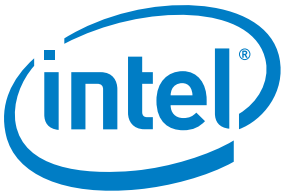

Intel® Rack Scale Design (Intel® RSD) Conformance and Software Reference Kit

Getting Started Guide
Software v2.4

April 2019

Revision 001

Document Number: 608487-001

Contents

Figures

Tables

Revision History

Revision	Description	Date
001	Initial release	April 2019

§

1.0 Introduction

The *Intel® Rack Scale Design (Intel® RSD) Conformance and Software Reference Kit Getting Started Guide* cover implementing the required functions of the Intel® RSD reference architecture, including hardware, software, system setup, and standards-based APIs.

1.1 Intended Audience

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.

1.2 Background and Prerequisite Information

This guide provides background and prerequisite information for the Intel® RSD v2.4 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. Read the Intel® Rack Scale Design API Software Specifications listed in Table 2. and then plan the configuration of the Intel® RSD software components across your hardware.
2. Once the configuration across your hardware has been established, decide which servers in your rack configurations will run which PSME and PODM components (also known as agents). For example:
    a. Dedicate an Ubuntu* v16.04 server with a Baseboard Management Controller (BMC); (for example, 1U on any rack) to run the PSME compute and core REST interface modules. For example, Dell* PowerEdge* and HP* ProLiant* have this capability.
    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) described in Step 4.a above, as the server could also have a storage controller and additional disks.

    Note: The Intel® RSD v2.4 reference software provides secure erase on decomposition security features. Intel® SSD DC P4500/4600 Series memory (previously known as Cliffdale) (4500/4600) or later solid-state drives (SSDs) are required to support the Intel RSD v2.4.
3. Contact an Intel® RSD account representative or visit the Intel® RSD website (refer to Table 2) to acquire the required code.
4. 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. 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.
    b. Intel® RSD PSME/RMM reference code provides fully functional implementation to communicate with 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.

c. 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.

5. Build, install, and modify the PSME components on the hardware configuration (refer to the *Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) User Guide* listed in Table 2). Modify the GAMI agents to interface with the hardware configuration.
6. 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).
7. Connect to an outside orchestration layer if one is used (for example, OpenStack*).
8. Read *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 (refer to Table 2) 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 6 (CIMI) Model and RESTful HTTP-based Protocol 7* 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
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

Term	Definition
iSCSI	Internet Small Computer System Interface
ISV	Independent Software Vendor
LAG	Link Aggregation Group
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
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

1.6 References and Resources

Table 2. Reference Documents and Resources

Doc ID	Title	Location
608486	*Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) User Guide Software v2.4*	Note: *https://www.intel.com/content/www/us/en/architecture-and-technology/rack-scale-design/rack-scale-design-resources.html*
608488	*Intel® Rack Scale Design (Intel® RSD) POD Manager (PODM) Release Notes Software v2.4*	
608489	*Intel® Rack Scale Design (Intel® RSD) POD Manager (PODM) User Guide Software v2.4*	
608490	*Intel® Rack Scale Design (Intel® RSD) Pooled System Management (PSME) Release Notes Software v2.4*	
608491	*Intel® Rack Scale Design Storage Services API Specification Software v2.4*	
608492	*Intel® Rack Scale Design (Intel® RSD) Architecture Specification Software v2.4*	
608493	*Intel® Rack Scale Design (Intel® RSD) Pod Manager (PODM) Representational State Transfer (REST) API Specification Software v2.4*	
608494	*Intel® Rack Scale Design (Intel® RSD) Rack Management Module (RMM) Representational State Transfer (REST) API Specification Software v2.4*	
608495	*Intel® Rack Scale Design (Intel® RSD) Generic Assets Management Interface (GAMI) API Specification v2.4*	
608496	*Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) REST API Specification Software v2.4*	
608497	*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
N/A	*NVM Express over Fabrics 1.0*	*http://nvmexpress.org/wp-content/uploads/NVMe_over_Fabrics_1_0_Gold_20160605-1.pdf*
N/A	*Intel® RSD Conformance Test Suite (CTS) tool binary*	*https://github.com/intel/intelRSD/tree/master/CTS*

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. Intel® RSD v2.4 (CTS_2.4) release can produce an Interoperability report between the *Rack Scale Metadata and Redfish* v2018.1, v2018.2, v2018.3 (with *Swordfish* v1.0.6), refer to Table 2 for details.

