

Intel[®] Rack Scale Design (Intel[®] RSD)

Getting Started Guide
Software v2.5

July 2019

Revision 001



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and noninfringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications.

This document contains information on products, services, and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications, and roadmaps.

Copies of documents that have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2019 Intel Corporation. All rights reserved.



Contents

1.0	Introduction	5
1.1	Intended Audience	5
1.2	Background and Prerequisite Information	5
1.3	Conventions	6
1.4	Notes and Symbol Convention	6
1.5	Terminology	6
1.6	References and Resources	7
2.0	Intel® RSD Conformance Overview	9
3.0	Intel® RSD Reference Code	11
4.0	Intel® RSD Product Design Prerequisites	12
5.0	New Intel RSD Conformance in v2.5	13
5.1	CTS Prerequisites (PODM)	13
5.2	CTS Prerequisites (PSME/RMM)	14

Figures

Figure 1.	Intel® RSD Co-Development Process and Conformance Process Flow	9
Figure 2.	Intel® RSD and Standards	10

Tables

Table 1.	Terminology	6
Table 2.	Reference Documents and Resources	7



Revision History

Revision	Description	Date
001	Intel® Rack Scale Design Software v2.5 initial release	July 2019

§



1.0 Introduction

This guide is the starting point for developers planning to work with Intel® RSD software and conformance testing. Intel recommends reading the entire guide before starting.

1.1 Intended Audience

The intended audience for this specification includes designers and engineers working with the Intel® RSD Software .v2.5 release, such as Independent Software Vendors (ISVs), Original Equipment Manufacturers (OEMs), and customers who are considering building their own RSD implementation

1.2 Background and Prerequisite Information

This guide provides background and prerequisite information for the Intel® RSD v2.5 documentation and the Intel® RSD conformance process. The information is provided to make sure the process goes smoothly and efficiently.

Note: The Intel® RSD code is reference software only. Developers are expected to modify the software and make it their own.

The following steps outline a summary of all available Intel® RSD materials:

1. Intel® RSD PSME/ Rack Management Module (RMM) reference code provides fully functional implementation to communicate with the PODM, northbound REST interface exposing Redfish*-aligned APIs, manage and report power/thermal data to the PODM, and RMM implementations. It also includes stubs for the PSME network, compute, and chassis agents.
2. Intel® RSD PSME Storage Service reference code includes fully functional remote storage service implementations with northbound REST APIs and creates initial internet Small Computer System Interface (iSCSI) targets on service initiation.
3. Read the Intel® Rack Scale Design API Software Specifications listed in [Table 2](#). Then plan the configuration of the Intel® RSD software components across your hardware.
4. Once the hardware configuration is established, decide which servers in the rack configurations will run which Pooled System Management Engine (PSME) and PODM components (also known as agents). For example:
 - a. Dedicate an Ubuntu* v16.04 server with a Baseboard Management Controller (BMC); to run the PSME compute and core Representational State Transfer (REST) interface modules.
 - b. Dedicate a 10 GbE Top of Rack (ToR) or another switch to run the PSME core and networking modules.
 - c. Dedicate another storage server (disk controller) to run the PSME core REST APIs and storage agents/modules. This storage server could be the same dedicated server (with a BMC, as the server could also have a storage controller and additional disks.

Note: For more information on PSME/RMM configurations and example code, refer to *Intel® Rack Scale Design Pooled System Management Engine User Guide* and *Intel® Rack Scale Design Rack Management Module Representational State Transfer API Specification* in [Table 2](#).

5. Read the Intel® Rack Scale Design API Software Specifications listed in [Table 2](#). Understanding these references help to work with the functional code provided, including:
 - a. The Intel® RSD PODM reference code includes a fully functional northbound interface exposing Redfish*-aligned APIs along with code to discover, compose, and manage Intel® RSD resources.



