2025 DESIGN AND VERIFICATION™ DVCON CONFERENCE AND EXHIBITION INDIA

Advancing Early Enablement, Validation, and Signoff of Manageability IP for Seamless dGPU-Server Paring

Shyam Kumar Choudhary, Paras Agarwal and Madhusoodanan Pudhucherry
Intel Technology India Pvt. Ltd. Bengaluru, Karnataka
shyam.kumar.choudhary@intel.com, paras.agarwal@intel.com,
madhusoodanan.pudhucherry@intel.com

***Abstract*- This paper presents a comprehensive approach for enablement, validation and enhancing manageability in server equipped with discrete GPUs (dGPUs) through Board management controller (BMC) Add-in management card (AMC) interaction, marking a pioneering effort to sign off on serial interfaces such as Management Component Transport Protocol (MCTP) over Improved-Inter-Integrated Circuit (I3C) for enabling dGPU server pairing. The integration of dGPUs into server environments is crucial for meeting the demands of high-performance computing, yet it introduces unique challenges in terms of manageability and communication between components. This paper addresses the challenges in validation of implementing MCTP over I3C bus, a critical step in achieving effective manageability.**

Keywords- Board management controller (BMC), Management Component Transport Protocol (MCTP), Discrete GPUs (dGPUs), Add-in management card (AMC), Improved-Inter-Integrated Circuit (I3C), Peripheral Component Interconnect Express (PCIe), Security Protocol and Data Model (SPDM), System Management Bus (SMBus), Inter-Integrated Circuit ((I2C), In-Band Interrupt (IBI).

I. Introduction

Manageability is a critical component of modern server systems. It enables remote engineering users and system administrators to maintain comprehensive control and visibility over servers and their internal metrics. Manageability encompasses a collection of tools and methods designed to simplify routine tasks, such as monitoring server status, configuring systems, performing software updates, troubleshooting issues, and ensuring security. The primary objective of manageability is to automate and streamline these operations, thereby allowing organizations to maintain server reliability and security while reducing manual effort and operational costs. In large-scale server environments and enterprise data centers, robust manageability features such as remote server control, automated alerts, and centralized management platforms are essential for efficient operations, minimizing downtime, and enabling rapid response to potential issues.

Manageability validation is important in post and pre-silicon environments because it ensures that all the manageability features in a server work as expected before the server is used in real environments. By testing and validating these features, we can catch problems early, such as issues with remote control, monitoring, or updates. This helps prevent failures or downtime when the server is running important tasks. Proper validation also makes sure that IT staff can rely on the tools to manage servers easily and securely,

which is especially important in large data centers where many servers need to be managed at once. In short, manageability validation helps deliver reliable, secure, and easy-to-use server systems.

Key challenges in manageability validation:

- Early-stage validation is impeded by unavailable production-grade drivers, AMC–BMC integration, and firmware readiness, creating circular HW/SW dependencies; a minimal-dependency validation framework is required to enable pre-driver/FW debug and accelerate SoC, software, and board bring-up.
- Increased complexity due to disaggregated, multi-die SoCs and multiple manageability domains/subsystems.
- Scalability concerns as telemetry and manageability data volume grows with more cores and components.
- Difficulty in ensuring standardization and interoperability across new interfaces and protocols (e.g., I3C[1], MCTP [2]).
- Need for robust security, authentication, and access control in manageability flows.
- Importance of negative testing to validate system robustness against invalid or malicious requests.
- Ensuring seamless integration and error handling across hardware, firmware, and software layers.
- Automatic detection and inventory of Peripheral Component Interconnect Express (PCIe) [3] cards, including dGPUs, for asset management.
- Real-time telemetry collection for proactive maintenance and reduced hardware failure risk.
- Power management and control to maintain system stability and prevent overloads.
- Support for remote firmware updates, recovery, and secure operation.
- Enhanced security, crash/debug data collection, hot-plug support, and unified management of diverse PCIe devices.

The I3C MCTP interface is an emerging standard for connecting dGPUs and other PCIe devices to server manageability systems. I3C is a high-speed, low-power serial bus that improves upon older protocols like Inter-Integrated Circuit ((I2C)[4] and System Management Bus (SMBus)[4]. MCTP is a protocol that allows management messages to be sent over various physical interfaces, including I3C and PCIe.

