Handler	Supported	<a href="#">p.73</a>
VAPI_set_async_event_handler	VAPI_query_hca_cap	Supported	<a href="#">p.74</a>



## 2 HCA Verbs

This chapter describes the Mellanox implementation of the following verbs handling basic HCA functionality:

- Open HCA ([p.7](#))
- Get HCA Handle ([p.9](#))
- Query HCA Capabilities ([p.10](#))
- Query HCA Port Properties ([p.13](#))
- Query HCA GID Tbl ([p.15](#))
- Query HCA Pkey Tbl ([p.16](#))
- Modify HCA Attributes ([p.17](#))
- Release HCA Handle ([p.19](#))
- Close HCA ([p.20](#))
- Allocate Protection Domain ([p.21](#))
- Deallocate Protection Domain ([p.22](#))
- Allocate Reliable Datagram Domain ([p.23](#))
- Deallocate Reliable Datagram Domain ([p.24](#))

### 2.1 Open HCA

#### SYNOPSIS:

```
VAPI_ret_t
VAPI_open_hca
(
    IN      VAPI_hca_id_t      hca_id,
    OUT    VAPI_hca_hndl_t    *hca_hndl_p
)

```

**ARGUMENTS:** **hca\_id:** HCA identifier.  
**hca\_hndl\_p:** Pointer to the HCA object handle.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN:** Insufficient resources.  
**VAPI\_EINVAL\_HCA\_ID:** Invalid HCA identifier.  
**VAPI\_EBUSY:** HCA already in use.

**DESCRIPTION:**

Creates a new HCA Object.  
The newly created object is assigned to the device described by  
hca\_id\_num.

## 2.2 Get HCA Handle

### SYNOPSIS:

```
VAPI_ret_t
EVAPI_get_hca_hndl
(
    IN      VAPI_hca_id_t      hca_id,
    OUT     VAPI_hca_hndl_t    *hca_hndl_p
)
```

**ARGUMENTS:** **hca\_id:** HCA identifier.  
**hca\_hndl\_p:** Pointer to the HCA object handle.

**RETURNS:** **VAPI\_OK**  
**VIP\_EINVAL\_HCA\_HNDL :** No such opened HCA.

**DESCRIPTION:** Gets the handle of an already opened HCA. All applications should use this function to get the HCA handle to work with.

## 2.3 Query HCA Capabilities

### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_hca_cap
(
    IN      VAPI_hca_hndl_t      hca_hndl,
    OUT    VAPI_hca_vendor_t    *hca_vendor_p
    OUT    VAPI_hca_cap_t       *hca_cap_p
)
```

**ARGUMENTS:** **hca\_hndl:** HCA object handle.  
**hca\_vendor\_p:** Pointer to HCA vendor-specific information object.  
**hca\_cap\_p:** Pointer to HCA capabilities object

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EAGAIN:** Insufficient resources.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Query HCA capabilities retrieves:

- A structure of type **VAPI\_hca\_vendor\_t**, providing a list of the vendor-specific information about the HCA,
- A structure of type **VAPI\_hca\_cap\_t**, providing a detailed list of the HCA capabilities.

Further information on the hca ports can be retrieved using the verbs **VAPI\_query\_hca\_port\_prop**, **VAPI\_query\_hca\_gid\_tbl** and **VAPI\_query\_hca\_pkey\_tbl**.

The following table lists the vendor-specific parameters, which are described in the structure **VAPI\_hca\_vendor\_t**:

**Table 2 VAPI\_hca\_vendor\_t**

vendor_id	Vendor ID
vendor_part_id	Vendor supplied part ID
hw_ver	Hardware Version

The **hca\_cap** structure describes the HCA capabilities and is of type **VAPI\_hca\_cap\_t**, as listed in the following table:

**Table 3 VAPI\_hca\_cap\_t**

max_num_qp	Maximum number of QP supported.
max_qp_ous_wr	Maximum number of outstanding work requests per QP supported does not appear in VIP.
flags	Different capabilities: <ul style="list-style-type: none"> <li>- VAPI_RESIZE_OUS_WQE_CAP - ability to resize maximum outstanding WR per QP.</li> <li>- VAPI_BAD_PKEY_COUNT_CAP - bad p_key counter support indication.</li> <li>- VAPI_BAD_QKEY_COUNT_CAP - q_key violation counter support indication.</li> <li>- VAPI_RAW_MULTI_CAP - supports raw packets multicast.</li> <li>- VAPI_AUTO_PATH_MIG_CAP - supports automatic path migration</li> <li>- VAPI_CHANGE_PHY_PORT_CAP - Ability to change the primary physical port when transitioning from SQD to RTS state.</li> <li>- VAPI_UD_AV_PORT_ENFORCE_CAP - Port check enforced for UD address vectors.</li> </ul>
max_num_sg_ent	Maximum number of scatter gather entries for descriptors other than reliable datagram.
max_num_sg_ent_rd	Maximum number of scatter gather entries for descriptors reliable datagram.
max_num_cq	Maximum number of completion queues supported by this HCA.
max_num_ent_cq	Maximum number of completion queue entries in each CQ.
max_num_mr	Maximum number of memory regions supported.
max_mr_size	Maximum size of contiguous memory region.
max_pd_num	Maximum number of protection domains supported.
page_size_cap	Memory maximum page size.
phys_port_num	Number of physical ports.
max_pkeys	Number of partitions supported.
node_guid	Node GUID for this HCA.
local_ca_ack_delay	Max expected time between message receive and ACK or NAK transmission.
max_qp_ous_rd_atom	Maximum number of read or atomic operation requests that can be outstanding on a Queue Pair.
max_ee_ous_rd_atom	Maximum number of read or atomic operation requests that can be outstanding on an EE.
max_res_rd_atom	Maximum number of resources that must be allocated to incoming RDMA Read or Atomic.
max_qp_init_rd_atom	Maximum number of RDMA Reads or Atomic operations that can be initiated on a specific QP.
max_ee_init_rd_atom	Maximum number of RDMA Reads or Atomic operations that can be initiated on a specific QP (hard coded to 1 <b><i>on IB v1.0</i></b> )

**Table 3 VAPI\_hca\_cap\_t (Continued)**

atomic_cap	Level of atomicity supported: <ul style="list-style-type: none"> <li>- VAPI_NOT_SUPPORTED - atomic operations not supported.</li> <li>- VAPI_LOCAL_QP - atomicity guaranteed only between QPs on this HCA.</li> <li>- VAPI_ALL - atomicity is guaranteed between this HCA and any other device.</li> </ul>
max_ee_num	Maximum number of EE contexts supported.
max_rdd_num	Maximum number of RDDs supported.
max_mw_num	Maximum Number of Memory Windows supported.
max_raw_ipv6_qp	Maximum number of raw IPV6 QPs supported.
max_raw_ethy_qp	Maximum number of raw Ethertype packets supported by this QP.
max_mcast_grp_num	Maximum number of multicast groups supported.
max_mcast_qp_attach_num	Maximum number of QPs that can be attached to a multicast group.
max_total_mcast_qp_attach_num	Maximum number of QPs that can be attached to any multicast group.
max_ah_num	Maximum number of address handles.

## 2.4 Query HCA Port Properties

### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_hca_port_prop
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      IB_port_t         port_num,
    OUT     VAPI_hca_port_t    *hca_port_p
)

```

**ARGUMENTS:** **hca\_hndl:** HCA object handle.  
**port\_num:** Port number  
**hca\_port\_p:** HCA port object describing the port properties.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EAGAIN:** Insufficient resources.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Query HCA port properties retrieves a structure of type [VAPI\\_hca\\_port\\_t](#) for the port specified in **port\_num**. The number of the HCA physical ports can be obtained using the verb [VAPI\\_query\\_hca\\_cap](#). Further information about the port p-key table and gid table can be obtained using the verbs [VAPI\\_query\\_hca\\_gid\\_tbl](#) and [VAPI\\_query\\_hca\\_pkey\\_tbl](#).

Upon successful completion, the verb returns a **hca\_port\_p** structure of type [VAPI\\_hca\\_port\\_t](#), which is described in the following table:

**Table 4 VAPI\_hca\_port\_t**

max_mtu	Max MTU supported on this port.
max_msg_sz	Max Message size supported on this port.
lid	Base LID
lmc	LMC
state	Port State
capability_mask	Various bits as defined in PortInfo CapabilityMask in IB spec. (IB_capability_mask_bits_t)
max_vl_num	Maximum number of VL supported by this port.
bad_pkey_counter	Bad PKey counter (if supported).

**Table 4 VAPI\_hca\_port\_t (Continued)**

qkey_viol_counter	QKey violation counter.
sm_lid	Subnet Manager LID on this port.
sm_sl	Subnet Manager SL on this port.
pkey_tbl_len	P-Key table length.
gid_tbl_len	GID table length.



