

Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME)

Release Notes
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 technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer.

The products described may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, 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.

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, Optane, 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	Conventions	5
1.3	Software Package Contents	5
1.4	Revision Numbers of Package Components	6
1.5	Terminology	6
1.6	References	7
2.0	Overview	8
2.1	New Features for Intel® RSD v2.5	8
2.2	Deprecation Notices	8
2.3	Limitations	8
2.4	Security Recommendations	9
3.0	Known Issues	11
4.0	Fixed Issues	16

Tables

Table 1.	Software package for Intel® RSD v2.5	5
Table 2.	Package Components	6
Table 3	Terminology	6
Table 4	Reference Documents and Resources	7
Table 5.	Security Recommendations	9
Table 6.	Status Definitions	11
Table 7.	Known Issues	11
Table 8.	Fixed Issues	16



Revision History

Revision	Description	Date
001	Initial release for Intel® RSD software release v2.5	July 2019

§



1.0 Introduction

These release notes are intended for the Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) v2.5 release. This software will be referred to as the PSME throughout this document.

1.1 Intended Audience

The intended audiences for this document include:

- Independent Software Vendors (ISVs) who make use of POD Manager (PODM) to discover, compose, and manage drawers, in a multivendor environment.
- Original Equipment Manufacturers (OEMs) who would like to provide Intel® RSD PODM API interface in their hardware platform on top of their hardware platform.

1.2 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*, RFC 2119 (refer to [Table 4](#)).


1.3 Software Package Contents

Table 1 lists the contents of the release package.

Table 1. Software package for Intel® RSD v2.5

Title	Description
Intel® Rack Scale Design (Intel® RSD) Pooled System Management (PSME) Release Notes Software v2.5	This document
Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) User Guide Software v2.5	User Guide
Intel® Rack Scale Design (Intel® RSD) Generic Assets Management Interface (GAMI) API Software v2.5	JSON-RPC* API specifications to communicate with GAMI Modules
Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) REST API Specification Software v2.5	PSME RESTful API Specifications
Intel® Rack Scale Design (Intel® RSD) Rack Management Module (RMM) Representational State Transfer (RESTful) API Specification Software v2.5	RMM API Specification
Intel® Rack Scale Design (Intel® RSD) Storage Services API Specification Software v2.5	Storage Service REST API Specifications
License.txt	Apache*License, Version 2.0

Customers should check <http://www.intel.com/intelRSD> to download the latest available onboard device drivers, system firmware, and system software. For further assistance, contact your Intel Field Representative.



1.4 Revision Numbers of Package Components

Table 2. Package Components

Subproject (component)	Revision
Intel® Rack Scale Design PSME	RSD_PSME_2.5

1.5 Terminology

Table 3 Terminology

Term	Definition
BDC	Bulldog Creek
BMC	Baseboard Management Controller
CA	Certificate Authority
CM	Control Module
DCBX	Data Center Bridging eXchange
DMTF	Distributed Management Task Force
ETS	Enhanced Transmission Selection
FPGA	Field Programmable Gate Array
FPGA-oF*	FPGA-over Fabrics* protocol
IPMI	Intelligent Platform Management Interface
IPsec	Internet Protocol Security
iSCSI	Internet Small Computer Systems Interface
ISV	Independent Software Vendor
LAG	Link Aggregation Group
MAC	Media Access Control
MMP	Management Mid Plane
NVMe	Non Volatile Memory Express
NVMe-oF*	NVMe over Fabrics*
Intel® PAC	Intel® Programmable Acceleration Card
PCIe*	PCI Express
PFC	Priority Flow Control
PSME	Pooled System Management Engine
PSU	Power Supply Unit
OEMs	Original Equipment Manufacturers
QoS	Quality of Service
QSFP	Quad Small Form-factor Pluggable
RDMA	Remote Direct Memory Access
RMM	Rack Management Module
RoCEv2	RDMA over Converged Ethernet v2
Intel® RSD	Intel® Rack Scale Design
SMBIOS	System Management BIOS
SPDK	Storage Performance Development Kit
TLS	Transport Layer Security
TORs	Top of Rack switches
TPM	Trusted Platform Module