6. Build, install, and modify the PSME components on the hardware configuration and modify the GAMI agents to interface with the hardware configuration listed in *Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) User Guide* [Table 2](#).
7. Build, install, and modify the PODM components to talk with the PSME agents and manage the racks (refer to the *Intel® Rack Scale Design (Intel® RSD) POD Manager (PODM) Representational State Transfer (REST) User Guide* listed in [Table 2](#)).
8. Read the *Intel® Rack Scale Design (Intel® RSD) PODM Release Notes* and *Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) Release Notes* to be aware of potential issues. Contact an Intel® RSD account representative or visit the Intel® RSD website (<http://intel.com/intelrdsd>) if you run into issues, have questions, or want to provide general feedback.

1.3 Conventions

The key words/phrases "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in *Key words for use in RFCs to Indicate Requirement Levels*, March 1997, RFC 2119, refer to [Table 2](#).

1.4 Notes and Symbol Convention

Symbol and note conventions are similar to typographical conventions used in the *Cloud Infrastructure Management Interface (CIMI) Model and RESTful HTTP-based Protocol and Interface for Managing Cloud Infrastructure*, DSP0263, [Table 2](#) notation used in JSON* serialization description:

- Values in italics indicate data types instead of literal values.
- Characters are appended to items to indicate cardinality:
 - ? (0 or 1)
 - * (0 or more)
 - + (1 or more)
 - Vertical bars, |, denote choice. For example, a|b means a choice between a and b.
 - Parentheses, (), indicate the scope of the operators ?, *, +, and |.
- Ellipses, ..., indicate points of extensibility. The lack of an ellipsis does not mean no extensibility point exists; rather, it is just not explicitly called out.

1.5 Terminology

Table 1. Terminology

Term	Definition
ACL	Access Control List
API	Application Programming Interface
BMC	Baseboard Management Controller
CA	Certificate Authority
CIMI	Cloud Infrastructure Management Interface
CTS	Conformance Test Suite
GAMI	Generic Assets Management Interface
IPMI	Intelligent Platform Management Interface
iSCSI	Internet Small Computer System Interface
ISV	Independent Software Vendor
LAG	Link Aggregation Group
MAC	Media Access Control



Term	Definition
NIC	Network Interface Card
NVMe*	Non-Volatile Memory express*
NVMe-oF*	NVMe over Fabrics*
OEM	Original Equipment Manufacturer
PCIe*	Peripheral Component Interconnect express*
PNC	Pooled NVMe Controller
POD	A physical collection of multiple racks
PODM	POD Manager
PSME	Pooled System Management Engine
REST	Representational State Transfer
RDMA	Remote Direct Memory Access
RMM	Rack Management Module
RSD	Rack Scale Design
SLED	Single Large Expensive Disk
SMBIOS	System Management BIOS
SSD	Solid-State Drive
ToR	Top of Rack
VLAN	Virtual Local Area Network

1.6 References and Resources

Table 2. Reference Documents and Resources

Doc ID	Title	Location
613314	Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) User Guide Software v2.5	Note: https://www.intel.com/content/www/us/en/architecture-and-technology/rack-scale-design/rack-scale-design-resources.html
613316	Intel® Rack Scale Design (Intel® RSD) POD Manager (PODM) Release Notes Software v2.5	
613317	Intel® Rack Scale Design (Intel® RSD) POD Manager (PODM) User Guide Software v2.5	
613318	Intel® Rack Scale Design (Intel® RSD) Pooled System Management (PSME) Release Notes Software v2.5	
613329	Intel® Rack Scale Design Storage Services API Specification Software v2.5	
613319	Intel® Rack Scale Design (Intel® RSD) Architecture Specification Software v2.5	
613320	Intel® Rack Scale Design (Intel® RSD) Pod Manager (PODM) Representational State Transfer (REST) API Specification Software v2.5	
613321	Intel® Rack Scale Design (Intel® RSD) Rack Management Module (RMM) Representational State Transfer (REST) API Specification Software v2.5	
613324	Intel® Rack Scale Design (Intel® RSD) Generic Assets Management Interface (GAMI) API Specification v2.5	
613325	Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) REST API Specification Software v2.5	
613326	Intel® Rack Scale Design (Intel® RSD) Conformance Test Suite (CTS) Release Notes	
608298	Field Programmable Gate Array (FPGA) over Fabric Protocol Architecture Specification	https://cdrdv2.intel.com/v1/dl/getContent/608298
596167	Intel® Rack Scale Design (Intel® RSD) for Cascade Lake Platform Firmware Extension Specification	https://cdrdv2.intel.com/v1/dl/getContent/596167