## 2.5 Query HCA Gid Tbl

### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_hca_gid_tbl
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      IB_port_t         port_num,
    IN      u_int16_t         tbl_len_in,
    OUT     u_int16_t         *tbl_len_out,
    OUT     IB_gid_t          *gid_tbl_p
)

```

**ARGUMENTS:**

- hca\_hndl:** HCA object handle.
- port\_num:** Port number.
- tbl\_len\_in:** Number of entries in given gid\_tbl\_p buffer.
- tbl\_len\_out:** Actual number of entries in this port GID table.
- gid\_tbl\_p:** The GID table buffer to return result in.

**RETURNS:**

- VAPI\_OK**
- VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.
- VAPI\_EAGAIN:** tbl\_len\_out > tbl\_len\_in.
- VAPI\_EINVAL\_PARAM:** Invalid port number.

### DESCRIPTION:

The GID table of the given port is returned in gid\_tbl\_p. If tbl\_len\_out (actual number of entries) is more than tbl\_len\_in, the function should be recalled with a larger buffer.

## 2.6 Query HCA Pkey Tbl

### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_hca_pkey_tbl
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      IB_port_t         port_num,
    IN      u_int16_t         tbl_len_in,
    OUT     u_int16_t         *tbl_len_out,
    OUT     VAPI_pkey_t       *pkey_tbl_p
)

```

**ARGUMENTS:**

- hca\_hndl:** HCA object handle.
- port\_num:** Port number.
- tbl\_len\_in:** Number of entries in given gid\_tbl\_p buffer.
- tbl\_len\_out:** Actual number of entries in this port GID table.
- pkey\_tbl\_p:** The PKEY table buffer to return result in.

**RETURNS:**

- VAPI\_OK**
- VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.
- VAPI\_EAGAIN:** tbl\_len\_out > tbl\_len\_in.
- VAPI\_EINVAL\_PARAM:** .Invalid port number.

### DESCRIPTION:

## 2.7 Modify HCA Attributes

### SYNOPSIS:

```
VAPI_ret_t
VAPI_modify_hca_attr
(
    IN      VAPI_hca_hdl_t      hca_hdl,
    IN      IB_port_t          port_num,
    IN      VAPI_hca_attr_t     *hca_attr_p,
    IN      VAPI_hca_attr_mask_t *hca_attr_mask_p
)

```

**ARGUMENTS:** **hca\_hdl:** Handle to HCA.  
**port\_num:** Port number.  
**hca\_attr\_p:** Pointer to the HCA attributes structure.  
**hca\_attr\_mask\_p:** Pointer to the HCA attributes mask

**ARGUMENTS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_COUNTER:** Invalid counter.  
**VAPI\_EINVAL\_COUNT\_VAL:** Invalid counter value.  
**VAPI\_EPERM:** Insufficient permissions to do this operation.

### DESCRIPTION:

Sets the HCA attributes specified in **hca\_attr\_p** to port number **port\_num**. Only the values specified in **hca\_attr\_mask\_p** are modified. **hca\_attr\_p** is a pointer to a structure of type **VAPI\_hca\_attr\_t**, which is specified in the following table:

Table 5 VAPI\_hca\_attr\_t

is_sm	Is Subnet manager
is_snmp_tun_sup	Is SNMP tunneling supported
is_dev_mgt_sup	Is device management supported
is_vendor_cls_sup	Is Vendor Class supported

The attributes to be changed are specified using a mask of type **VAPI\_hca\_attr\_mask\_t**. The mask can be any combination of the following flags:

**Table 6 HCA Attributes Flags**

HCA_ATTR_IS_SM	Is Subnet manager
HCA_ATTR_IS_SNMP_TUN_SUP	Is SNMP tunneling supported
HCA_ATTR_IS_DEV_MGT_SUP	Is device management supported
HCA_ATTR_IS_VENDOR_CLS_SUP	Is Vendor Class supported

To set up the attributes mask, the following macros should be used:

**Table 7 HCA Attributes Mask Macros**

HCA_ATTR_MASK_SET(hca_attr, flag)	set a specific flag in attributes mask
HCA_ATTR_MASK_CLR(hca_attr, flag)	clear a specific flag in attributes mask
HCA_ATTR_MASK_IS_SET(hca_attr, flag)	check if a flag is set in the mask
HCA_ATTR_MASK_CLR_ALL(hca_attr)	clear all flags in the attributes mask.
HCA_ATTR_MASK_SET_ALL(hca_attr)	set all flags in the attributes mask
HCA_ATTR_IS_FLAGS_SET(hca_attr)	check any of the capability_mask bits is set

## 2.8 Release HCA Handle

**Note:** This is an Extended VAPI function

### SYNOPSIS:

```
VAPI_ret_t
EVAPI_release_hca_hndl
(
    IN      VAPI_hca_hndl_t    hca_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** Handle for which process resources are released.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** No such opened HCA.

**DESCRIPTION:** Releases all resources used by this process for an opened HCA. To release all resources, applications *must* use this function.

## 2.9 Close HCA

### SYNOPSIS:

```
VAPI_ret_t
VAPI_close_hca
(
    IN      VAPI_hca_hndl_t    hca_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** Handle to HCA.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

This call deallocates all the structures allocated during the call to [VAPI\\_open\\_hca](#) and any other resource in the domain of the CI.

It is the responsibility of the consumers to free resources allocated for the HCA that are under its scope.

## 2.10 Allocate Protection Domain

### SYNOPSIS:

```
VAPI_ret_t
VAPI_alloc_pd
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    OUT    VAPI_pd_hndl_t     *pd_hndl_p
)

```

### ARGUMENTS:

**hca\_hndl:** Handle to HCA.

**pd\_hndl\_p:** Pointer to handle to Protection Domain object.

### RETURNS:

**VAPI\_OK**

**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.

**VAPI\_EAGAIN:** Insufficient resources.

**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

This call registers a new protection domain.

## 2.11 Deallocate Protection Domain

### SYNOPSIS:

```
VAPI_ret_t
VAPI_dealloc_pd
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_pd_hndl_t    pd_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** Handle to HCA.  
**pd\_hndl:** Handle to Protection Domain Object.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_PD\_HNDL:** Invalid Protection Domain  
**VAPI\_EBUSY:** Protection Domain in use.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:** Deregisters the Protection Domain from the **PDA**. It is the **PDA's** responsibility to validate that there are no objects associated with the Protection Domain being deallocated.



## 2.12 Allocate Reliable Datagram Domain

### SYNOPSIS:

```
VAPI_ret_t
VAPI_alloc_rdd
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    OUT    VAPI_rdd_hndl_t    *rdd_hndl_p
)
)
```

**ARGUMENTS:** **hca\_hndl:** HCA Handle.  
**rdd\_hndl\_p:** Pointer to Reliable Datagram Domain object handle.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EAGAIN:** Out of resources.  
**VAPI\_EINVAL\_RD\_UNSUPPORTED:** RD is not supported.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:**  
Allocates an RD domain.

## 2.13 Deallocate Reliable Datagram Domain

### SYNOPSIS:

```
VAPI_ret_t
VAPI_dealloc_rdd
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_rdd_hndl_t    rdd_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** Handle to HCA.  
**rdd\_hndl:** Reliable Datagram Domain object.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_RDD\_HNDL:** Invalid RDD handle.  
**VAPI\_EAGAIN:** Out of resources.  
**VAPI\_EBUSY:** RDD is busy.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:**  
Deallocates an RD domain.

## 3 Address Management Verbs

This chapter describes the Mellanox implementation of the following verbs handling basic address management functionality:

- Create Address Handle ([p.25](#))
- Modify Address Handle ([p.27](#))
- Query Address Handle ([p.28](#))
- Destroy Address Handle ([p.29](#))

### 3.1 Create Address Handle

#### SYNOPSIS:

```
VAPI_ret_t
VAPI_create_addr_hndl
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_pd_hndl_t    pd_hndl,
    IN      VAPI_ud_av_t      *av_p,
    OUT     VAPI_ud_av_hndl_t  *av_hndl_p
)

```

**ARGUMENTS:** **hca\_hndl:** Handle to HCA.  
**pd\_hndl:** Protection domain handle  
**av\_p:** Pointer to Address Vector structure.  
**av\_hndl\_p:** Handle of Address Vector.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_PD\_HNDL:** Invalid Protection Domain handle.  
**VAPI\_EAGAIN:** Insufficient resources.  
**VAPI\_EPERM:** Insufficient permissions.

#### DESCRIPTION:

Creates a new Address Vector Handle that can be used later when posting a WR to a UD QP.

