DMTF
Technical Note
January 2023

Platform Management Communications Infrastructure (PMCI) – Scalable Architecture for Modern Platforms

Introduction

Over the last several years platforms have become increasingly complex. Standardized communication between devices and management elements is critical to obtain platform adoption across multiple implementations while reducing common complaints from system integrators.

DMTF's Platform Management Communications Infrastructure (PMCI) Working Group defines standards to address "inside the box" communication interfaces between the components of the platform management subsystem. These specifications provide a comprehensive, common architecture for improved communication between management subsystem components. In a nutshell, **it is a scalable architecture designed for modern platforms** that addresses the needs of system integrators through a simple standardized way.

By using PMCI-based standards, customers can expect reduced downtime, a secure and reliable platform, lower total cost of ownership, and interoperability at both the system and component levels.

This Technical Note provides an overview of the PMCI architecture and the benefits of each standard. Developers, implementers, and end users are encouraged to visit the PMCI's Working Group Page - https://www.dmtf.org/standards/pmci - for more in-depth information.

Why PMCI?

The group develops several specifications -- Management Component Transport Protocol (MCTP), Network Controller Sideband Interface (NC-SI), Platform Level Data Model (PLDM), and Security Protocol and Data Model (SPDM) -- that provide a comprehensive, common architecture for improved communication between management subsystem components. These specifications do not depend on a host or Operating System (OS) interaction thus enabling monitoring and control independent of the OS. For example, this greatly simplifies the firmware inventory and update process while accelerating delivery. Additionally, PMCI's internal-facing standards and technologies can be used by implementations of external-facing standards, including Redfish, for improved interoperability.

PMCI's multiple standards for these communications are detailed in separate specifications as noted above, to deliver maximum flexibility in implementation. Together, they provide a comprehensive, common architecture for platform management subsystem communication, which is an essential component for an interoperable end-to-end management solution -- ***a scalable architecture designed for modern platforms.***

DMTF's PMCI Standards

In the following section, each standard will be explained at a high level while highlighting the benefits of each one.

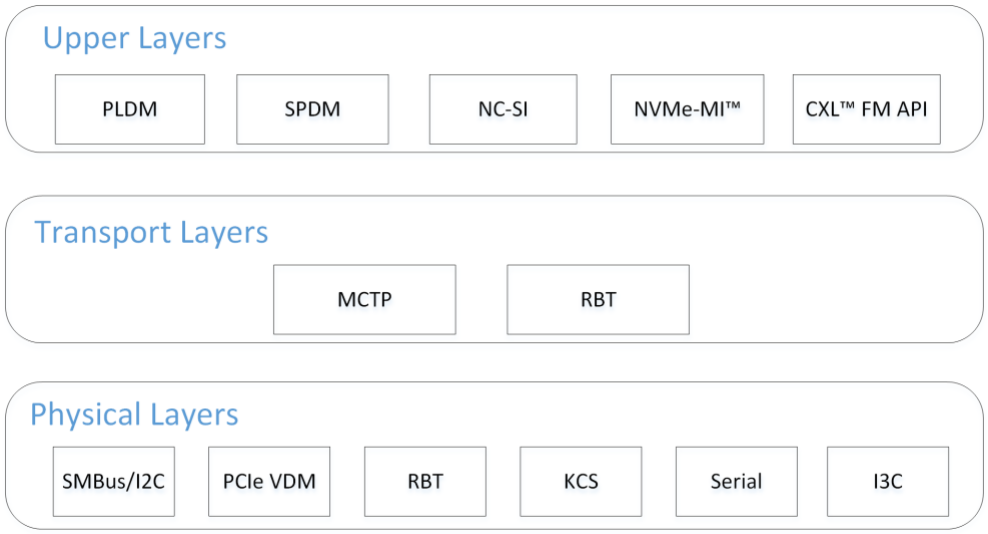

Management Component Transport Protocol (MCTP)

The Management Component Transport Protocol (MCTP) defines a communication model and can be used over multiple types of physical layers (e.g., SMBus/I2C, PCIe). It is intended to be used for intercommunication between elements of platform management subsystems used in computer systems and is a structural transport layer that is truly extendable and suitable for use in mobile, desktop, workstation, and server platforms. Management controllers such as a baseboard management controller (BMC) can use this protocol for communication between one another, as well as for accessing managed devices (e.g., sensors, power supplies, GPUs, NICs and disk drives) within the platform thus increasing system availability and business operational performance.

MCTP is used by the management controller, which sits in a server, appliance, or similar device, to interact with devices being managed to monitor everything from device state and statistics, and hold management parameters (e.g., speed, power state and utilization, link state, error count, uncorrectable error count, etc.). Therefore, the management controller aggregates parameters from one or more devices and gives access to those parameters available to local or remote software, or to other management controllers.

Providing out-of-band management, MCTP removes dependence on in-band agents, allowing the BMC to monitor devices without the presence of an operating system. This has an added benefit of releasing CPU cycles for other computing tasks since the impact to host performance is removed. Using MCTP can provide the most relevant data through an intuitive user interface and a centralized location to access the monitoring data. This can eliminate the need for multiple monitoring agents.

MCTP is a critical component for remote system management – it ubiquitously connects anything inside the platform - providing flexibility and performance benefits as well as helping to maximize uptime, reduce costs and provide a sense of sanity for busy IT professionals.