Doc ID	Title	Location
DSP2046	<i>Redfish* Resource and Schema Guide v2018.3</i>	https://www.dmtf.org/sites/default/files/standards/documents/DSP2046_2018.3.pdf
DSP0263	<i>Cloud Infrastructure Management Interface (CIMI) Model and RESTful HTTP-based Protocol</i>	https://www.dmtf.org/sites/default/files/standards/documents/DSP0263_2.0.0.pdf
DSP0266	<i>Redfish* Scalable Platforms Management API Specification v1.6.1</i>	https://www.dmtf.org/sites/default/files/standards/documents/DSP0266_1.6.1.pdf
DSP8010	<i>Redfish* Schema 2018.3</i>	https://www.dmtf.org/sites/default/files/standards/documents/DSP8010_2018.3.zip
RFC2119	<i>Key Words for Use in RFCs to Indicate Requirement Levels</i>	https://www.ietf.org/rfc/rfc2119.txt
N/A	<i>NVM Express over Fabrics 1.0</i>	http://nvmexpress.org/wp-content/uploads/NVMe_over_Fabrics_1_0_Gold_20160605-1.pdf
N/A	<i>Intel® RSD Conformance Test Suite (CTS) tool binary</i>	https://github.com/intel/intelRSD/tree/master/CTS
N/A	<i>Swordfish* Scalable Storage Management API Specification</i>	https://www.snia.org/sites/default/files/technical_work/Swordfish/Swordfish_v1.0.7a_Specification.pdf

NOTE: Documents referenced in this table which have a Doc ID, but cannot be accessed, can be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm obtain a copy.



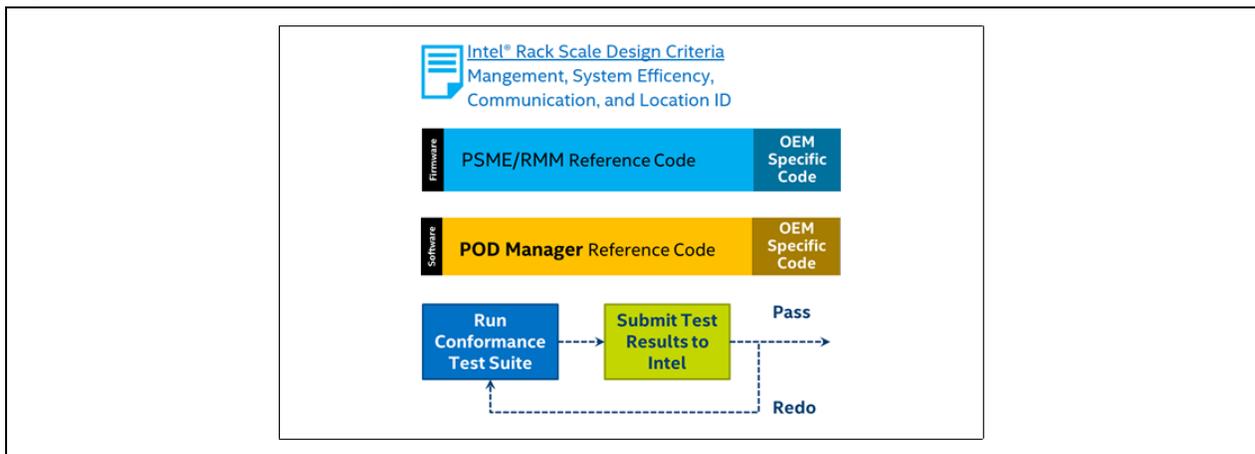
2.0 Intel® RSD Conformance Overview

The Intel® Rack Scale Design (Intel® RSD) Conformance and Software Reference Kit covers implementing the required functions of the Intel® RSD reference architecture, including hardware, software, system setup, and standards-based APIs. Conformance drives alignment to industry standards (driving those standards into the market), builds the Intel® RSD ecosystem, builds end customer assurance, and reduces product development costs by improving time to market of new technologies.