The fields of the address vector are specified in the following table:

**Table 8 VAPI\_ud\_av\_t**

sl	Service Level
grh_flag	Send GRH flag
dlid	Destination LID
traffic_class	TClass (for global routing)
flow_label	Flow Label (for global routing)
hop_limit	Hop Limit (for global routing)
sgid_index	SGID index in SGID table (for global routing)
dgid	Destination GID
static_rate	Maximum static rate: - 1 - 1 Gb/s - 2 - 2.5 Gb/s - 3 - 10 Gb/s - 4 - 30 Gb/s
src_path_bits	Source Path Bits (with which to create the src LID)
port	Egress port

## 3.2 Modify Address Handle

### SYNOPSIS:

```
VAPI_ret_t
VAPI_modify_addr_hndl
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_ud_av_hndl_t  av_hndl,
    IN      VAPI_ud_av_t      *av_p
)
```

**ARGUMENTS:** **hca\_hndl:** Handle to HCA.  
**av\_hndl :** Handle of Address Vector handle  
**av\_p:** Pointer to Address Vector structure.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_AV\_HNDL:** Invalid Address Vector handle.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:** Modifies existing address vector handle to point to a different new address vector. For address vector fields, refer to [Table 8](#), “VAPI\_ud\_av\_t,” on page 26.

### 3.3 Query Address Handle

#### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_addr_hndl
(
    IN      VAPI_hca_hndl_t      hca_hndl,
    IN      VAPI_ud_av_hndl_t    av_hndl,
    OUT     VAPI_ud_av_t         *av_p
)

```

**ARGUMENTS:** **hca\_hndl:** Handle to HCA.  
**av\_hndl :** Handle of Address Vector .  
**av\_p:** Pointer to Address Vector structure.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_AV\_HNDL:** Invalid address vector handle.  
**VAPI\_EPERM:** Insufficient permission.

**DESCRIPTION:** Returns pointer to ADDR\_VECP with information about the UD Address Vector represented by AddrVecHandle. For address vector fields, refer to [Table 8, “VAPI\\_ud\\_av\\_t,” on page 26.](#)

## 3.4 Destroy Address Handle

### SYNOPSIS:

```
VAPI_ret_t
VAPI_destroy_addr_hndl
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_ud_av_hndl_t  av_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** Handle to HCA.  
**av\_hndl:** Handle to Address Vector

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_AV\_HNDL:** Invalid address vector handle.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:**  
Destroys address vector structure.





# 4 Queue Pair Verbs

This chapter describes the Mellanox implementation of the following verbs handling basic Queue Pair functionality:

- Create Queue Pair (p.31)
- Modify Queue Pair (p.34)
- Query Queue Pair (p.37)
- Destroy Queue Pair (p.38)
- Get Special Queue Pair (p.39)

## 4.1 Create Queue Pair

### SYNOPSIS:

```
VAPI_ret_t
VAPI_create_qp
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_qp_init_attr_t *qp_init_attr_p,
    OUT     VAPI_qp_hndl_t     *qp_hndl_p,
    OUT     VAPI_qp_prop_t     *qp_prop_p
)

```

**ARGUMENTS:** **hca\_hndl** : HCA Handle.  
**qp\_init\_attr\_p**: Pointer to QP attribute to used for initialization.  
**qp\_hndl\_p**: Pointer to returned QP Handle number.  
**qp\_prop\_p**: Pointer to properties of created QP.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN**: Insufficient resources.  
**VAPI\_EINVAL\_HCA\_HNDL**: Invalid HCA handle.  
**VAPI\_EINVAL\_CQ\_HNDL**: Invalid CQ handle.  
**VAPI\_E2BIG\_WR\_NUM**: Number of WR exceeds HCA cap.  
**VAPI\_E2BIG\_SG\_NUM**: Number of SG exceeds HCA cap.  
**VAPI\_EINVAL\_PD\_HNDL**: Invalid protection domain handle.  
**VAPI\_EINVAL\_SERVICE\_TYPE**: Invalid service type for this QP.  
**VAPI\_EINVAL\_RDD**: Invalid RD domain handle.  
**VAPI\_EPERM**: Insufficient permissions.

**DESCRIPTION:**

This call interacts with the **QPM** in order to allocate a **QPO (Queue Pair Object)**. The **QPO** is responsible for the resource allocation and management of the **Queue Pair Context**.

When **VAPI\_create\_qp** is invoked, the QP will be in the reset state. To make the Queue Pair useful to the consumer, it has to modify its state so reception and/or transmission is enabled, e.g., RTS. This is done by calling to **VAPI\_modify\_qp**.

When a QP is opened, the properties specified by **qp\_init\_attr\_p** are applied to the QP. **qp\_init\_attr\_p** type **VAPI\_qp\_init\_attr\_t** and includes the following fields:

**Table 9 VAPI\_qp\_init\_attr\_t**

sq_cq_hndl	CQ handle to be associated with the SQ.
rq_cq_hndl	CQ handle to be associated with the RQ.
cap	QP Capabilities (see <a href="#">Table 10, “VAPI_qp_cap_t,” on page 32</a> ).
rdd_hndl	RD Domain handle.
sq_sig_type	SQ Signaling type: <ul style="list-style-type: none"> <li>- VAPI_SIGNAL_ALL_WR - signal every submitted WR.</li> <li>- VAPI_SIGNAL_REQ_WR - consumer should specify on a WR if it wants signalling</li> </ul>
rq_sig_type	RQ Signaling type <sup>1</sup> .
pd_hndl	Protection Domain handle.
ts_type	Transport Service Type <ul style="list-style-type: none"> <li>- VAPI_TS_RC</li> <li>- VAPI_TS_RD</li> <li>- VAPI_TS_UC</li> <li>- VAPI_TS_UD</li> </ul>

1. Signaling type on RQ is Mellanox specific.

The QP capability structure, **VAPI\_qp\_cap\_t**, contains the following fields:

**Table 10 VAPI\_qp\_cap\_t**

max_oust_wr_sq	Maximum number of outstanding WRs in the SQ.
max_oust_wr_rq	Maximum number of outstanding WRs in the RQ.
max_sg_size_sq	Maximum number of scatter/gather elements in a WR in the SQ.
max_sg_size_rq	Maximum number of scatter/gather elements in a WR in the RQ.

If completed successfully, the verb returns the handle to the QP and the actual properties that were assigned to the QP in `qp_prop_p`, of type `VAPI_qp_prop_t`. The fields of `VAPI_qp_prop_t` are specified in the following table:

**Table 11** `VAPI_qp_prop_t`

[Redacted]	
<code>qp_num</code>	QP number
<code>cap</code>	Actual QP capabilities (see <a href="#">Table 10, “VAPI_qp_cap_t,” on page 32</a> )

## 4.2 Modify Queue Pair

### SYNOPSIS:

```
VAPI_ret_t
VAPI_modify_qp
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_qp_hndl_t    qp_hndl,
    IN      VAPI_qp_attr_t    *qp_attr_p,
    IN      VAPI_qp_attr_mask_t *qp_attr_mask_p,
    OUT     VAPI_qp_cap_t     *qp_cap_p
)
```

**ARGUMENTS:**

- hca\_hndl:** HCA handle.
- qp\_hndl:** QP handle.
- qp\_attr\_p:** Pointer to QP attributes to be modified.
- qp\_attr\_mask\_p:** Pointer to the attributes mask to be modified.
- qp\_cap\_p:** Pointer to QP actual capabilities returned.

### RETURNS:

**VAPI\_OK**  
**VAPI\_EAGAIN:** Out of resources.  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_QP\_HNDL:** Invalid QP handle.  
**VAPI\_ENOSYS\_ATTR:** QP attribute is not supported.  
**VAPI\_EINVAL\_ATTR:** Cannot change QP attribute.  
**VAPI\_ENOSYS\_ATOMIC:** Atomic operation not supported.  
**VAPI\_EINVAL\_PKEY\_IX:** PKey index out of range.  
**VAPI\_EINVAL\_PKEY\_TBL\_ENTRY:** Pkey index points to an invalid entry in pkey table.  
**VAPI\_EINVAL\_QP\_STATE:** Invalid QP state.  
**VAPI\_EINVAL\_RDD\_HNDL:** Invalid RDD domain handle.  
**VAPI\_EINVAL\_MIG\_STATE:** Invalid path migration state.  
**VAPI\_E2BIG\_MTU:** MTU exceeds HCA port capabilities.  
**VAPI\_EINVAL\_PORT:** Invalid port.  
**VAPI\_EINVAL\_SERVICE\_TYPE:** Invalid service type.  
**VAPI\_E2BIG\_WR\_NUM:** Maximum number of WRs requested exceeds HCA capabilities.  
**VAPI\_EINVAL\_RNR\_NAK\_TIMER:** Invalid RNR NAK timer value.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Modifies the QP attributes and transitions the QP to a new state. Note that only a subset of all the attributes can be modified while changing the QP state. The **qp\_attr\_mask\_p** specifies the actual attributes to be modified.

