Server Tech Talk Chat Transcript (Day 2)
April 27, 2022

00:30:25 Michael Schill: OCP Server Project Wiki: https://www.opencompute.org/wiki/Server
00:30:41 Kate Hendle: Welcome to the Server Tech Talk (Day 2)! Please input any questions you have for the speakers here in the chat and keep yourself muted if you are not talking. Thanks!
00:45:11 Ravi B: Are there multiple outlines for HPM specified to capitalize on the open rack's wider chassis vs. EIA 19" rack chassis that are slightly narrower ?
00:58:23 Shawn Dube: The current HPM widths are for a superset of 19" and 21" rack applications. For 21" racks, we envision (1) utilization of the extra width for platform infrastructure, or (2) Adaptations of the Base HPM designs to enable additional capabilities for targeted designs.
01:00:25 Andrew Junkins: Seems HBM definition would also need to include a common backplane architecture in terms of power delivery and management approach. Anything you can highlight on this topic.
01:01:45 Tim Lambert: Andrew, M-XIO and M-PIC include the HPM interfaces for very flexible peripheral card subsystems....including backplanes with hotplug support.
01:16:23 sheng: What is the HPM type 3 Power Ranges for CPU + DDR memory?
01:19:59 Phil: Liquid cooling ports location?
01:21:12 Michael Gregoire:To be more specific each Power Connector supports 864W
01:21:37 Michael Gregoire:We have 2 placements in each Power Zone - we suspect a single power zone would be used for Ingress in most designs so 864x2 maximum
01:22:09 Michael Gregoire:Although we do not prohibit use of both power zones for Ingress (although the system layout would be potentially challenging) which could in theory get you 864x4
01:22:36 Shawn Dube: Liquid cooling ports location is an artifact of specific HPM designs, so is not part of the HPM form-factor spec. Industry standardization could show up in subsequent Design Specs
01:23:16 Eduardo Estrada: 864W nominal, before whatever derating spec is used right?
01:24:34 Michael Gregoire:Yea, 864 Nominal
01:24:57 Eduardo Estrada: thanks
01:26:54 Corey Hartman: Re, question on Liquid Cooling Port Locations: Liquid cooling solutions were considered in determining certain aspects of HPM form factor specs (overall sizes, height restrictions). But the variety of potential solutions is large (for example, Near Edge exit, vs Far Edge exit.) So, as Shawn mentions, liquid interconnects are out of scope for these Base Specs, and could be detailed in generational design specs, and /or future Standardization efforts.
01:30:08 Ravi B: could you comment on how 864W nominal is arrived at
01:30:43 Phil: Will the liquid locations be fixed or not on future designs?
01:30:49 Chris Scott: Anyone demonstrating a system with the DC-SCM 2.0 managing multiple hosts?

01:32:00 Shawn Dube: Phil, that depends of what's appropriate for the future design. It's not immediate scoped for the base specs.
01:32:24 Phil: Thanks
01:32:38 Michael Gregoire: Chris - The DNO team has discussed a number of ways this can be done - so far we don't feel like it has real ties into the base form factor specification... but the team has reviewed a few system concepts to try and ensure we are not prohibiting anything. We have one example in our appendix.
01:32:45 Eduardo Estrada: for a 6x2 Power Pins at 12A per pin. 6*12A*12V= 864W
01:33:08 Shawn Dube: Chris, good question for the DC-SCM 2 topic later on.
01:33:08 Eduardo Estrada: 6pins for PWR and 6 Pins for Return GND
01:33:27 KWANG04: If the power cable is pinched and shorted, what is protecting the 12V source?
01:37:40 Sanjiv Sinha: KWANG04, fusing on the HPM is allowed to be implemented to protect connectors, copper planes, power sources. The fusing mechanism is not explicitly defined.
01:39:15 Tim Lambert: A design spec could call out fuses upstream of the PICPWR source connector. When choosing fuse limits, be aware of new PCIe CEM High Power excursion changes (e.g., up to 3X nominal for short periods on a logarithmic scale)
01:41:58 Phil: SMBus protocol i2c or i3c ?
01:43:07 Tim Lambert: M-XIO, PICPWR and CP call out SMBUS and optionally I3C Basic 1.1.1 (i.e. the same discovery / negotiation to lower voltage to follow PCIe SIG internal cable and CEM specs).
01:45:53 Pekon Gupta (SMART Modular): Is there a reason why x16 -> 4 (x4) is not allowed. with increase in speeds (x4) may be more popular than (x8) ?
01:46:23 bo.baudrexl: could you please show the connectivity slide once more?
02:17:18 Kali Burdette: Yesterday's Server tech talk is now available on demand: https://www.opencompute.org/events/past-events/ocp-tech-talk-series-server-day-1
02:31:15 Tim Lambert: Before anyone asks, we are looking to merge the best of DSSI and M-PESTI protocols. Both deterministic / real-time with similar latency needs.
02:33:18 Phil: Does the PS provide real time telemetry on power demand ?
02:34:22 Phil: Thank you
02:57:39 Rob Nance: What is the best way to follow the development of this standard? Will there be a Server sub-project created specifically for DC-MHS?
02:58:31 sheng: any DC-SCM chiplet (BMC,ROT,TPM) that one can include in their SoC socket?
03:01:18 Tim Lambert: There are BMCs with platform ROT today. TPM remains a discrete subsystem.
03:15:11 Chandra Krishnaswamy: We already defined some standard interconnects (connectors and cables) within a board. Is OCP going to define/select any standard interconnects between the boards? If so when?
03:25:42 sheng: Host Processor Module = (CPU|GPU|FPGA) + (DDR|HBM) + IO::HighSpeed::Connectors::PCIeSwitch|NVlinkSwitch
03:27:47 Shawn Dube: Sheng, that's a valid set of combinations. There are further combinations choices that are valid.

03:29:56 Mark Shaw: I have to drop. Thanks everybody for a great Tech Talk!
03:31:18 bo.baudrexl: great concepts and dialog, thank you!
03:31:22 Fang Yang: Great talk