Design Guide for Open Rack Management Backplane Connection (v. 0.52):

Q & A from Boston Open Rack Workshop Presentation
October 9th, 2015

Compiled by: Ethan Long, Intel

Mechanical Alignment

1) Can this MBP design support ½ OU servers?

No, the chosen connector is too long to support this density. This design assumes minimum 1OU increment servers. This detail will be added to the design guide.

**2) What are the tolerance requirements for the IT equipment making the MBP connection?**

The overall tolerances required to maintain a reliable connection vary with rack and IT equipment design, and it was not the intent of the design guide to specify a requirement here. However, Intel can help define a requirement if there is continued interest in including a management backplane in the Open Compute standard.

The connectors themselves can resolve up to 2 mm of X-Y misalignment during docking. However, the Intel/Quanta implementation is designed to tolerate up to 4 mm of X-Y misalignment through the use of the additional coarse alignment pins. The design guide will be updated to include the allowable misalignment in the Intel/Quanta implementation.

Mechanical Alignment (cont…)

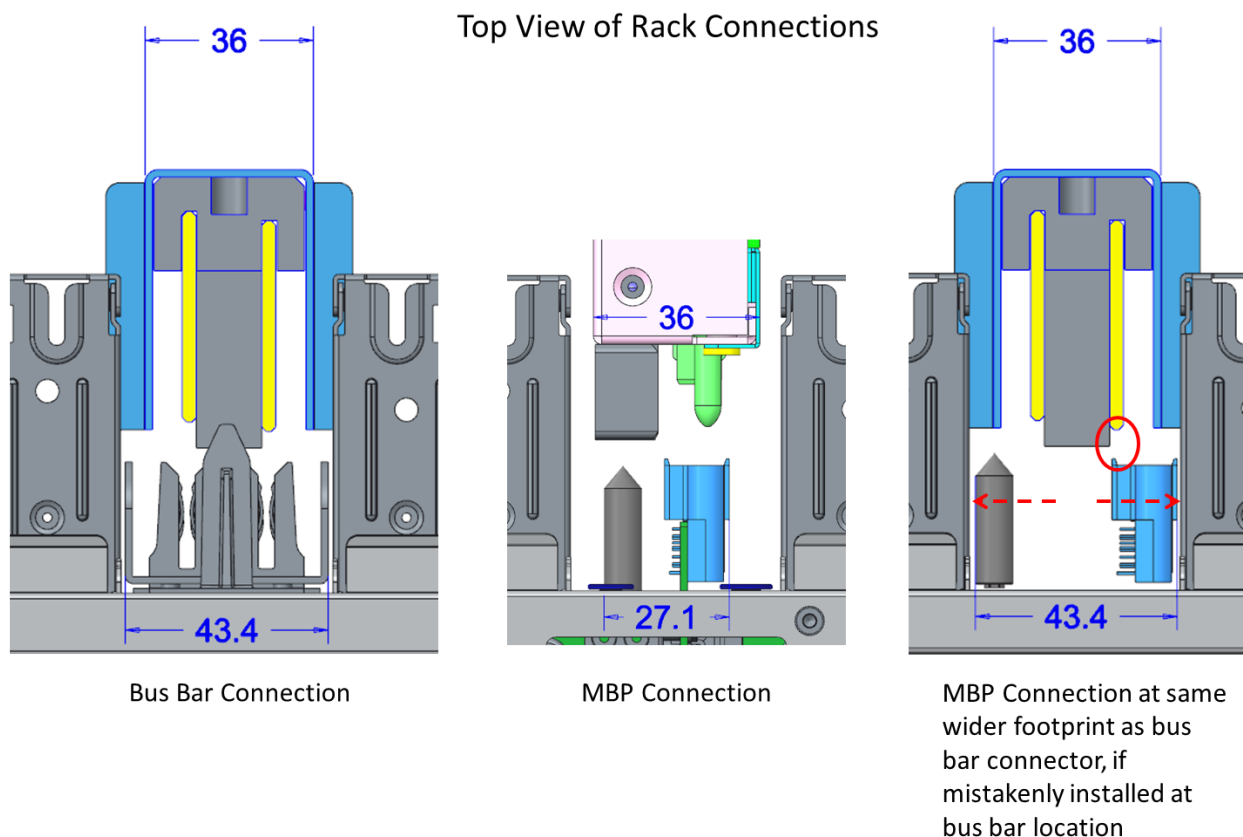**3) Are the alignment pins going to touch a hot bus bar if mistakenly installed into a 3 bus bar rack?**