Engagement with the Intel® RSD ecosystem, illustrated in [Figure 1](#), involves the following:

- Releasing the feature complete “beta” Intel® RSD software v2.5 to the open-source community.
- Support for Original Equipment Manufacturer (OEM) partners to complete development of the Intel® RSD software with OEM hardware-specific code and optional OEM features.
- Working with third-party independent software vendors (ISVs) to enable orchestration solutions to work with the Intel® RSD software.

Figure 1. Intel® RSD Co-Development Process and Conformance Process Flow

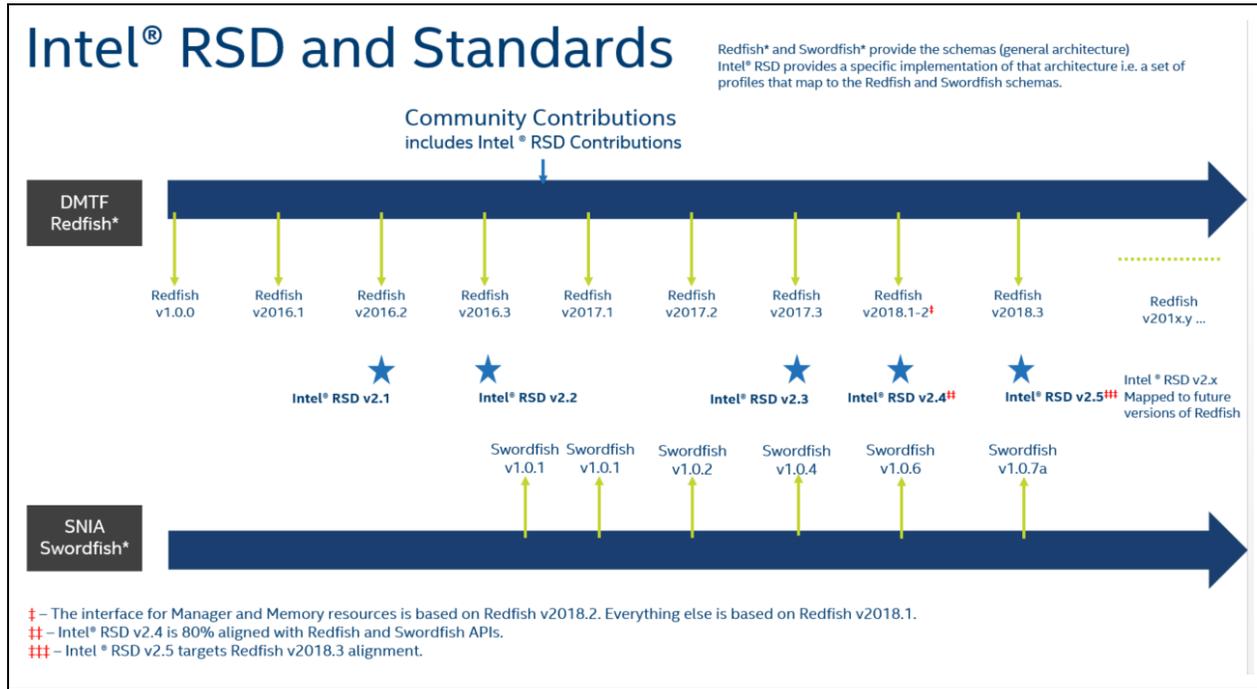


The Intel® RSD Conformance Test Suite (CTS) tool verifies the northbound API schema conformance of the PSME/RMM and PODM through automated techniques. The tool can also test some hardware and software parameters.

Note: Intel recommends completing the full conformance testing process with a manual review of the remaining system and hardware parameters. Also, validate the rack architecture to confirm the hardware implementation is completed, as described in the *Intel® Rack Scale Design (Intel® RSD) Architecture Specification* (refer to [Table 2](#)).

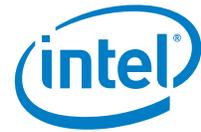
Intel® RSD is mapped to the schema and models standardized by specific Redfish* versions. [Figure 2](#) shows how Intel® RSD contributes to the Redfish* community and aligns with Redfish* releases.