Engagement with the Intel® RSD ecosystem, illustrated in Figure 1, involves the following:

- Releasing the feature complete "beta" Intel® RSD software v2.4 to the open-source community.
- Support for Original Equipment Manufacturer (OEM) Partners to complete development of the Intel® RSD software (PSME/RMM and PODM) 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

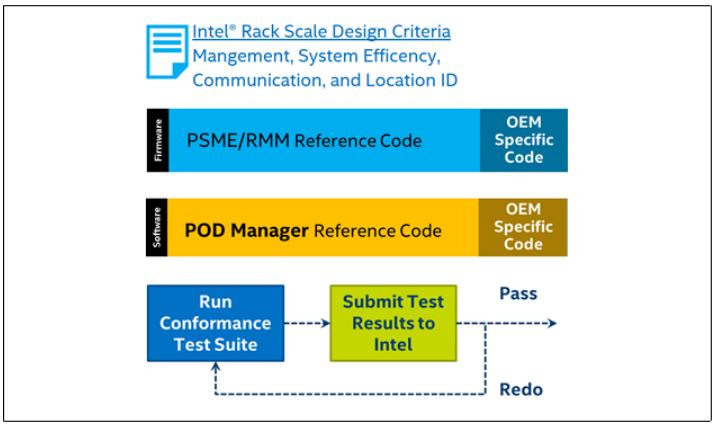

The Intel® RSD CTS tool verifies the northbound API schema conformance of the PSME/RMM and PODM through automated techniques. The tool can also test some hardware and software parameters.

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 Specificat*ion (refer to Table 2).

When discussing Intel® RSD conformance, many questions arise regarding how Intel® RSD APIs align with Redfish* APIs. 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 Mapping to Redfish* Releases