In this scenario, the plastic housing of the management connector would contact the negative bus bar first, and prevent further travel (see top view image below). That would leave 2.2 mm of nominal clearance between the alignment pin and the positive bus bar. If one takes manufacturing variation into account, it is possible that the pin could contact the hot bus bar on some systems.
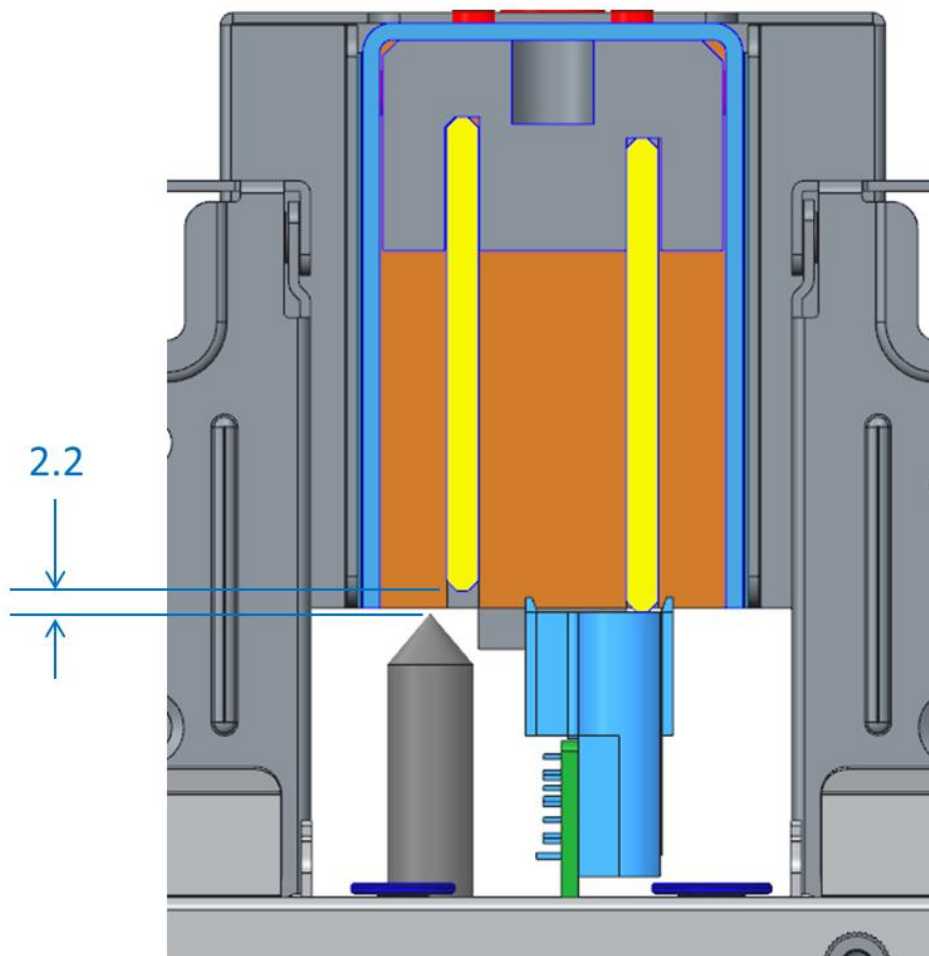
Mechanical Alignment (cont…)

4) The bus bar connectors extend to the sides past the bus bar space. Can that space also be used to possibly avoid bus bar contact (see previous question)?

No. The bus bar connector assemblies used on the demo systems, at 43.4mm, are wider than the MBP connector assembly, and the rack's 36mm-wide bus bar zone. Given this, the server-side connector and alignment pin mating locations may theoretically be spread wider to avoid bus bar contact altogether, but the management backplane in the rack will then exceed the 36 mm zone width specified in the Open Rack standard. Also, the Intel/Quanta implementation has optional fan mounting locations on either side of the management connection that prevent a wider backplane footprint. If the footprint could be set to the same width as the bus bar connector, the alignment pin would be at less risk of contacting the positive bus bar, but the MBP connector would still impact the negative bus bar. See the images below for a comparison of the various connections.

Top View of Rack Connections

Bus Bar Connection

MBP Connection

MBP Connection at same wider footprint as bus bar connector, if mistakenly installed at bus bar location

5) Does the MBP connector allow for the bus bar negative to contact first when inserting a tray?

Yes, the center bus bar connector contacts the negative first. The MBP connector does not establish a connection until after the bus bar power connection has been established. This detail will be added to the design guide.

Connector and Connector Pin Out

6) Is there more than one supplier for the MBP connector?