II. Proposed Methodology

To enable early development and thorough validation of dGPU telemetry integration with server platforms, we established a comprehensive and robust testing framework utilizing specialized exerciser tools. This framework was designed to execute MCTP transactions and facilitate access to telemetry data from dGPU SoCs. Two primary types of exercisers were employed in this process: first, an I3Cexerciser [5] capable of operating in both master and slave modes, which was used to encode and decode MCTP protocol packets; and second, an evaluation board featuring the same chipset intended for customer deployment as an AMC. This evaluation board [6] was connected to a validation card running firmware that emulated MCTP transactions, closely mirroring real-world operational scenarios.

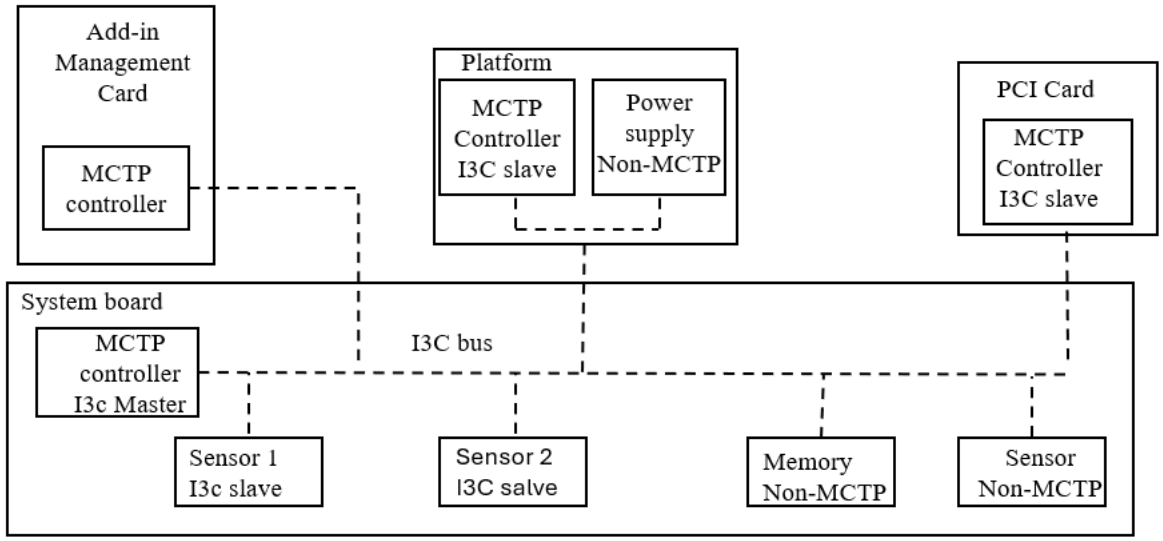

Figure 1. Physical Topology of the I3C Bus with MCTP

This setup enabled comprehensive testing and validation of the MCTP implementation, ensuring not only protocol compliance but also robust performance and compatibility across a variety of server configurations. Implementing MCTP over the I3C bus on server platforms significantly enhances the efficiency and speed of telemetry data extraction from graphics cards. It allows for the collection of critical metrics such as energy consumption, temperature, performance statistics, and PCIe link speed, all without the need for traditional driver-based solutions.