The QP attributes specified are of type **VAPI\_qp\_attr\_t** and are specified in the following table:

**Table 12 VAPI\_qp\_attr\_t**

qp_state	Next QP State: RST, INIT, RTR, RTS, SQD, SQER, ERR	Always
en_sqd_asyn_notif	Enable SQD asynchronous affiliated event notification.	Any->SQD
qp_num	QP Number. This is relevant only for query_qp, and not modify_qp.	
remote_atomic_flags	Enable/Disable RDMA and Atomic Operations: - VAPI_EN_REM_WRITE - enable RDMA WR - VAPI_EN_REM_READ - enable RDMA RD - VAPI_EN_REM_ATOMIC_OP - enable RDMA atomic operations	RST->INIT INIT->RTR RTR->RTS RTS->RTS SQER->RTS
pkey_ix	PKey Index	RST->INIT INIT->RTR
port	Physical Port	RST->INIT
qkey	QKey (for unconnected service types)	RST->INIT INIT->RTR RTR->RTS RTS->RTS SQER->RTS SQD->RTS
path_mtu		
av	Remote Node Address Vector (connected service types), see <a href="#">Section Table 8“VAPI_ud_av_t,” on page 26.</a>	INIT->RTR SQD->RTS
timeout	Timeout (RC only)	
retry_count	Retry count (RC only)	
rnr_retry	RNR retry count (RC only)	
rq_psn	RQ PSN	INIT->RTR
qp_ous_rd_atom	Number of responder resources for RDMA read/atomic	INIT->RTR SQD->RTS
alt_av	Alternate Destination Node Address Vector (RC, UC only), see <a href="#">Section Table 8“VAPI_ud_av_t,” on page 26.</a>	INIT->RTR RTR->RTS RTS->RTS SQD->RTS
alt_timeout	Timeout (RC only)	
alt_retry_count	Retry count (RC only)	
alt_rnr_retry	RNR retry count (RC only)	
alt_pkey_ix	Alternate Path PKey index	
min_rnr_timer	Minimum RNR NAK timer field (RC only)	INIT->RTR RTR->RTS RTS->RTS SQER->RTS SQD->RTS

**Table 12 VAPI\_qp\_attr\_t (Continued)**

timeout	Timeout (RC only)	RTR->RTS SQD->RTS
retry_count	Retry Count (RC Only)	RTR->RTS SQD->RTS
rnr_retry	RNR Retry Count (RC only)	RTR->RTS
sq_psn	SQ PSN	RTR->RTS
ous_dst_rd_atom	Number of outstanding RDMA read/atomic operations at destination.	RTR->RTS SQD->RTS
path_mig_state	Path migration state: - VAPI_MIGRATED - VAPI_REARM	RTR->RTS RTS->RTS SQER->RTS SQD->RTS
cap	Required QP capabilities (only max_sq_ous_wr, max_rq_ous_wr are valid). For more detail, see <a href="#">Table 10, “VAPI_qp_cap_t,” on page 32</a> .	
dest_qp_num	Destination QP number.	

When successful, the output of this verb is the actual QP capability of type **VAPI\_qp\_cap\_t** (for more detail, see [Table 10, “VAPI\\_qp\\_cap\\_t,” on page 32](#)). Only the **max\_sq\_ous\_wr** and **max\_rq\_ous\_wr** fields are valid.

The QP attributes to be changed are specified using a mask of type **VAPI\_qp\_attr\_mask\_t**. The mask can be any combination of the flags, which are denoted by the **HCA\_ATTR\_** prefix and the parameter name listed in [Table 12, “VAPI\\_qp\\_attr\\_t,” on page 35](#) in capital letters (for example, **HCA\_ATTR\_QP\_STATE**, **HCA\_ATTR\_SQD\_ASYNC\_NOTIF**, **HCA\_ATTR\_REMOTE\_ATOMIC\_FLAGS**, etc.).

To set up the attributes mask, the following macros should be used:

**Table 13 QP Attributes Mask Macros**

QP_ATTR_MASK_SET(hca_attr, flag)	Set a specific flag in the attributes mask.
QP_ATTR_MASK_CLR(hca_attr, flag)	Clear a specific flag in the attributes mask.
QP_ATTR_MASK_IS_SET(hca_attr, flag)	Check whether a flag is set in the mask.
QP_ATTR_MASK_CLR_ALL(hca_attr)	Clear all flags in the attributes mask.
QP_ATTR_MASK_SET_ALL(hca_attr)	Set all flags in the attributes mask.

## 4.3 Query Queue Pair

### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_qp
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_qp_hndl_t    qp_hndl,
    OUT     VAPI_qp_attr_t     *qp_attr_p,
    OUT     VAPI_qp_attr_mask_t *qp_attr_mask_p,
    OUT     VAPI_qp_init_attr_t *qp_init_attr_p
)

```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**qp\_hndl:** QP handle.  
**qp\_attr\_p:** Pointer to QP attributes.  
**qp\_attr\_mask\_p:** Pointer to QP attributes mask.  
**qp\_init\_attr\_p:** Pointer to init attributes

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_QP\_HNDL:** Invalid QP handle.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Returns a `VAPI_qp_attr_t` structure to the application with all the relevant information that applies to the QP matching `qp_hndl`, see [Table 12, “VAPI\\_qp\\_attr\\_t,” on page 35](#). The verb also returns a pointer to `VAPI_qp_init_attr_t`, which includes further information of the QP init attributes, see [Table 9, “VAPI\\_qp\\_init\\_attr\\_t,” on page 32](#).

Note that only the relevant fields in `qp_attr_p` and `qp_init_attr_p` are valid. The valid fields in `qp_attr_p` are marked in the mask returned by `qp_attr_mask_p`. For further description of the `VAPI_qp_attr_mask_t`, see [Section “VAPI\\_modify\\_qp,” on page 34](#).

## 4.4 Destroy Queue Pair

### SYNOPSIS:

```
VAPI_destroy_qp
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_qp_hndl_t    qp_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**qp\_hndl:** QP handle.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_QP\_HNDL:** Invalid QP handle.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Releases all the resources allocated by the CI for this QP.

It is the responsibility of the consumers to release all the resources allocated for all the WRs posted to the QP that are no longer under the responsibility of the CI.



## 4.5 Get Special QP

### SYNOPSIS:

```
VAPI_ret_t
VAPI_get_special_qp
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      IB_port_t         port,
    IN      VAPI_special_qp_t qp,
    IN      VAPI_qp_init_attr_t *qp_init_attr_p,
    OUT     VAPI_qp_hndl_t     *qp_hndl_p,
    OUT     VAPI_qp_cap_t      *qp_cap_p
)
```

**ARGUMENTS:**

- hca\_hndl:** HCA Handle.
- phy\_port:** Physical port (valid only for QP0 and QP1).
- qp:** QP type.
- qp\_init\_attr\_p:** Pointer to init attributes structure.
- qp\_hndl\_p:** Pointer to the QP handle.
- qp\_cap\_p:** Pointer to the actual QP capabilities.

**RETURNS:**

- VAPI\_OK**
- VAPI\_EAGAIN:** Out of resources.
- VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.
- VAPI\_EINVAL\_QP\_TYPE:** Invalid special QP type.
- VAPI\_EBUSY:** QP already in use (GSI, SMI QPs only).
- VAPI\_E2BIG\_RAW\_DGRAM\_NUM:** number of raw. datagram QPs exceeded.
- VAPI\_EINVAL\_PORT:** Invalid port number.
- VAPI\_EINVAL\_CQ\_HNDL:** Invalid CQ handle
- VAPI\_E2BIG\_WR\_NUM:** Maximum number of work requests exceeds HCA capabilities.
- VAPI\_E2BIG\_SG\_NUM:** Maximum number of scatter/gather elements exceeds HCA capabilities.
- VAPI\_EINVAL\_PD\_HNDL:** Invalid protection domain handle.
- VAPI\_ENOSYS\_RAW:** Raw datagram QPs are not supported.
- VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Creates a special QP that can be used to generate MADs, RAW IPV6 or Ethertype packets.

The valid QP types are (defined in VAPI\_special\_qp\_t):

- VAPI\_SMI\_QP - QP 0
- VAPI\_GSI\_QP - QP 1
- VAPI\_RAW\_ETY\_QP - Raw ethertype QP – not supported

- **VAPI\_RAW\_IPV6\_QP** - Raw IPv6 QP

The physical port is valid only for QP 0 and QP 1, and is otherwise ignored.

The **VAPI\_qp\_init\_attr\_t** is described in [Table 9](#), “[VAPI\\_qp\\_init\\_attr\\_t](#),” on page 32. Note that the **rdd\_hndl** and the transport service (**ts\_type**) are not used.

Upon successful completion, the QP returns the QP handle and the actual QP capabilities. The **VAPI\_qp\_cap\_t** is described in [Table 10](#), “[VAPI\\_qp\\_cap\\_t](#),” on page 32.