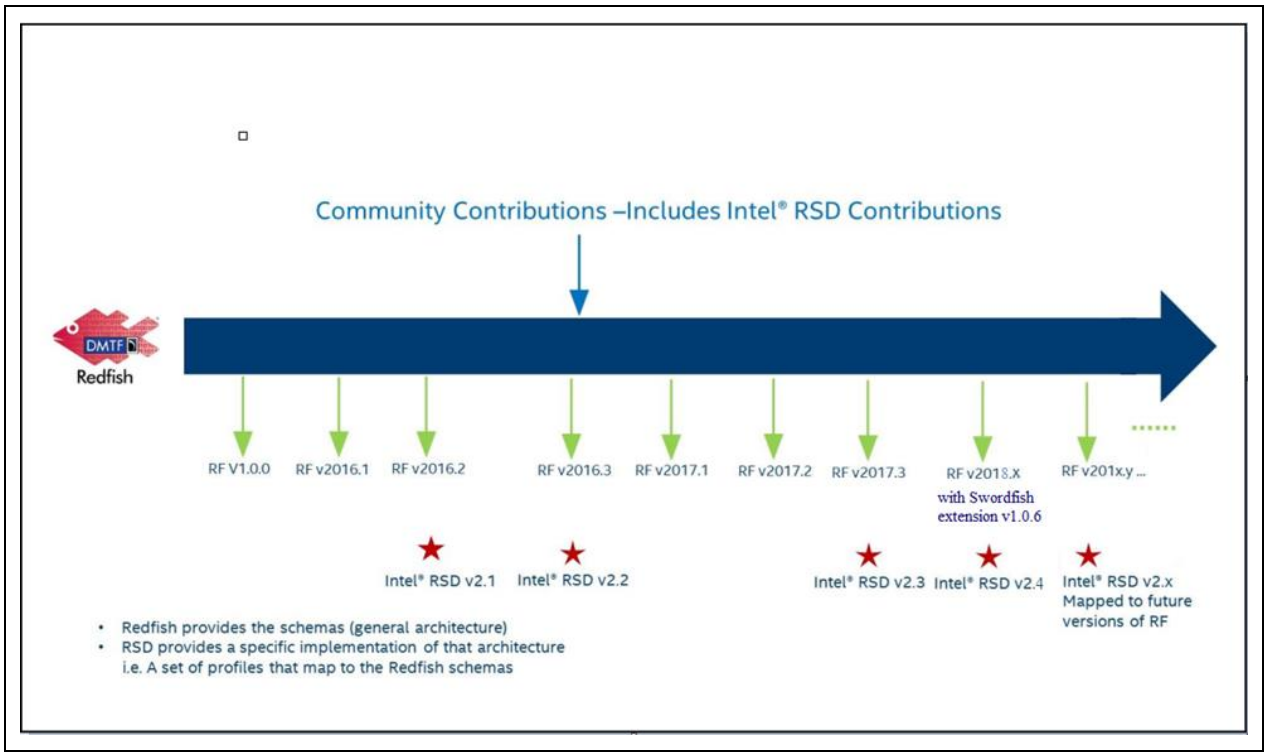

§

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 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
- RMM Implementations
- 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
- Create initial Internet Small Computer System Interface (iSCSI) targets on service initiation
- Intel® RSD Reference code provides stubs for PSME Network, Compute, and Chassis agents

The Intel® RSD Reference code does NOT include all required Intel® RSD elements noted in the Intel® RSD Architecture Specification Software (refer to Table 2). Contact your Intel account representative with questions regarding the Intel® RSD reference code and architecture requirements.

§

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 hardware conformance, Intel engages with partners in the following ways:

- Architects from both Intel and partners collaborate and confirm the implementation of the Intel® RSD architecture.
- The Intel ecosystem engineers 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, node reset support, and so forth.

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

- **Software Architecture**: Architects from both Intel and partners collaborate and confirm the implementation of the Intel® RSD software architecture. Implementation details include (but are not limited to):
  - Where the PSME/RMM/BMC reside in the rack.
  - APIs the plans to provide in addition to the Intel® RSD APIs (needed extensions for a specific design).
- **Compute Single Large Expensive Disk (SLED) Telemetry**: For partners using Intel® RSD reference code, Intel includes a system management BIOS (SMBIOS) implementation to gather the needed telemetry data required for conformance (CPU, memory, BIOS, and network).

  For partners not using the Intel® RSD reference code, an equivalent implementation must be developed to report this data to pass conformance testing.
- **Location Hierarchy**: The PODM requires a location hierarchy between the chassis and SLED resources to locate the composed node resources.
- **Power and Thermal Telemetry**: Functionality to display the power and thermal matrix must be implemented for conformance. The parameters in the power and thermal matrix include active power supplies, total watts consumed, and so forth.
- **PSME Network Functionality**: PSME networking conformance depends on the implementation of the required PSME network APIs.
- **Non-Volatile Memory Express* (NVMe*) over Fabric (NVMe-oF*)**: Functionality depends on the implementation of the NVMe* PSME agent and NVMe discovery PSME agent. NVMe-oF* requires Ethernet as the fabric that uses Remote Direct Memory Access (RDMA) as the transport protocol between the compute and storages nodes.

Several prerequisites are required to ensure Intel® RSD API conformance. Intel® RSD v2.4 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 Test for Intel® RSD Conformance

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 v2.4 Conformance Test Suite (CTS) tool binary from GitHub* (refer to Table 2).
2. 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.
3. Enable Certificate Authority (CA) Authentication to establish a secure communication connection between the PODM and PSME/RMM.
4. Execute GET/PATCH/CRUD actions via CTS against implemented RSD modules (PSME, PODM, RMM)
5. Complete the PODM and PSME/RMM 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 the PDOM 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 the PDOM, complete the following tasks:

1. Create an Intel® RSD node with:
   - Remote storage (iSCSI targets and NVMe drives),
   - Associate all NIC with the nodes (1 G and 10 G)
2. Create a VLAN.
3. Create 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), Ethernet type parameter (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 snipped to show the Rack ID change is reflected in the Chassis `"ParentID"` field.
   b. Create 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 nodes MUST be created to verify unique chassis location ID and to verify unique chassis hierarchy requirement.

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

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)

Before running the CTS tool for the PSME/RMM, they will need to perform the following items:

1. 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 Access control list (ACL)
   - Link ACL rule to a switch port. This operation can be performed from PSME Ethernet switch API
   - ACL rule type MUST have an action (For instance: Deny, Mirror or Forward)
3. Create VLAN

   This VLAN is created on Ethernet switch. The user should associate RSD logical node(s) with this VLAN. This operation can be performed from Ethernet CLI

4. Create 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 Ethernet switch API

6. The RSD logical node MAC address should be visible as NeighborMAC on the switch port
   - Initiate the traffic from the RSD node and let the switch learn `NeightborMAC.`

   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.
7. If pooled the NVMe/FPGA functions are implemented, the related service API MUST be supported.
8. 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 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.

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 escalation.

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.

§