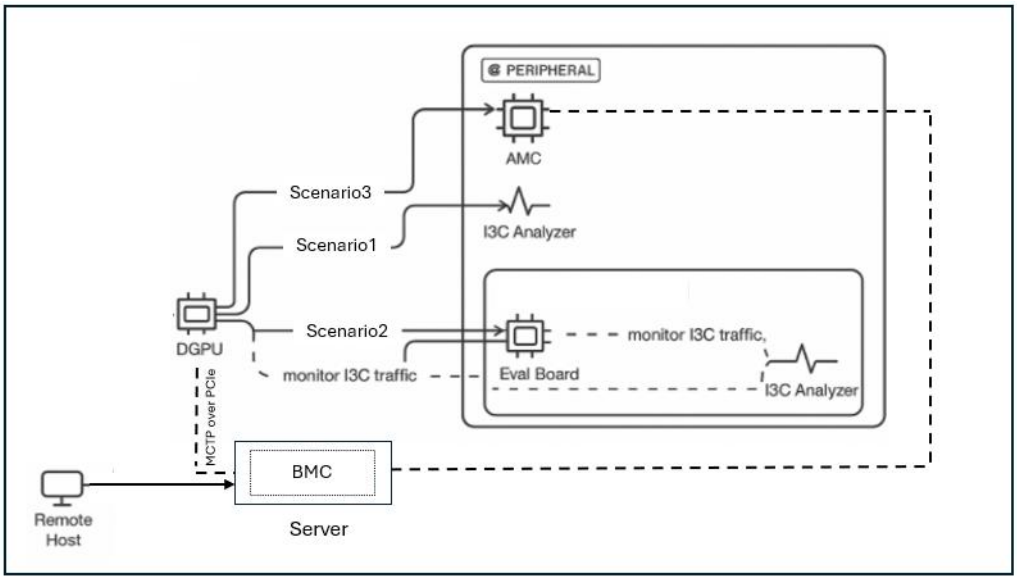

Figure 2. Telemetry Access from dGPU Across Various Scenarios

Platform Monitoring Technology (PMT) [7] and Platform Environment Control Interface (PECI) [7] payloads are encapsulated within Vendor-Defined Messages (VDMs), streamlining the data acquisition process. Our approach of using an I3C exerciser and evaluation boards to encapsulate VDMs over MCTP reduces dependencies on host drivers and firmware, thereby improving operational efficiency and simplifying system integration. In this architecture, the exerciser acts as the master device while the graphics card functions as the slave, with MCTP vendor-defined formats used to capture and structure telemetry data. This provides a reliable and scalable framework for validating dGPU resources within server environments.

Furthermore, the paper delves into the methodology of writing MCTP messages to the graphics card and retrieving telemetry data using In-Band Interrupt (IBI). This technique represents an innovative approach to graphics card telemetry, enabling asynchronous and efficient data transfer. The results from this validation framework demonstrate notable improvements in both validation cycle and quality.