# 5 Completion Queue Verbs

This chapter describes the Mellanox implementation of the following verbs handling completion queue functionality:

- Create Completion Queue ([p.41](#))
- Query Completion Queue ([p.43](#))
- Resize Completion Queue ([p.44](#))
- Destroy Completion Queue ([p.45](#))

## 5.1 Create Completion Queue

### SYNOPSIS:

```
VAPI_ret_t
VAPI_create_cq
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_cqe_num_t    cqe_num,
    OUT     VAPI_cq_hndl_t    *cq_hndl_p,
    OUT     VAPI_cqe_num_t    *num_of_entries_p
)

```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**cqe\_num:** Minimum number of entries required in CQ.  
**cq\_hndl\_p:** Pointer to the created CQ handle.  
**num\_of\_entries\_p:** Actual number of entries in CQ.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN:** Out of resources.  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_E2BIG\_CQ\_NUM:** Number of entries in CQ exceeds HCA capabilities.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Allocates the required data structures for administration of a completion queue, including completion queue buffer space sufficient for the maximum number of entries in the completion.

Completion queue entries are accessed directly by the application.

## 5.2 Query Completion Queue

### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_cq
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_cq_hndl_t    cq_hndl,
    OUT     VAPI_cqe_num_t     *num_of_entries_p
)

```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**cq\_hndl:** Completion Queue handle.  
**num\_of\_entries\_p:** Pointer to actual number of entries in CQ.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle  
**VAPI\_EINVAL\_CQ\_HNDL:** Invalid CQ handle.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:** Retrieves the number of entries in the CQ.

## 5.3 Resize Completion Queue

**Note:** not currently supported.

### SYNOPSIS:

```
VAPI_ret_t
VAPI_resize_cq
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_cq_hndl_t    cq_hndl,
    IN      VAPI_cqe_num_t    cqe_num,
    OUT     VAPI_cqe_num_t    *num_of_entries_p
)

```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**cq\_hndl:** CQ handle.  
**cqe\_num:** Minimum number of entries required in CQ.  
**num\_of\_entries\_p:** Pointer to actual number of entries in CQ.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN:** out of resources  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_CQ\_HNDL:** Invalid CQ handle.  
**VAPI\_E2BIG\_CQ\_NUM:** Number of entries in CQ exceeds HCA capabilities.  
**VAPI\_E2BIG\_OUS\_ENT:** Number of current outstanding entries in CQ exceeds required size.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:** Requests the CEM to increase or decrease the size of the buffer that holds the completion queue entries for a specific CQ.

## 5.4 Destroy Completion Queue

### SYNOPSIS:

```
VAPI_ret_t
VAPI_destroy_cq
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_cq_hndl_t    cq_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**cq\_hndl:** CQ handle.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_CQ\_HNDL:** Invalid CQ handle.  
**VAPI\_EBUSY:** one or more work queues are still associated with this CQ.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Destroys a CQ and releases all resources associated with it.

Destruction of the completion is not allowed if there are any QPs still associated to the CQ. The CQM holds a list of QP's associated with each completion queue.





# 6 EE Context Verbs

This chapter describes the Mellanox implementation of the following verbs handling basic address management functionality:

- Create EE Context ([p.47](#))
- Modify EE Context Attributes ([p.48](#))
- Query EE Context ([p.49](#))
- Destroy EE Context ([p.50](#))

## 6.1 Create EE Context

**Note:** not currently supported.

### SYNOPSIS:

```
VAPI_ret_t
VAPI_create_eec
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_rdd_t         rdd,
    OUT     VAPI_eec_hndl_t    *eec_hndl_p
)

```

**ARGUMENTS:** **hca\_hndl** : HCA handle.  
**rdd**: RD domain.  
**eec\_hndl\_p**: Pointer to EE Context Handle.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN**: out of resources.  
**VAPI\_EINVAL\_HCA\_HNDL**: Invalid HCA handle.  
**VAPI\_EPERM**: Insufficient permissions.

**DESCRIPTION:**  
 Creates an EE context.

## 6.2 Modify EE Context Attributes

**Note:** not currently supported.

### SYNOPSIS:

```
VAPI_ret_t
VAPI_modify_eec_attr
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_eec_hndl_t    eec_hndl,
    IN      VAPI_eec_attr_t    *eec_attr_p
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**eec\_hndl:** EE Context handle  
**eec\_attr\_p:** Pointer to EE Context Attributes structure.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN:** out of resources.  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_EEC\_HNDL:** Invalid EEC handle.  
**VAPI\_EINVAL\_EEC\_STATE:** Invalid EEC state.  
**VAPI\_EINVAL\_RDD:** Invalid RD domain.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:**  
 Modifies attributes of an EE context.

## 6.3 Query EE Context

**Note:** not currently supported.

### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_eec_attr
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_eec_hndl_t    eec_hndl,
    OUT     VAPI_eec_attr_t    *eec_attr_p
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**eec\_hndl:** EE context handle.  
**eec\_attr\_p:** Pointer to EE Context Attributes structure.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_EEC\_HNDL:** Invalid EEC handle.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:**  
Returns attributes of an EE context.

## 6.4 Destroy EE Context

**Note:** not currently supported.

### SYNOPSIS:

```
VAPI_ret_t
VAPI_destroy_eec
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_eec_hndl_t    eec_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**eec\_hndl:** EE context handle.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_EEC\_HNDL:** Invalid EEC handle.  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:**  
Destroys an EE context.

# 7 Memory Management Verbs

This chapter describes the Mellanox implementation of the following verbs handling basic memory management functionality:

- Register Memory Region ([p.51](#))
- Query Memory Region ([p.53](#))
- Deregister Memory Region ([p.54](#))
- Reregister Memory Region ([p.55](#))
- Register Shared Memory Region ([p.56](#))
- Allocate Memory Window ([p.57](#))
- Query Memory Window ([p.58](#))
- Bind Memory Window ([p.59](#))
- Deallocate Memory Window ([p.60](#))

## 7.1 Register Memory Region

### SYNOPSIS:

```
VAPI_ret_t
VAPI_register_mr
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_mrw_t         *req_mrw_p,
    OUT     VAPI_mr_hndl_t     *mr_hndl_p,
    OUT     VAPI_mrw_t         *rep_mrw_p
)

```

**ARGUMENTS:** **hca\_hndl** : HCA handle.  
**req\_mrw\_p**: Pointer to the requested memory region properties.  
**mr\_hndl\_p**: Pointer to the memory region handle.  
**rep\_mrw**: Pointer to the responded memory region properties.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN**: out of resources.  
**VAPI\_EINVAL\_HCA\_HNDL**: Invalid HCA handle.  
**VAPI\_EINVAL\_PD\_HNDL**: Invalid PD handle.  
**VAPI\_EINVAL\_VA**: Invalid virtual address.  
**VAPI\_EINVAL\_LEN**: Invalid length.

**VAPI\_EINVAL\_ACL:** Invalid ACL specifier.

**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Registers a memory region.

The caller should fill the **req\_mrwr\_p** structure fields with the **type**, **virtual start address**, **size**, protection domain handle (**pd\_hdl**) and the access control list (**acl**). Upon successful completion, the **rep\_mrwr\_p** will include the **l\_key**, the **r\_key** (if remote access was requested). The memory region handle is returned in **mr\_hdl\_p**. Note that the **type** in **req\_mrwr\_p** should be initialized with **MR** (Memory Region). The **VAPI\_mrwr\_t** is described in the following table:

**Table 14 VAPI\_mrwr\_t**

FIELD	DESCRIPTION
type	Memory Region Type: <ul style="list-style-type: none"> <li>- VAPI_MR = Memory Region</li> <li>- VAPI_MW = Memory Windows</li> <li>- VAPI_MPR = Memory Physical Region</li> <li>- VAPI_MSHAR = Memory Shared Region</li> </ul>
l_key	For regions: Local protection key. For windows: l_key of region to bind to.
r_key	Remote protection key given for this window/region
start	Start address of region/window (64 bit address). <sup>1</sup> For physical region, this is the start IOVA.
size	Length of the region/window (64 bits). <sup>2</sup>
pd_hdl	Protection domain handle (regions only)
acl	Region access control list. <ul style="list-style-type: none"> <li>- VAPI_EN_LOCAL_WRITE</li> <li>- VAPI_EN_REMOTE_WRITE</li> <li>- VAPI_EN_REMOTE_READ</li> <li>- VAPI_EN_REMOTE_ATOM</li> <li>- VAPI_EN_MEMREG_BIND</li> </ul>
pbuf_list_len	Physical buffers list (pbuf_list_p) length. For physical memory region only (type==VAPI_MPR)
*pbuf_list_p	Physical buffers addresses list. For physical memory region only (type==VAPI_MPR)
iova_offset	Offset of "start" in first buffer. For physical memory region only (type==VAPI_MPR)

1. May differ from requested address by aligning down to page boundary.
2. May differ from requested address by aligning up to page boundary.

