
InfiniBand Network Architecture

Last updated on: September 1, 2004

InfiniBand Network Architecture, First Edition Errata

Page	Severity	Date	Description
269	Moderate	9/1/03	In the "Number of Buffers" row, the description should read "The number of <u>Scatter</u> buffers in the list (note that 0 is valid)."
276	Minor	9/1/03	In the "Scatter Buffer List" row, the final sentence in the description should read "The <u>presence</u> of a GRH is indicated by a bit in the CQE."
299	important	9/1/04	Page 299, last paragraph should be changed to read: "When software tells the verb layer to create a virtual memory region, the verb layer (which is part of the OS kernel) creates a virtual-to-physical address mapping table for the use of the HCA. The table maps QP-generated virtual addresses (in the range supplied by the calling application) to the same physical pages that are accessed by the calling application."
300	important	9/1/04	Delete the last sub-bullet under the first major bullet item.

InfiniBand Network Architecture

InfiniBand Network Architecture, First Edition Errata

Page	Severity	Date	Description
301	important	9/1/04	<p>Refer to Figure 1 on page 9. Replace the numbered list with the following:</p> <ol style="list-style-type: none">1. The application calls the OS malloc routine with a request to allocate a buffer within the application's memory space.2. The malloc routine locates a series of physical pages in memory that are available and provide at least the requested amount of memory.3. The malloc routine creates a series of page table entries that map the application accesses within its virtual space to the physical pages in memory.4. Malloc returns control to the application.5. The application executes the <i>Register Memory Region</i> verb to create a region of memory that can be accessed by the QP(s) the application uses to send and receive messages. It supplies the start virtual address and length of the region, as well as the ID of the calling application.6. Using the application's ID and supplied virtual address range, the <i>Register Memory Region</i> verb obtains the list of physical pages that are mapped into the application's virtual buffer space.7. The verb creates a page table within the HCA that maps the same virtual buffer address space into the same physical memory pages.8. The verb returns control to the calling application, supplying it with a local access key (L_Key) and, if remote access rights were requested, a remote access key (R_Key).

InfiniBand Network Architecture

InfiniBand Network Architecture, First Edition Errata

Page	Severity	Date	Description
305	important	9/1/04	Delete the following text: As a result, the verb makes a call to the OS memory management routine specifying the following: <ul style="list-style-type: none">• The requested 64-bit start virtual address is 141200h.• The length of the region is 10000d bytes.• The pages allocated by the OS must be pinned in memory.• The requested access control attributes can be any desired.
305	important	9/1/04	In the middle of the page <ul style="list-style-type: none">• replace “The OS memory management routine takes the following actions:”• with “The <u>Register Memory Region verb</u> takes the following actions:”
305	important	9/1/04	In item 1, <ul style="list-style-type: none">• replace “Based on their availability, it allocates the following three 4KB pages to the memory region being created:”• with “Using the application’s ID and supplied virtual address range, the <i>Register Memory Region verb</i> obtains the list of physical pages that are mapped into the application’s virtual buffer space.”
305	important	9/1/04	In item 2, <ul style="list-style-type: none">• replace “The OS memory management routine creates the following virtual-to-physical address mappings (in a processor-specific manner, typically by creating Page Table entries):”• with “The verb creates a page table within the HCA that maps the same virtual buffer address space into the same physical memory pages.”
305	important	9/1/04	On the 2nd text line, delete “(the verb)”.

InfiniBand Network Architecture

InfiniBand Network Architecture, First Edition Errata

Page	Severity	Date	Description
306	important	9/1/04	Replace all of the text in item 4 with the following: “The verb returns the following information to the calling application: <ul style="list-style-type: none">• The region’s handle.• The local access key (L_Key).• If remote access rights were requested, the remote access key (R_Key). ”
306	important	9/1/04	Delete the sentence immediately above the figure.
307	important	9/1/04	Delete the entire first sentence (begins with “When a virtual memory region is created...”).
422	Important	9/1/03	In the first paragraph, the third sentence should read “The first WQE posted has an SSN of <u>000001h</u> , the second one an SSN of <u>000002h</u> , and so on. ”
467	Important	9/1/03	In the numbered list, change all references to “QP5” to “QP2”.
473	Minor	9/1/03	Item 2 at the bottom of the page reads should read “If the DGID is the loopback address (<u>address 1</u>), then it is invalid.”
474	Minor	9/1/03	Item 2 should read “If the SGID is the loopback GID (<u>address 1</u>), then it is invalid.