TE Connectivity also offers the following connector part numbers, which match the FCI spec:

Header: 8-1926730-1
Receptacle: 1-1892782-2

TE's part numbers will be added to the design guide for reference.

7) Are the MBP connector pins sequenced?

Yes, there are three stages of pin engagement. The power pins engage first, followed by most signal pins. There are two signal pins (A2/E2) that engage last. This information will be added to the design guide.

8) Is there a dedicated pin to detect full connector insertion?

There are two short pins that engage last and could be used to define full connector insertion. In this implementation, Pin A2 is used to detect if the backplane is present, which could be used as a measure of full insertion since it enables the flow of other management traffic. The other short pin, E2, is not used. Further analysis would be required to determine if E2 must be included in full insertion detection.

9) Are more pins required to specify tray ID's in backplanes larger than 8OU?

The Quanta/Intel system does limit the backplane to 8 connector locations. Intel considered backplane size when designing the demo rack. First, it was determined that 8OU is an optimal size for a cooling zone when using shared cooling. Second, the assumed High Volume Manufacturing limit for backplane board length is approximately 20 inches. An 8OU backplane requires at least 15 inches of length.

A spec could consider an option for adding additional tray locations to backplanes by making pin "TRAY_ID[0]" a serial bus that communicates location from the MBP controller. This approach would also free up additional spare pins for future use.

10) Can the spare pins be grouped together?

Perhaps, but the intended purpose of such a grouping is not evident to Quanta/Intel.

Power

11) Why is backplane power routed through the MBP connector from the IT equipment? Is this a requirement?

It is not required for the IT equipment to route power to the management backplane. Intel and Quanta chose to offer the option for use of the MBP connection because it eliminated the need for additional cables running from the power source to the backplane. The Intel/Quanta implementation also allows for a direct supply of power to the MBP. There are enough power pins assigned in the MBP connector to provide enough power for the MBP circuitry and 6 fans from a single MBP connection.

12) This specification discusses rack-level power capping only. How would power get managed at the server level?

Intel's assumption is that power will be managed at the server level by RSA API's operating through the chassis PSME. The THROTTLE_N signal was added for low-latency critical power situations.

Management backplane and management traffic

13) How does the management backplane know its physical location within a rack?

In order to minimize the number of connector pins, a dip switch was used to set the backplane location within the Quanta/Intel implementation.

Intel has a couple of ideas on how to assign a physical location to a backplane:
   a) Adopters can use a DIP switch
   b) Adopters may incorporate an interface to a physical rack feature that provides varying locations of grounding to the MBP depending on MBP location in the rack. As the MBP is installed in the rack, the rack feature then grounds different points on the MBP such that you get the equivalent of the DIP switch.

14) What is the advantage of using a 1 Gb rear management backplane interface over 10 Gb connections out the front?

In general, a rear-mounted management network is advantageous because of its proximity to the power source and fans, and because IT equipment may connect via a simple blind-mate interface. With regard to speed, 1 Gb is the current low-cost, high volume Ethernet standard, and more than enough bandwidth for the intended purpose of the interface (management). Intel did not predict any use case that would require speeds in excess of 1Gb.

15) Does this design require that Ethernet and UART serial lines are run from the motherboard to the MBP connection, rather than to the front of the server?

This is not an either/or question. The management backplane would require lines run to the motherboard. However, there may be opportunity to architect the server to accommodate a selectable front serial/Ethernet connections or a rear MBP connection.

For the serial lines, Intel believes that it is possible and low-cost to provide a console at both the front and MBP connections because the proposed spec uses TTL signal levels. This is accomplished by sending host transmit via a buffer to both locations, and using open collector buffers for host serial receive. Alternatively, you can route the signals to one or the other with a build time option.

For management Ethernet, Intel recommends placement of the Ethernet KX source at a location that allows a direct connection to either the front or the MBP.

16) Why did Intel/Quanta not employ wireless for management traffic?

The Intel approach was to implement the existing low-cost KX standard components.

17) How does one multiplex Ethernet and UART serial signals from all sleds through one MBP connector? What does that do to traffic flow?

This is done through the use of low-cost Ethernet switching and/or microcontrollers within the chassis. We believe that for management traffic this provides an opportunity for low-latency event notification. Further, we think this is part of the tray management design specification.

18) Do we need to explore other protocols for full rack implementation? For instance, can KX traffic be routed all the way from bottom to top without signal loss or speed hit?

The KX spec is well-defined by IEEE802.3 and the signal loss of a particular implementation (ie. rack) is the responsibility of the rack architect. In Intel's case, we utilized a low-cost Ethernet switch on the MBP.