## 7.2 Query Memory Region

### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_mr
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_mr_hndl_t     mr_hndl,
    OUT     VAPI_mrwr_t        *rep_mrwr_p,
    OUT     VAPI_virt_addr_t    *remote_start_p,
    OUT     VAPI_virt_addr_t    *remote_size_p
)

```

**ARGUMENTS:**

- hca\_hndl:** HCA handle.
- mr\_hndl:** Memory Region handle.
- rep\_mrwr\_p:** Pointer to Memory Region attributes
- remote\_start\_p:** Pointer to returned value of the remote start address.
- remote\_size\_p:** Pointer to returned value of the size of the remote region.

**RETURNS:**

- VAPI\_OK**
- VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.
- VAPI\_EINVAL\_MR\_HNDL:** Invalid Memory Region handle.
- VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Queries a memory region handle and returns a **VAPI\_mrwr\_t**, which includes all the memory region properties: protection domain handle, ACL, LKey, RKey and actual protection bounds. The protection bounds returned in **rep\_mrwr\_p** are the **local** protection bounds enforced by the HCA. The **remote** protection bounds are returned in **remote\_start\_p** and **remote\_size\_p** and are valid only when remote access is requested.

## 7.3 Deregister Memory Region

### SYNOPSIS:

```
VAPI_ret_t
VAPI_deregister_mr
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_mr_hndl_t    mr_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**mr\_hndl:** Memory Region Handle

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle  
**VAPI\_EINVAL\_MR\_HNDL:** Invalid memory region handle  
**VAPI\_EBUSY:** memory region still has bound window(s)  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:** Destroys a registered memory region. The memory region deregistering has to be invalidated from the CI.



## 7.4 Reregister Memory Region

**Note:** not currently supported.

### SYNOPSIS:

```
VAPI_ret_t
VAPI_reregister_mr
(
    IN      VAPI_hca_hndl_t      hca_hndl,
    IN      VAPI_mr_hndl_t      mr_hndl,
    IN      VAPI_mr_change_t    change_type,
    IN      VAPI_mrwr_t         *req_mrwr_p,
    OUT     VAPI_mr_hndl_t      *rep_mr_hndl_p,
    OUT     VAPI_mrwr_t         *rep_mrwr_p
)
```

**ARGUMENTS:**

- hca\_hndl:** HCA handle.
- mr\_hndl:** Old Memory Region handle.
- change\_type:** Requested change type.
- req\_mrwr\_p:** Pointer to the requested memory region properties.
- rep\_mr\_hndl\_p:** Pointer to the returned new memory region handle.
- rep\_mrwr\_p:** Pointer to the returned memory region properties.

### RETURNS: VAPI\_OK

- VAPI\_EAGAIN:** out of resources.
- VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.
- VAPI\_EINVAL\_MR\_HNDL:** Invalid memory region handle.
- VAPI\_EINVAL\_VA:** Invalid virtual address.
- VAPI\_EINVAL\_LEN:** Invalid length.
- VAPI\_EINVAL\_PD\_HNDL:** Invalid protection domain handle.
- VAPI\_EINVAL\_ACL:** Invalid ACL specifier.
- VAPI\_EBUSY:** memory region still has bound window(s).
- VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Reregisters the memory region associated with the **mr\_hndl**.

## 7.5 Register Shared Memory Region

### SYNOPSIS:

```
VAPI_ret_t
VAPI_register_smr
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_mr_hndl_t    orig_mr_hndl,
    IN      VAPI_mrwr_t        *req_mrwr_p,
    OUT     VAPI_mr_hndl_t    *mr_hndl_p,
    OUT     VAPI_mrwr_t        *rep_mrwr_p
)

```

**ARGUMENTS:**

- hca\_hndl** : HCA handle.
- orig\_mr\_hndl**: Original memory region handle.
- req\_mrwr\_p**: Pointer to the requested memory region properties.
- mr\_hndl\_p**: Pointer to the returned memory region handle.
- rep\_mrwr**: Pointer to the returned memory region properties.

**RETURNS:**

- VAPI\_OK**
- VAPI\_EAGAIN**: out of resources.
- VAPI\_EINVAL\_HCA\_HNDL**: Invalid HCA handle.
- VAPI\_EINVAL\_MR\_HNDL**: Invalid MR handle.
- VAPI\_EINVAL\_PD\_HNDL**: Invalid PD handle.
- VAPI\_EINVAL\_ACL**: Invalid ACL specifier.
- VAPI\_EPERM**: Insufficient permissions.

**DESCRIPTION:**

Registers a shared memory region associated with the physical buffers of an existing memory region referenced by **orig\_mr\_hndl**. The **req\_mrwr\_p** is a pointer to the requested memory region properties and is of type **VAPI\_mrwr\_t**. The requested type should be set to **MSHAR** (shared memory region), and the struct should contain the requested start virtual address (**start** field), the protection domain handle and the ACL.

## 7.6 Allocate Memory Window

**Note:** not currently supported.

### SYNOPSIS:

```
VAPI_ret_t
VAPI_alloc_mw
(
    IN      VAPI_hca_hdl_t      hca_hdl,
    IN      VAPI_pd_hdl_t      pd,
    OUT     VAPI_mw_hdl_t      *mw_hdl_p,
    OUT     VAPI_rkey_t         *rkey_p
)
```

**ARGUMENTS:** **hca\_hdl:** HCA handle.  
**pd:** Protection Domain handle.  
**mw\_hdl\_p:** Pointer to new allocated windows handle.  
**rkey\_p:** Pointer to windows unbounded Rkey

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN**  
**VAPI\_EINVAL\_HCA\_HNDL**  
**VAPI\_EINVAL\_PD\_HNDL**  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:** Allocates an MWO object that later can be bound to an RKey.

## 7.7 Query Memory Window

**Note:** not currently supported.

### SYNOPSIS:

```
VAPI_ret_t
VAPI_query_mw
(
    IN      VAPI_hca_hdl_t      hca_hdl,
    IN      VAPI_mw_hdl_t      mw_hdl,
    OUT     VAPI_rkey_t         *rkey_p,
    OUT     VAPI_pd_hdl_t      *pd_p
)
```

**ARGUMENTS:** **hca\_hdl:** HCA handle.  
**mw\_hdl:** Windows handle.  
**\*rkey\_p:** Pointer to unbound Rkey of window.  
**\*pd\_p:** Pointer to Protection Domain handle of window.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL**  
**VAPI\_EINVAL\_PD\_HNDL**  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:**  
Returns the current PD associated with the memory domain retrieved from the PDA (no access to HW required).

## 7.8 Bind Memory Window

**Note:** not currently supported.

### SYNOPSIS:

```
VAPI_ret_t
VAPI_bind_mw
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_mw_hndl_t    mw_hndl,
    IN      const VAPI_mrw_t   *bind_prop_p,
    IN      VAPI_qp_hndl_t    qp_hndl,
    IN      VAPI_wr_id_t       id,
    IN      VAPI_comp_type_t   comp_type,
    OUT     VAPI_rkey_t        *new_rkey_p
)
```

**ARGUMENTS:**

- hca\_hndl:** HCA handle.
- mw\_hndl:** Handle of memory window.
- bind\_prop\_p:** Pointer to binding properties.
- qp\_hndl:** QP handle for binding this window.
- id:** WQE ID for binding operation (to track on completion).
- comp\_type:** Completion type for binding WQE.
- new\_rkey\_p:** Pointer to RKey of bound window.

**RETURNS:**

- VAPI\_OK
- OUT\_OF\_RESOURCES
- INVALID\_HCA\_hndl
- INVALID\_QP\_hndl
- INVALID\_MEMORY\_WINDOWS\_hndl
- INVALID\_R\_KEY
- INVALID\_MEMORY\_REGION\_hndl
- INVALID\_L\_KEY
- INVALID\_VIRTUAL\_ADRESS
- INVALID\_LENGTH
- INVALID\_ACCESS\_REQUEST
- INVALID\_COMP\_REQUEST\_STATUS
- WORK\_REQUEST\_OK
- WORK\_REQUEST\_PROTECTION\_ERROR
- VAPI\_EPERM: Insufficient permissions.

### DESCRIPTION:

Binds a memory window. This call is performed entirely in user mode. The posted descriptor returns an RKey that can be used in subsequent remote access to the bound memory region.

## 7.9 Deallocate Memory Window

**Note:** not currently supported.

### SYNOPSIS:

### SYNOPSIS:

```
VAPI_ret_t
VAPI_dealloc_mw
(
    IN      VAPI_hca_hdl_t    hca_hdl,
    IN      VAPI_mw_hdl_t    mw_hdl
)
```

**ARGUMENTS:** **HCAHndl:** HCA handle.  
**MEMWWINHandle:** Memory Windows handle.

**RETURNS:** **VAPI\_OK**  
**INVALID\_HCA\_hdl**  
**INVALID\_MEMORY\_WINDOWS\_hdl**  
**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:**  
Destroys memory windows.