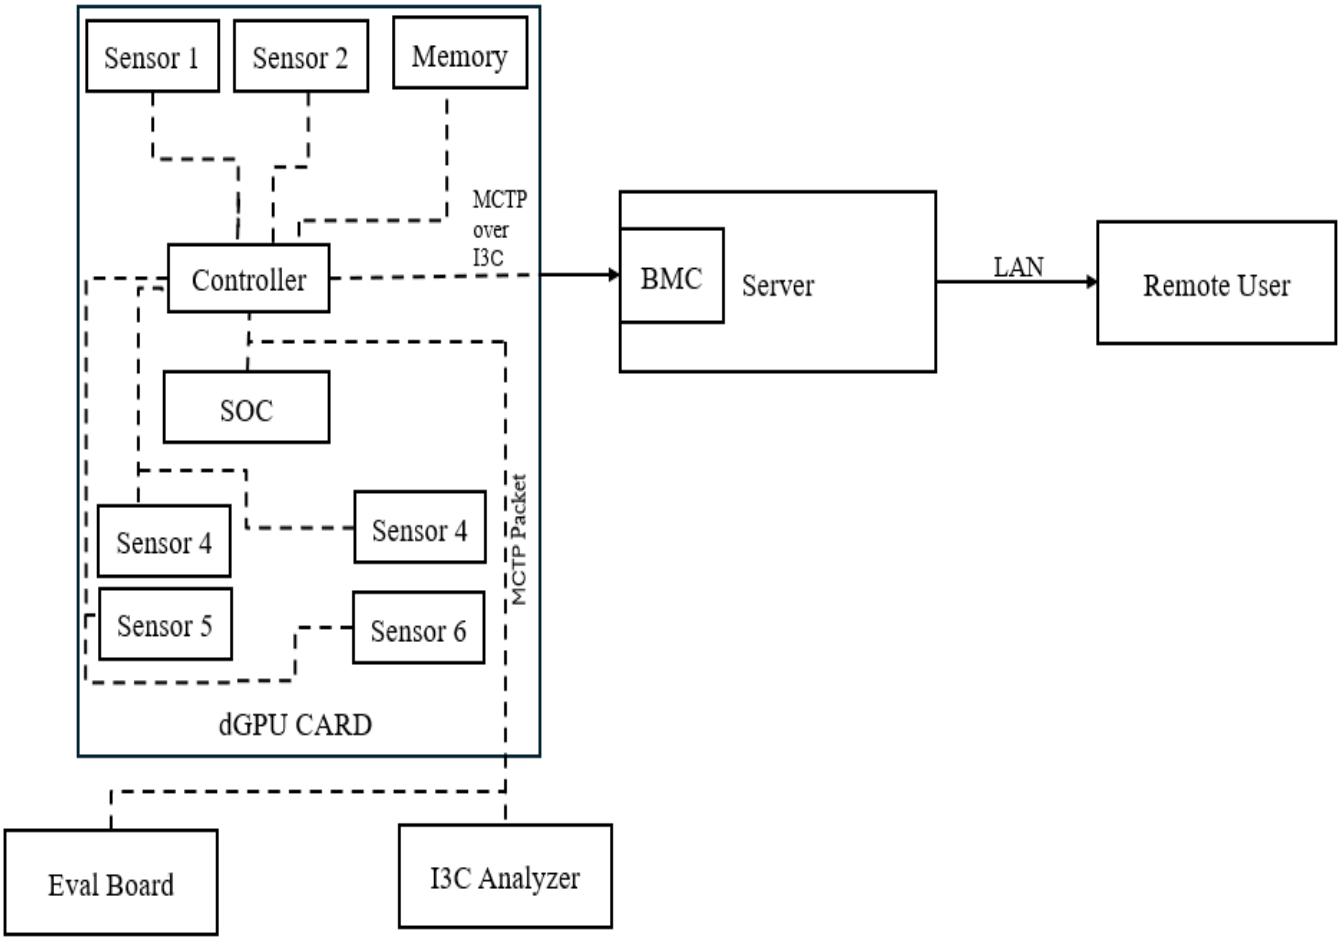

Figure 3. Data Access through dGPU Using Evaluation Board and Bus Analysis with Analyzer

Key solutions points :

- Developed a robust validation framework using exercisers to simulate and test MCTP transactions for dGPU telemetry integration in servers.
- Utilized two main exercisers:
  - An I3C exerciser operating in both master and slave modes for encoding and decoding MCTP packets.
  - An evaluation board with the same chipset as the intended AMC, connected to a validation card running firmware that mimics MCTP transactions.
- Enabled comprehensive testing of MCTP implementation, ensuring protocol compliance, performance, and compatibility across diverse server configurations.
- Implemented MCTP over the I3C bus using exerciser, allowing fast and efficient Validation of telemetry data (energy, temperature, performance, PCIe link speed) from graphics cards without traditional drivers.
- Using exerciser, demonstrated the process of writing MCTP messages and retrieving telemetry using IBI, enabling asynchronous and efficient data transfer.
- Validated that integrating MCTP over I3C is a viable and scalable solution for enhancing server manageability and gaining deeper insights into dGPU performance.

III. APPLICATIONS

Discrete GPUs Manageability Validation with Server Telemetry is essential for optimizing high-performance computing tasks such as AI , machine learning, and data analytics [8]. It provides valuable insights into various aspects of system performance and resource management.

Application	Description
Fault Detection & Recovery	Enables rapid identification and mitigation of hardware faults, thermal excursions, and power anomalies.
Security & Compliance	Supports validation of secure telemetry data and enforcement of access controls for ensuring regulatory compliance.
Performance Optimization	Provides insights on real-time telemetry for dynamic resource allocation, workload balancing, and energy efficiency tuning.
Automated Orchestration	Supplies validated data for orchestration frameworks to enable agile workload placement, scaling, and migration.
Debug & Root Cause Analysis	Delivers reliable telemetry and error logs for accelerated root cause analysis and reduced time to market.
Error isolation	Assists in isolating and containing faults to prevent systemic propagation and maintain platform reliability.

IV. RESULTS

The validation of dGPU telemetry over MCTP/I3C was executed in a phased approach, beginning with the exposure of I3C master headers to enable early MCTP transaction testing. In the initial phase, an I3C exerciser was used to inject synthetic MCTP payloads, enabling rapid prototyping and early bug identification in both the dGPU firmware and the manageability stack. This approach facilitated the detection and resolution of protocol compliance issues, timing mismatches, and error handling gaps.