Network Controller Sideband Interface (NC-SI)

NC-SI is a common interoperable sideband interface and protocol used to transfer management traffic between a management controller and a network controller (NC). NC-SI defines the functionality and behavior of the Sideband Interface responsible for connecting the NC (including Ethernet, Fibre Channel, and InfiniBand controllers) to the management controller and can be transported via MCTP or RMII based Transport. The NC-SI enables a shared NIC model, which eliminates the need for separate management of network traffic as well as the need for additional switches, cables or RAC infrastructure.

NC-SI enables a shared NIC model by leveraging the pass-through communication as defined in the specification and control protocol. By combining your management LAN and your data LAN together an end user will not need to separate management traffic and user data, hence lowering the overall total cost of ownership and infrastructure reduction for the end user. By implementing NC-SI, system management will be more efficient, consume lower power, etc. -- equaling cost savings.

Platform Level Data Model (PLDM)

PLDM is a set of complementary specifications where each one can be used independently or together to provide a more robust systems management capability. These specifications actively help solve end-user concerns in a common, standardized way. It's an effective interface and data model providing efficient access to low-level platform inventory, BIOS, configuration data, monitoring/control, alerting, event log, data/parameters transfer functions, firmware update, Redfish enablement for managed devices (RDE), etc. For example, temperature, voltage, or fan sensors can have a PLDM representation that can be used to monitor and control the platform by using a set of PLDM messages thus allowing for better operational control over fans and power reduction, etc. By utilizing PLDM standards, customers can expect reduced downtime, a secure and reliable platform, lower total cost of ownership, and interoperability at both the system and component levels.

PLDM is a family of specifications providing a variable set of configuration options and capabilities including:

- SMBIOS – describes a model for transferring the tables or SMBIOS data structure
- Monitoring and Control – addresses sensor models, device events, and configuration
- Firmware Update – supports firmware updates of devices by using a standard method for obtaining current firmware version details, transferring a new code image to the device, and a consistent packaging format regardless of what type of device is being updated
- Redfish Device Enablement (RDE) – extends the Redfish interface to devices by using PLDM. It enables a management controller to present Redfish-conformant management of I/O adapters in a server, without the need for code specifics to each adapter family/vendor/model.

The complimentary specifications (used together or independently) enable systems to be more efficient, more effective, and ultimately, drive lower costs for end users.

Security Protocol and Data Model (SPDM)

Platform security is becoming increasingly important. As platform firmware components have become a new area for attacks, DMTF has developed SPDM to address these challenges. Developed by the Security Task Force within the PMCI Working Group, DMTF has created **THE** platform security protocol.

What does this mean exactly? When implemented properly, SPDM can ensure a complete chain of trust for the platform. SPDM defines messages, data objects, and sequences for performing message exchanges between devices over a variety of transport and physical media. The description of message exchanges includes authentication of hardware identities, measurement for firmware identities and session key exchange protocols to enable confidentiality and integrity protected data communication. SPDM enables efficient access to low-level security capabilities and operations.

By using SPDM, management traffic inside the box over MCTP can be encrypted allowing management data inside the platform to be encrypted – like TLS/HTTP encrypts your traffic over the Internet – since this has been proven to be an attack surface.

Additionally, SPDM is leveraged by other industry specifications to create a common security framework. As part of our Alliance Partner program, organizations including CXL Consortium, HDBaseT Alliance, MIPI Alliance, Open Compute Project, PCI-SIG, and the Trusted Computing Group provide crucial input to SPDM thus benefiting the industry at large.

Other mechanisms, including both non-DMTF and DMTF-defined mechanisms, can use SPDM specifications. A sample implementation from PMCI's SPDM Code Task Force can be found here: https://github.com/DMTF/libspdm.

Interoperability and Test Tools

As you can see, PMCI has an entire stack of protocols, many of which are dependent on each other -- and all of which are fundamentally dependent on both sides of a wire, a device and a management controller, or a device and another device -- speaking the same protocol. To obtain interoperability, it is critical that DMTF ensures devices can speak to these protocols correctly, hence the need for test tools. The PMCI's Test Tools Task Force is actively developing tools that will be used to demonstrate whether a device, has in fact, implemented the protocol correctly.

These tools will be useful for not only firmware developers of the device that is under test, but will also be useful for quality assurance groups, device manufacturers, and systems integration OEMs as they will provide acceptance tests for firmware drops and vendors.

DMTF is striving to enable heterogeneous support and interoperable solutions so that the end user can rest assured that when a company specification sheet says, "NC-SI (or another PMCI standard) has been tested at the same level as another industry spec sheet" it is verifiable and true. This allows end users the ability to determine their data center selections based on other merits, and not lower-level considerations.

Conclusion

The PMCI suite of standards enables complete platform and system management by providing standardized options that are easy to manage and more efficient. These capabilities are the 'invisible gems' that exist within data center systems that end users may not even know about. Because the specifications are not externally exposed but rather managing everything 'inside the box,' PMCI standards fundamentally increase customer and end user satisfaction by providing support and solutions such as a shared NIC, firmware updates, sensors, etc. in a simplified, standardized way – **providing scalable architecture for modern platforms.**

Acknowledgements

To learn more about the PMCI Working Group and the platform management standards it defines, or to get involved in this work, please visit https://www.dmtf.org/standards/pmci. Detailed information on all DMTF standards can be found at www.dmtf.org/standards. For a list of our current Alliance Partners visit https://www.dmtf.org/about/registers. Those interested in supporting and joining DMTF's efforts can learn more at www.dmtf.org/join.