# 8 Multicast Verbs

This chapter describes the Mellanox implementation of the following verbs handling basic multicast functionality:

- Attach Qp to Multicast Group ([p.61](#))
- Detach QP from Multicast Group ([p.62](#))

## 8.1 Attach QP to Multicast Group

### SYNOPSIS:

```
VAPI_ret_t
VAPI_attach_to_multicast
(
    IN  VAPI_hca_hndl_t  hca_hndl,
    IN  IB_gid_t         mcg_dgid,
    IN  VAPI_qp_hndl_t   qp_hndl,
    IN  IB_lid_t         mcg_dlid
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**mcg\_dgid:** DGID of multicast group.  
**qp\_hndl:** QP handle.  
**mcg\_dlid:** DLID of multicast group - currently ignored

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN:** out of resources.  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_QP\_HNDL:** Invalid QP handle.  
**VAPI\_E2BIG\_MCG\_SIZE:** Too many QPs attached to multicast group.  
**VAPI\_EINVAL\_MCG\_GID:** Invalid multicast DGID.  
**VAPI\_EINVAL\_SERVICE\_TYPY:** Invalid Service Type for this QP.

**DESCRIPTION:** Attaches QP to multicast group.

## 8.2 Detach QP from Multicast Group

### SYNOPSIS:

```
VAPI_ret_t
VAPI_detach_from_multicast
(
    IN  VAPI_hca_hndl_t  hca_hndl,
    IN  IB_gid_t         mcg_dgid,
    IN  VAPI_qp_hndl_t   qp_hndl,
    IN  IB_lid_t         mcg_dlid
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**mcg\_dgid:** DGID of multicast group.  
**qp\_hndl:** QP handle.  
**mcg\_dlid:** DLID of multicast group - currently ignored

**RETURNS:** **VAPI\_OK**  
**VAPI\_EAGAIN:** Out of resources.  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_QP\_HNDL:** Invalid QP handle.  
**VAPI\_E2BIG\_MCG\_SIZE:** Number of QPs attached to multicast groups exceeded.  
**VAPI\_EINVAL\_MCG\_GID:** Invalid multicast DGID.  
**VAPI\_EINVAL\_SERVICE\_TYPY:** Invalid Service Type for this QP.

**DESCRIPTION:**  
 Detaches QP from a multicast group.



# 9 Work Request Verbs

This chapter describes the Mellanox implementation of the following verbs handling work requests:

- Post Send request ([p.63](#))
- Post Receive Request ([p.66](#))
- Poll for Completion ([p.67](#))
- Request Completion Notification ([p.69](#))
- Request Completion Notification for a Completion Queue ([p.70](#))
- Clear Completion Notification for a Completion Queue ([p.71](#))

## 9.1 Post Send Request

### SYNOPSIS:

```
VAPI_ret_t
VAPI_post_sr
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_qp_hndl_t    qp_hndl,
    IN      VAPI_sr_desc_t    *sr_desc_p
)

```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**qp\_hndl:** QP handle.  
**sr\_desc\_p:** Pointer to the send request descriptor attributes structure.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_QP\_HNDL:** Invalid QP handle.  
**VAPI\_E2BIG\_WR\_NUM:** Too many posted work requests.  
**VAPI\_EINVAL\_OP:** Invalid operation.  
**VAPI\_EINVAL\_QP\_STATE:** Invalid QP state.  
**VAPI\_EINVAL\_NOTIF\_TYPE:** Invalid completion notification type.  
**VAPI\_EINVAL\_SG\_FMT:** Invalid scatter/gather list format.  
**VAPI\_EINVAL\_SG\_NUM:** Invalid scatter/gather list length.  
**VAPI\_ENOSYS\_ATOMIC:** Atomic operations not supported.

**VAPI\_EINVAL\_AH:** Invalid address handle.

**VAPI\_EPERM:** Insufficient permissions.

**DESCRIPTION:**

Posts a send queue work request, the properties of which are specified in the structure pointed to by **sr\_desc\_p**, which is of type **VAPI\_sr\_desc\_t**:

**Table 15 VAPI\_sr\_desc\_t**

id	User defined work request ID
opcode	Transaction OPCODE: <ul style="list-style-type: none"> <li>- RDMA_WRITE</li> <li>- RDAM_WRITE_WITH_IMM</li> <li>- SEND</li> <li>- SEND_WITH_IMM</li> <li>- RDMA_READ</li> <li>- ATOMIC_CMP_AND_SWAP</li> <li>- ATOMIC_FETCH_AND_ADD</li> </ul>
comp_type	Set if completion will be generated: <ul style="list-style-type: none"> <li>- SIGNALLED</li> <li>- UNSIGNALLED (default)</li> </ul>
sg_lst_p	Pointer to the gather list. See <a href="#">Table 16, “VAPI_sg_lst_entry_t,”</a> on page 65.
sg_lst_len	Length of the gather list.
imm_data	dword immediate data.
fence	Complete execution before next descriptor is executed.
remote_ah	Remote node address handle (valid only for datagrams).
remote_qp	Remote node QP (valid only for datagrams).
remote_qkey	Remote node QKey (valid only for datagrams).
ethertype	Ethertype associated with the WR (valid only for RAW datagrams).
eec_hndl	EEC handle (valid only for UD).
set_se	Set solicited bit (SE) on last packet of message: <ul style="list-style-type: none"> <li>- true</li> <li>- false - default</li> </ul>
remote_addr	Remote Address (valid only for RDMA write).
r_key	Remote Address Rkey (valid only for RDMA write).

**Table 15 VAPI\_sr\_desc\_t (Continued)**

operand1	First 64 bit operand for atomic operation.
operand2	Second 64 bit operand for atomic operation.
local_data_seg	A local data segment.
if_gen	Generate event upon description completion: - DO_EVENT - NO_EVENT (default)
ack_req	Set ack-required bit in last packet: - SET_ACK - NO_ACK

The **sg\_lst\_p** points to a gather list (of length **sg\_lst\_len**), which is an array of local buffers used as the source of the data to be transmitted in this Work Request. Each entry in this array has the following format.

**Table 16 VAPI\_sg\_lst\_entry\_t**

src_addr	64 bit source data address
len	Length of buffer
l_key	Lkey used by the CI when accessing this array.

## 9.2 Post Receive Request

### SYNOPSIS:

```
VAPI_ret_t
VAPI_post_rr
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_QP_hndl_t    qp_hndl,
    IN      VAPI_rr_desc_t    *rr_desc_p
)

```

**ARGUMENTS:**

- hca\_hndl:** HCA handle.
- qp\_hndl:** QP handle.
- rr\_desc\_p:** Pointer to the receive request descriptor attributes structure.

**RETURNS:**

- VAPI\_OK**
- VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.
- VAPI\_EINVAL\_QP\_HNDL:** Invalid QP handle.
- VAPI\_E2BIG\_WR\_NUM:** Too many posted work requests.
- VAPI\_EINVAL\_OP:** Invalid operation.
- VAPI\_EINVAL\_QP\_STATE:** Invalid QP state.
- VAPI\_EINVAL\_SG\_FMT:** Invalid scatter/gather list format.
- VAPI\_EINVAL\_SG\_NUM:** Invalid scatter/gather list length.
- VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Posts a receive request descriptor to the receive queue. The receive request descriptor is of type [VAPI\\_rr\\_desc\\_t](#) and is described in the following table:

**Table 17 VAPI\_rr\_desc\_t**

id	User defined work request ID.
opcode	Transaction OPCODE: - VAPI_RECEIVE
comp_type	Set if completion will be generated: <sup>1</sup> - VAPI_SIGNALED - VAPI_UNSIGNALED (default)
sg_lst_p	Pointer to the scatter list. See <a href="#">Table 16</a> , “ <a href="#">VAPI_sg_lst_entry_t</a> ,” on page 65.
sg_lst_len	Length of the scatter list.

1. This is Mellanox specific

## 9.3 Poll for Completion

### SYNOPSIS:

```
VAPI_ret_t
VAPI_poll_cq
(
    IN      VAPI_hca_hndl_t    hca_hndl,
    IN      VAPI_cq_hndl_t    cq_hndl,
    OUT     VAPI_wc_desc_t     *comp_desc_p
)

```

**ARGUMENTS:** **hca\_hndl:** Handle to HCA.  
**cq\_hndl:** CQ handle.  
**\*comp\_desc\_p:** Pointer to work completion descriptor structure.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_CQ\_HNDL:** Invalid CQ handle.  
**VAPI\_EAGAIN:** CQ is empty.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Retrieves an ICQE (Independent Completion Queue Entry), which is a device independent structure used to retrieve completion status of a WR posted to the Send/Receive Queue including [VAPI\\_bind\\_mw](#).