Following iterative feedback and firmware stabilization, the validation environment was upgraded to utilize an advanced evaluation board equipped with production-grade firmware identical to the intended AMC. This transition enabled a high-fidelity emulation of customer deployment scenarios, ensuring that the telemetry path, protocol stack, and error recovery mechanisms were exercised under realistic conditions.

Key technical outcomes and metrics include:

- **Protocol Robustness:** >99.9% MCTP success over 10K+ cycles; error recovery, zero data loss.
- **Coverage:** Validated power, thermal, PCIe link, and performance telemetry; >95% firmware code coverage.
- **Defect Resolution:** Early discovery of error identification, packet misalignment, and state-machine bugs; all fixed pre-production.
- **Schedule Impact:** Enabled early telemetry for board bring-up, accelerated debug, informed Reliability, Availability, and Serviceability (RAS) strategies.
- **Time-to-Market:** Staged test flow (synthetic → eval boards → full stack) ensured rapid maturity and production-ready deployment.

V. Conclusion

The deployment of MCTP over I3C bus marks a transformative leap in server platform manageability, especially for architectures with dGPUs. The validation methodology presented in here enables early validation of MCTP over I3C which enables driver-independent telemetry validation, crucial for monitoring metrics like power consumption, thermal status, PCIe link integrity, and performance counters.

Incorporating VDM, PMT, and PECI payloads within the MCTP framework significantly enhances data collection and refines system health monitoring. The staged validation approach, from synthetic exerciser-based testing to full-stack evaluation board deployment, ensures comprehensive protocol compliance, robust error handling, and high interoperability across diverse server configurations and dGPU silicon revisions.

This work lays the groundwork for scalable, out-of-band manageability solutions resilient to platform topology changes and late-binding device configurations. Future research will aim to optimize MCTP transaction throughput validation, Validation of advanced telemetry classes, and security protocols like Security Protocol and Data Model (SPDM) [9]. Additionally, the framework will be evaluated for broader applicability in heterogeneous server environments, including AI accelerators and other PCIe-based devices, to further enhance platform reliability, serviceability, and operational efficiency in next-generation data centers.

VI. REFERENCES

[1] DMTF, "Management Component Transport Protocol (MCTP) I3C Transport Binding Specification," Revision 1.0, [Online].
Available: https://www.dmtf.org/documents/pmci/management-component-transport-protocol-mctp-i3c-transport-binding-specification-100

[2] DMTF, "Management Component Transport Protocol (MCTP) Base Specification," Revision 1.3.1, [Online].
Available: https://www.dmtf.org/documents/pmci/management-component-transport-protocol-mctp-base-specification-131

[3] PCI Express," Wikipedia, The Free Encyclopedia, [Online].
Available: https://en.wikipedia.org/wiki/PCI_Express

[4] DMTF, "Platform Level Data Model (PLDM) for Firmware Update Specification," DSP0237, version 1.2.0, [Online].
Available: https://www.dmtf.org/sites/default/files/standards/documents/DSP0237_1.2.0.pdf

[5] Prodigy Technovations, "MCTP Protocol Exerciser/Analyzer over I3C," [Online].
Available: https://www.prodigytechno.com/i3c/mctp-protocol-exerciser-analyzer-over-i3c

[6] ASPEED Technology Inc., "AST2750 Server Management Processor," [Online].
Available: https://www.aspeedtech.com/server_ast2750/

[7] Intel Corporation, "Intel® Platform Monitoring Technology Technical Specification," Rev. 3.2, [Online].
Available: https://cdrdv2-public.intel.com/710389/710389_Intel%C2%AE_Platform_Monitoring_Technology_TS_Rev_3_2.pdf

[8] Datadog, "Monitor NVIDIA GPUs with Datadog," [Online].
Available: https://www.datadoghq.com/blog/monitor-nvidia-gpus-with-datadog/

[9] DMTF, "Security Protocol and Data Model (SPDM)," DMTF Standards, [Online].
Available: https://www.dmtf.org/standards/spdm