1.6 References

Table 4 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
613315	Intel® Rack Scale Design (Intel® RSD) Conformance and Software Reference Kit Getting Started Guide v2.5	
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	
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	
613329	Intel® Rack Scale Design Storage Services API Specification Software v2.5	
613326	Intel® Rack Scale Design (Intel® RSD) Conformance Test Suite (CTS) Release Notes	See note
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
DSP8010	Redfish Schema v2018.3	https://www.dmtf.org/sites/default/files/Redfish_2018_Release_3_0verview.pdf
DSP8010	Redfish Schema v2018.2	https://www.dmtf.org/sites/default/files/standards/documents/DSP8010_2018.2.zip
DSP0266	Scalable Platforms Management API Specification v1.5.0	https://www.dmtf.org/sites/default/files/DSP0266_1.5.0.pdf
N/A	Swordfish* Scalable Storage Management API Specification Version 1.0.6	https://www.snia.org/sites/default/files/SMI/swordfish/v106/Swordfish_v1.0.6_Specification.pdf
N/A	Date and time format - ISO 8601	https://www.iso.org/iso-8601-date-and-time-format.html
RFC2068	Hypertext Transfer Protocol - HTTP/1.1	https://tools.ietf.org/html/rfc2616
RFC2119	Key Words for Use in RFCs to Indicate Requirement Levels, March 1997	https://ietf.org/rfc/rfc2119.txt
	Storage Performance Development Kit	https://github.com/spdk/spdk

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 Overview

This section provides an overview of the new features and Security Recommendations for the Intel® RSD v2.5 release.

2.1 New Features for Intel® RSD v2.5

- Log Service – Service that exposes BMC System Event Log (SEL) for each BMC in the form of Redfish Log Entries.
- Redfish 2018.3 Schema alignment:
 - Moved Telemetry Service, Metric Definition, Metric Definition Collection, and Processor Metrics from OEM to Redfish namespace.
 - Processor FPGA sub object was moved from OEM to Processor namespace.
 - Endpoint: `PciFunctionNumber` and `PciClassCode` are deprecated in favor of `EntityPciId`
- FPGA over-Fabrics support for multiple initiators.

2.2 Deprecation Notices

The following features are considered deprecated in the Intel® RSD v2.5 release and may be dropped in future releases:

PSME GPT NVMe – storage service which uses GUID Partition Table as volume management.

2.3 Limitations

The following list describes all the limitations for Intel® RSD v2.5 release. Described limitations are targeted for future releases:

- The code was verified with Intel® RSD Software Development Vehicle Rack and Arista* switch hardware. For details on Bulldog Creek (BDC) and Arista* firmware versions, contact your Intel representative.
- During initialization, the PSME SPDK NVMe agent will not try to automatically restore NVMe targets which were created via PSME if targets are missing.
- REST API:

Note: Some specific information about chassis hierarchy and identifiers may need to be provided in the configuration files due to HW limitations. which the user must set up according to the *Intel® Rack Scale Design (Intel® RSD) Pooled System Management Engine (PSME) User Guide*, Software v2.5 (refer to [Table 4](#)).

- Network Limitations:

Note: Arista's* startup-configuration file cannot contain any settings for the interfaces with postfix /2, /3 and /4 (for instance, Ethernet 25/2) if the Quad Small Form-factor Pluggable (QSFP) to 4xSFP+ splitter-cable is not plugged in.

Users must remove all settings for such interfaces, perform copy runtime configuration to startup configuration, and restart the switch.