The verb retrieves a completion queue entry into the descriptor pointed by **wc\_desc\_p**, which is of type [VAPI\\_wc\\_desc\\_t](#) and described in the following table:

**Table 18 VAPI\_wc\_desc\_t**

id	User defined work request ID
opcode	Transaction OPCODE: <ul style="list-style-type: none"> <li>- VAPI_RDMA_WRITE</li> <li>- VAPI_RDAM_WRITE_WITH_IMM</li> <li>- VAPI_SEND</li> <li>- VAPI_SEND_WITH_IMM</li> <li>- VAPI_RDMA_READ</li> <li>- VAPI_ATOMIC_CMP_AND_SWAP</li> <li>- VAPI_ATOMIC_FETCH_AND_ADD</li> <li>- VAPI_MEM_WND_BIND</li> <li>- VAPI_SEND_DATA_RCV (for post_receive requests)</li> <li>- VAPI_RDMA_DATA_RCV (for RDMA with imm)</li> </ul>
len	The actual number of bytes transferred <sup>1</sup> .

**Table 18 VAPI\_wc\_desc\_t (Continued)**

imm_data	dword immediate data (valid only when <b>RDMA_DATA_RCV</b> is indicated in the opcode field).
remote_node_addr	Remote node address (datagram services only), see <a href="#">Table 19, “VAPI_remote_node_addr_t,”</a> on page 68.
grh_flag	GRH present indication.
pkey_ix	PKey index.
status	Completion status: <ul style="list-style-type: none"> <li>- VAPI_SUCCESS</li> <li>- VAPI_LOC_LEN_ERR - local length error</li> <li>- VAPI_LOC_OP_ERR - local operation error</li> <li>- VAPI_LOC_PROT_ERR - local protection error</li> <li>- VAPI_WR_FLUSH_ERR - WR flushed error</li> <li>- VAPI_MW_BIND_ERR - memory window bind error</li> <li>- VAPI_REM_INV_REQ_ERR - remote invalid request error</li> <li>- VAPI_REM_ACCESS_ERR - remote access error</li> <li>- VAPI_REM_OP_ERR -remote operation error</li> <li>- VAPI_RETRY_EXC_ERR - transport retry exceeded error</li> <li>- VAPI_RNR_RETRY_EXC_ERR - RNR retry exceeded error</li> <li>- VAPI_REM_INV_RD_REQ_ERR - remote invalid RD request</li> <li>- VAPI_INV_EECN_ERR - invalid EEC number error</li> <li>- VAPI_INV_EEC_STATE_ERR - invalid EEC state error</li> </ul>
free_res_count	Freed resource count (RD RQ only).

1. In cases of a RAW datagram or UD, the **len** includes 40 extra bytes, regardless of whether a GRH is included.

The **remote\_node\_address** is of type [VAPI\\_remote\\_node\\_addr\\_t](#) and is valid only for Datagram services. It is specified in the following table:

**Table 19 VAPI\_remote\_node\_addr\_t**

type	Remote node address type: <ul style="list-style-type: none"> <li>- VAPI_RNA_RD</li> <li>- VAPI_RNA_UD</li> <li>- VAPI_RNA_RAW_IPV6</li> <li>- VAPI_RNA_RAW_ETY</li> </ul>
slid	Source LID
sl	Service Level
qp_ety (Union)	<ul style="list-style-type: none"> <li>- qp - Queue pair (RD and UD)</li> <li>- ety - Ethertype (RAW Ethertype only)</li> </ul>
ee_dlid (Union)	<ul style="list-style-type: none"> <li>- loc_eecn - Local EEC number (RD)</li> <li>- dst_path_bits - DLID path bits (UD, RAW IPv6 and RAW Ethertype)</li> </ul>

## 9.4 Request Completion Notification

### SYNOPSIS:

```
VAPI_ret_t
VAPI_req_comp_notif
(
    IN      VAPI_hca_hndl_t      hca_hndl,
    IN      VAPI_cq_hndl_t      cq_hndl,
    IN      VAPI_cq_notif_type_t notif_type
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**cq\_hndl:** CQ handle.  
**notif\_type:** CQ Notification Type.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_CQ\_HNDL:** Invalid CQ handle.  
**VAPI\_EINVAL\_NOTIF\_TYPE:** Invalid notification type.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Requests one of the following types of notification (specified in **notif\_type**):

- **VAPI\_NEXT\_COMP** - notify on next completion.
- **VAPI\_SOLIC\_COMP** - notify on solicited completion.

## 9.5 Request Completion Notification for a Completion Queue

**Note:** This is an Extended VAPI function

### SYNOPSIS:

```
VAPI_ret_t
EVAPI_set_comp_eventh
(
    IN    VAPI_hca_hndl_t      hca_hndl,
    IN    VAPI_cq_hndl_t      cq_hndl,
    IN    VAPI_completion_event_handler_t completion_handler,
    IN    void *               private_data,
    OUT   EVAPI_compl_handler_hndl_t *completion_handler_hndl
)

```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**cq\_hndl:** CQ handle.  
**completion\_handler:** handler to call for completions on Completion Queue cq\_hndl.  
**private\_data:** pointer to data for completion handler.  
**completion\_handler\_hndl:** returned handle to use for clearing this completion handler.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_CQ\_HNDL:** Invalid CQ handle.

### DESCRIPTION:

Registers a specific completion handler to handle completions for a specific completion queue. The private data specified here is provided to the completion callback when a completion occurs on the specified CQ. If the private data is a pointer, it should point to static or "malloc'ed" data; the private data must be available until this completion handler instance is cleared (with EVAPI\_clear\_comp\_eventh).



## 9.6 Clear Completion Notification for a Completion Queue

**Note:** This is an Extended VAPI function

### SYNOPSIS:

```
VAPI_ret_t
EVAPI_clear_comp_eventh
(
    IN    VAPI_hca_hndl_t      hca_hndl,
    OUT   EVAPI_compl_handler_hndl_t  completion_handler_hndl
)
```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**completion\_handler\_hndl:** handle to use for clearing this completion handler.

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EINVAL\_CQ\_HNDL:** Invalid CQ handle.

### DESCRIPTION:

Clears a completion handler that was registered to handle completions for a specific completion queue. If a handler was not registered, returns OK anyway.



# 10 Event Handling Verbs

This chapter describes the Mellanox implementation of the following verbs handling events:

- Set Completion Event Handler ([p.73](#))
- Set Asynchronous Event Handler ([p.74](#))

## 10.1 Set Completion Event Handler

### SYNOPSIS:

```
VAPI_ret_t
VAPI_set_comp_event_handler
(
    IN      VAPI_hca_hndl_t      hca_hndl,
    IN      VAPI_completion_event_handler_t handler,
    IN      void                 *private_data
)

```

**ARGUMENTS:** **hca\_hndl:** HCA handle.  
**handler:** Completion Event Handler function address.  
**private\_data:** Pointer to handler context (handler specific).

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Registers a completion event handler. Only one CQ event handler can be registered per HCA.

The CQ event handler function prototype is as follows:

```
void
VAPI_completion_event_handler
(
    IN      VAPI_hca_hndl_t      hca_hndl,
    IN      VAPI_cq_hndl_t      cq_hndl,
    IN      void                 *private_data
)

```

## 10.2 Set Asynchronous Event Handler

### SYNOPSIS:

```
VAPI_ret_t
VAPI_set_async_event_handler
(
    IN      VAPI_hca_hdl_t      hca_hdl,
    IN      VAPI_async_event_handler_t handler,
    IN      void                *private_data
)

```

**ARGUMENTS:** **hca\_hdl:** HCA handler.  
**handler:** Asynchronous event handler function address.  
**private\_data:** Pointer to handler context (handler specific).

**RETURNS:** **VAPI\_OK**  
**VAPI\_EINVAL\_HCA\_HNDL:** Invalid HCA handle.  
**VAPI\_EPERM:** Insufficient permissions.

### DESCRIPTION:

Sets the specified Asynchronous Event Handler.

The handler is a pointer to a function as follows:

```
void
VAPI_async_event_handler
(
    IN      VAPI_hca_hdl_t      hca_handle,
    IN      VAPI_event_record_t *event_record_p,
    IN      void                *private_data
)

```

The `event_record_p` is a pointer to a structure of type `VAPI_event_record_t`, which contains the following fields:

**Table 20 VAPI\_event\_record\_t**

type	Asynchronous event type: <ul style="list-style-type: none"> <li>- VAPI_QP_PATH_MIGRATED</li> <li>- VAPI_EEC_PATH_MIGRATED</li> <li>- VAPI_QP_COMM_ESTABLISHED</li> <li>- VAPI_EEC_COMM_ESTABLISHED</li> <li>- VAPI_SEND_QUEUE_DRAINED</li> <li>- VAPI_CQ_ERROR</li> <li>- VAPI_LOCAL_WQ_CATASTROPHIC_ERROR</li> <li>- VAPI_PATH_MIG_REQ_ERROR</li> <li>- VAPI_LOCAL_CATASTROPHIC_ERROR</li> <li>- VAPI_PORT ERROR</li> </ul>
modifier	Union of the following four fields:
qp_hndl	The affiliated QP handle (if applicable)
eec_hndl	The affiliated EEC handle (if applicable)
cq_hndl	The affiliated CQ handle (if applicable)
port_num	The affiliated port number (if applicable)