Figure 2. Intel® RSD and Standards



Note: The Intel® RSD release v2.5 is based on the Redfish* Specification v1.6.1. It uses the Redfish* Schema v2018.3 and Swordfish* Schema v1.0.7a (refer to [Table 2](#)).

§



3.0 Intel® RSD Reference Code

Along with alignment to Redfish* schemas, Intel provides fully functional Intel® RSD reference code for the following items:

- Intel® RSD PODM:
- A northbound REST interface to expose Redfish*-aligned APIs
- Discover, compose, and manage Intel® RSD resources

Intel® RSD PSME/RMM:

- PSME/RMM implementation to communicate with the Intel® RSD PODM
- A northbound REST interface to expose Redfish-aligned APIs
- Manage and report the power and thermal matrix to the PODM
- Firmware (FW) extensions (conversion of Intelligent Platform Management Interface)
- Intelligent Platform Management Interface (IPMI) to Redfish

Intel® RSD PSME Storage Service:

- Remote storage service implementation with northbound REST APIs
- Intel® RSD Reference code provides stubs for PSME Network, Compute, and Chassis agents

Note: The Intel® RSD reference code is reference software only. The Intel® RSD Reference code does NOT implement all required criteria that are listed in the *Intel® RSD Architecture Specification* (refer to [Table 2](#)). Additionally, backward compatibility of RSD v2.5 with previous RSD version racks and components have not been tested. For any further questions, contact your Intel account representative.



4.0 Intel® RSD Product Design Prerequisites

This section reviews information to keep in mind at the start of the Intel® RSD product design process. This information is designed to help ensure smooth conformance testing of API design, hardware, and software at the end of the development cycle.

For Intel® RSD conformance, Intel engages with partners in the following ways:

- Both Intel and partners collaborate and confirm the implementation of the Intel® RSD architecture.
- Intel and partners collaborate to complete the Intel® RSD hardware checklist document.

The checklist covers the required items from the Intel® RSD Architecture Specification. For example, Intel audits the power supply and fan numbering consistency, compute blade serviceability, and node reset support.

Intel® RSD v2.5 includes required functionalities across the PODM and PSME/RMM. Each required functionality may include the implementation of one or more APIs. To complete Intel® RSD conformance, all required APIs are expected to be implemented for all required functionality.

§



5.0 New Intel RSD Conformance in v2.5

Note: This section contains more advanced information on performing an Intel RSD® conformance test. If you are a beginner on Intel® RSD or Intel® RSD conformance, you may consider this chapter as future reference.

The next step is to test for Intel® RSD conformance. Some preparation is required before testing. To begin the process, complete the following tasks:

1. Download the *Intel® RSD v2.5 Conformance Test Suite (CTS) tool binary* from GitHub by using the following link, selecting “CTS” and clicking on “Download ZIP” <https://github.com/intel/intelRSD>.
Alternatively, clone the entire Intel® RSD repository using the command `git clone`:
<https://github.com/intel/intelRSD.git>
2. Install CTS by following the steps documented in the README.md here <https://github.com/intel/intelRSD/blob/master/CTS/README.md>
3. Populate the rack with the required hardware and software components:
 - **Hardware:** Compute blades, storage servers, PCIe* devices, ToR (Ethernet) switch, PCIe switch, cables, and power and thermal units
 - **Software:** PODM and PSME (network, compute, storage, RMM, PNC, and NVMe*) agents.
4. Enable the Certificate Authority (CA) Authentication to establish a secure communication connection (refer to the Appendices of *Intel® RSD PODM User Guide* and *Intel® RSD PSME User Guide* in [Table 2](#) for details on the PODM CA).
5. Execute the `GET/PATCH/CRUD` actions through the CTS against implemented RSD PSME modules
6. Complete the PODM conformance testing prerequisites described in this section.

5.1 CTS Prerequisites (PODM)

For the CTS test to run successfully, setup tasks are required for both PODM and PSME/RMM before testing. These tasks streamline the testing and feedback processes for all users, refer to [Table 2](#), PODM and PSME API specifications for details.