InfiniBand Network Architecture

InfiniBand Network Architecture, First Edition Errata

Page	Severity	Date	Description
505	Clarification	9/1/03	<p>Figure 19-6 and its description on pages 502 and 503 are correct as written. However, the text (and the spec) does not explain how the Resync packet issued <u>after</u> the request packet with a PSN of 3 manages to arrive at the destination EEC <u>before</u> the long-delayed request packet. Although the spec doesn't say so, the scenario assumes that the two EECs reside in different IBA subnets and that the packets exchanged between the two must cross a stretch of WAN (Wide Area Network) that tunnels IBA packets through the non-IBA, IP network. Since the routing taken by each packet in an IP network can be different, there are two possible scenarios:</p> <ul style="list-style-type: none"> • Case 1 is the one illustrated, wherein the Resync packet arrives before the long-delayed request packet. • In case 2, the Resync packet arrives after the request packet with a PSN of 3. <p>In either case, after transmitting the Resync packet, the requester EEC's Send Logic will not transmit any additional request packets until it receives the Ack corresponding to the Resync packet.</p>
609	Moderate	9/1/03	<p>The illustration on page 603 shows three possible exits from the LinkActive state, while the text on 608-609 only shows two. Add the following as item 3 on the top of page 609:</p> <p>"3. LinkActive to LinkArm. An SMP packet is received setting <i>PortInfo.PortState</i> = to Arm."</p>
612	Trivial	9/1/03	<p>Under "LNH Field" heading, the first sentence should read "The 2-bit Link Next Header (LNH) field identifies what header <u>immediately</u> follows"</p>
656	Trivial	9/1/03	<p>In the first bullet item, the empty quotes should contain "Link Packet (Flow Control Packet) Check" on page 661."</p>
658	Trivial	9/1/03	<p>In the second paragraph, "ealier" should be "earlier".</p>

InfiniBand Network Architecture

InfiniBand Network Architecture, First Edition Errata

Page	Severity	Date	Description
679	Minor	9/1/03	Item 8 should read "VL arbitration is performed in the exit port's Link Layer to determine which data VL transmit <u>buffer</u> gets to transmit next".
698	Important	9/1/03	The heading of the left-most column should read "RcvControl State".
746	Trivial	9/1/03	In Step 3, the first sentence should read "Both the local <u>and</u> the remote ports..."
750	Important	9/1/03	<ul style="list-style-type: none"> Item 7 should read "The receiver must be capable of de-skewing a minimum of six symbol times of total <u>lane-to-lane</u> skew." Item 8 should read "After successful completion of <u>lane-to-lane</u> de-skew,...".
751	Trivial	9/1/03	In the first paragraph, "(10 X X 400ps = 4ns)" should read "(10 X 400ps = 4ns)".
765	Important	9/1/03	The last sentence should read "The receiving port's Link Layer accepts the response SMP into its VL15 <u>receive</u> buffer and forwards the response packet back to the SM."
767	Minor	9/1/03	Step 5 should read "The receiving port's SMI SQ Logic swaps the LID addresses (SLID and DLID) that were in the request SMP, sets the BTH:DestQP field to QP0, and passes the packet <u>to</u> the receiving port's Link Layer where it is accepted into the VL15 transmit buffer and is transmitted to the SM."
767	Trivial	9/1/03	In the first bullet item at the bottom of the page, the first sentence should read "Earlier, the SM had <u>sent</u> a request SMP to a remote device port."
770	Trivial	9/1/03	The first sentence under the heading "Port States SMPs Can Be Sent and Received In" should read "The SMI can transmit and receive SMPs when the Link Layer on the port is <u>in</u> any of the following states:"

InfiniBand Network Architecture

InfiniBand Network Architecture, First Edition Errata

Page	Severity	Date	Description
788	Important	9/1/03	In table 28-5, the final Reserved field is 128 bytes in length (not 64 bytes).
789	Important	9/1/03	In table 28-6, the last row should be replaced with the following two rows: <ul style="list-style-type: none"> • 32-47 64 bytes—InitialPathArray • 48-63 64 bytes—ReturnPathArray
824	Important	9/1/03	Under the “PortInfo Attribute” heading, the second sentence should read “The Master SM accesses the <i>PortInfo</i> attribute by sending a <i>SubnGet(PortInfo)</i> or <i>SubnSet(PortInfo)</i> SMP MAD to the port.”
849	Trivial	9/1/03	In the first bullet item, the third sentence should read “If the port supports <i>Notices</i> as indicated by the <i>Port-Info.CapabilityMask.IsNoticeSupported</i> bit, the SMA must log a notice for P_Key and Q_Key mismatches.”
857	Important	9/1/03	Item 3 should read “Each of the Master <u>SMs</u> issues a series of <i>SubnGet(PortInfo)</i> SMPs to check the <i>Port-Info.PortState</i> of each of the switch ports.”
873	Important	9/1/03	In table 31-1, the final Reserved field is 128 bytes in length (not 64 bytes).
877	Important	9/1/03	In table 31-2, the last row should be replaced with the following two rows: <ul style="list-style-type: none"> • 32-47 64 bytes—InitialPathArray • 48-63 64 bytes—ReturnPathArray In addition, the lone little table cell at the left end of byte 0 in the last three rows is incorrect (just remove the left border of the cell).
880	Trivial	9/1/03	Under item 1’s second subbullet group, in the second subbullet, the second sentence should read “If this port’s LID address has not yet been assigned, the SLID is set <u>to</u> the PLID address (FFFFh).”