Before running the CTS tool for PODM, complete the following tasks:

1. Create an Intel® RSD node with:
 - Remote storage (iSCSI targets and NVMe drives),
 - Associate all Network Interface Cards (NICs) with the nodes (1 G and 10 G)
2. Create a VLAN if it is not already created. The VLAN needs to be able to communicate with the PODM from the CTS PC.
3. Compose nodes with a remote FPGA.
4. Make sure the CPU/Memory/Drive/Ethernet matrix from a computer system is populated under the PODM north bound APIs after composing the RSD node.
 - The RSD specification requires SMBIOS implementation to share the above information with the PODM north bound APIs.
 - Create multiple logical nodes with parameters defined in Section 6 of the *Intel® PODM API Specification* (refer to [Table 2](#)) with CPU, memory, drive (local and remote), and Ethernet type parameters (for example link speed, etc.)
5. Verify parent/child hierarchy (relationship in the rack).
 - a. Patch the rack “ID” through the PODM API, and paste the snippet to show the Rack ID change is reflected in the Chassis “ParentID” field.



- b. Compose a new Intel® RSD node, and verify that the new Rack ID is displayed in the "ParentID" field.
 - c. Patch an existing node with an NVMe drive to verify that a new RackID is displayed in the ParentID field.
 - d. Provide two separate CTS logs to demonstrate Rack location ID is changed.¹
6. At least two logical chassis MUST exist to verify unique chassis location to meet the hierarchy requirement.
 7. The user MUST cover the following two scenarios and submit separate CTS logs for each of these actions:
 - a. Create a brand new node and verify the node has a unique chassis location ID and hierarchy.
 - b. Patch an existing RSD node with an NVMe drive to verify the node has a unique chassis location ID and hierarchy.

5.2 CTS Prerequisites (PSME/RMM)

Perform the following steps before running the CTS tool for the PSME/RMM:

1. Certificate Authority (CA) authentication enabled
With CTS log of PSME/RMM service, it needs to show CA authentication be used instead of user name/password authentication only
 2. Create an Access Control List (ACL)
 - a. Link ACL rule to a switch port. This operation can be performed from the PSME Ethernet switch API
 - b. ACL rule type MUST have an action (For instance: Deny, Mirror or Forward)
 3. Create a Virtual Local Area Network (VLAN)
This VLAN is created on the Ethernet switch. The user should associate RSD logical node(s) with this VLAN. This operation can be performed from Ethernet CLI
 4. Create a Link Aggregation Group (LAG)
Verify the LAG information is visible through Ethernet switch APIs & on Ethernet switch ports
 5. Create `StaticMAC`.
Post creation, verify `StaticMAC` is created. This operation can be performed from the Ethernet switch API
 - a. The RSD logical node Media Access Control (MAC) address should be visible as "NeighborMAC" on the switch port
 - b. Initiate the traffic from the RSD node and let the switch learn `NeighborMAC`.
- Note:** In case leaf and spine switch architecture are implemented, the user MUST support the features (as mentioned in the Intel® RSD PSME API specification, Section 4.25, refer to [Table 2](#)) for all connected switches.
6. If pooled the NVMe/FPGA functions are implemented, the related service API MUST be supported.
 7. Demonstrate PSME Event Service subscription to verify `EventService` implementation.

When the tasks are completed, run the CTS tool and tests. Execute the `GET`, `PATCH`, and `CRUD` options with an individual command and confirmation. Also, generate a separate log file for review through the CTS tool against the implemented Intel® RSD modules (PODM, PSME, and RMM). Send the data and any support questions to an Intel® RSD representative.

¹ Provide CTS logs for all scenarios and the CTS tool will need to be re-initiated to achieve all actions.



Intel representatives review the Intel® RSD conformance test results and provide feedback and troubleshooting guidance. If you find any errors with the CTS tool or process, contact your Intel® RSD representative for troubleshooting and technical assistance.

As an exception, Intel does have an Intel® RSD conformance waiver process. Conformance waivers can be investigated on a case-by-case basis. Work with your Intel® RSD representative for waiver questions and support.

Once all conformance tests pass (manual and automated), Intel grants Intel® RSD Conformance.

§