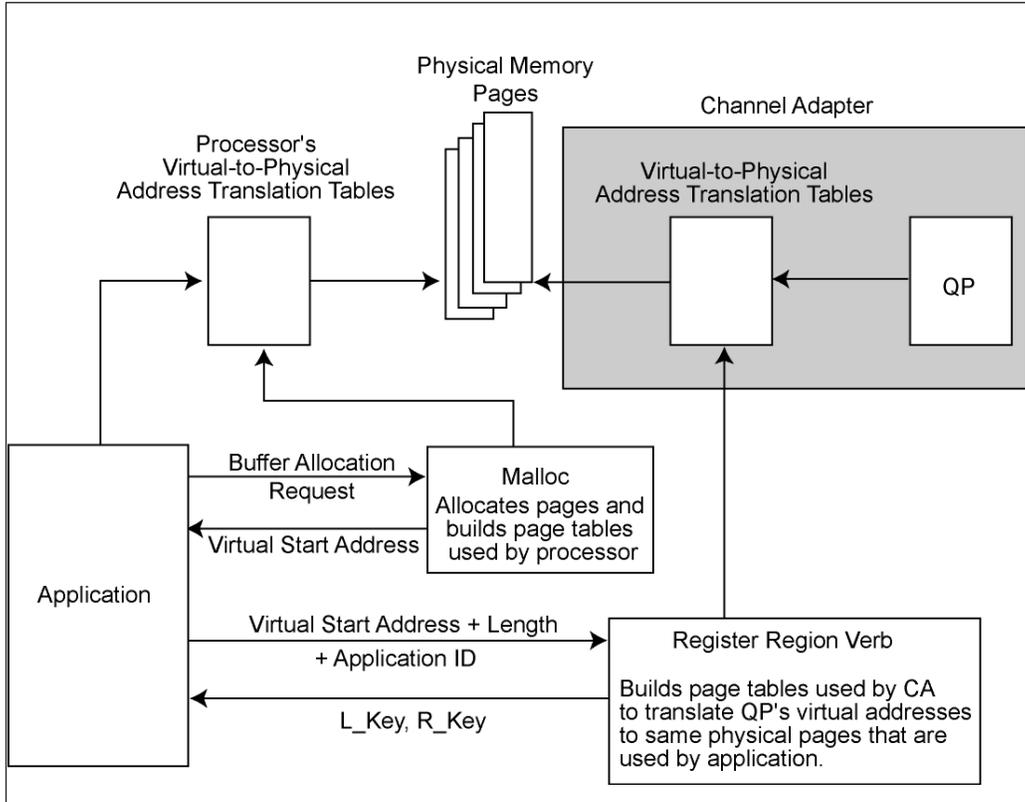
InfiniBand Network Architecture

InfiniBand Network Architecture, First Edition Errata

Page	Severity	Date	Description
923	Trivial	9/1/03	Under the "Location of the SA" heading, first paragraph, the second sentence should read "The location (i.e., port LID address) of the Master SM is stored <u>in</u> each port's <i>PortInfo.MasterSMLID</i> attribute element."
944	Moderate	9/1/03	In the third bullet item from the bottom of the page, the text in parentheses should read "(an Attribute Modifier of 0 and an End RID of FFFFFFFFh specifies all records in the table)".
957	Moderate	9/1/03	The last arrow in the figure should be annotated "..., Segment = N + 3".
1080	Moderate	9/1/03	In the right column of the first row, "RD. 1-bit value." should read "RC. 1-bit value."
1088	Important	9/1/03	The last sentence should read "...back to the requester with information about the newly created QP <u>or</u> EEC and QP."
1112	Trivial	9/1/03	Under the heading "Active Client to Passive Server", the first sentence should read "This <u>is</u> the most straightforward scenario ..."

InfiniBand Network Architecture

Figure 1: The Register Region Verb Call



InfiniBand Network Architecture