In the *Intel® Rack Scale Design (Intel® RSD) Architecture Specification Software v2.5*, Section 10.1.2 Switch PSME Support for Base Network Services, (refer to [Table 4](#)) the PSME APIs are required to support specific network services, but the following features were not implemented in the Intel reference code:

- Link Aggregation Group (LAG) configuration using several ports



- Access Control List (ACL) configuration
- Media Access Control (MAC) Address - view and configure switch MAC address tables to define rules about which packets are forwarded or discarded
- Telemetry Limitations

The following PSME APIs are defined in *Intel® RSD PSME REST API Specification Software v2.5* (refer to [Table 4](#)), but are not implemented in the Intel reference code:


- Baseboard Ambient Temperature
- Power Supply Unit (PSU) Temperature
- Fan Speed Revolutions Per Minute (RPM)

2.4 Security Recommendations

It is recommended to implement the following isolation, administration, and security procedures to use the Intel® RSD v2.5 reference code in a production environment.

Table 5. Security Recommendations

Issue	Recommendation
Unauthorized log on to Discovery Service, Storage Server or FPGA-oF Target	The Intel® RSD v2.5 provisioning of the Discovery Service Host, Storage Server, and FPGA-oF Target requires admin/root access to the host for placing manager credentials. Confirm the admin account is secured.
Network adversary provisions bogus credentials to PSME	To withstand network attacks, the communication between PODM and the Storage PSME or an FPGA-oF Target PSME must be secured. To provide this security, the communication channel between the PODM and PSME uses Transport Layer Security (TLS).
Compromise discovery/storage service/ FPGA-oF target authentication credential provisioning process	Internet Protocol Security (IPsec) is not used in Intel® RSD v2.5, and this credential will not be provisioned to either the storage client or the FPGA-oF client. This credential will also not be used between the storage client and the Discovery Service or Storage server.
Authorized admin is fooled into installing/updating a compromised image	There must be mechanisms in place to verify that every firmware and software element within the Intel® RSD trust boundary has not changed from the original version delivered by the author. Typically, the firmware/software should be delivered using a cryptographically signed file to ensure the code has not been altered. Refer to Table 4 , <i>Intel® Rack Scale Design (Intel® RSD) Architecture Specification Software v2.5</i> , Platform Security section for more details.
Attacker impersonates a PODM	Access to NVMe-oF target management APIs is only allowed when the PODM and NVMe-oF targets establish a Transport Layer Security (TLS) connection before they can communicate. The same recommendation stands for the FPGA-oF target. Refer to Table 4 , <i>Intel® Rack Scale Design (Intel® RSD) Architecture Specification Software v2.5</i> , Ethernet Pooled Storage - Security section, or more details.
Attacker eavesdrops on communication between client and Storage Server	Protection is for the datacenter to allow only trusted entities to access the storage access network.
Attacker eavesdrops on communication between FPGA-oF client and server	Protection is for the datacenter to allow only trusted entities to access the FPGA access network. Establishing a secure channel between the client and the Storage Server, for example, an IPsec tunnel is recommended as the FPGA over RoCEv2 protocol is not secure.
The attacker modifies data flowing between client and storage server	Protection for the datacenter to only allow trusted entities to access the storage access network.
The attacker modifies data flowing between FPGA-oF client and server	Protection is for the datacenter to allow only trusted entities to access the FPGA access network.
The attacker obtains a drive that has been discarded and accesses data in that drive by mounting it on its system	The default policy in the PODM shall be to erase the drive during decomposition where the PSME must securely erase the drive. As always, ensure security measures are enacted to safeguard the security of your physical drives during and after use.



Issue	Recommendation
The attacker obtains an FPGA that has been discarded and accesses the Acceleration Function in that FPGA	The default policy in the PODM shall be to erase the FPGA during decomposition where the PSME must remove the user's Acceleration Function from the FPGA, for example by reprogramming the device with a default function.
Attacker gains access to a drive previously assigned to a different user with the old user's data still in it.	The default policy in the PODM shall be to erase the drive during decomposition when the PSME must securely erase the drive.
Attacker gains access to an FPGA previously assigned to a different user with the old user's Acceleration Function still in it.	The default policy in the PODM shall be to erase the FPGA during decomposition when the PSME must remove the user's Acceleration Function from the FPGA.

§



3.0 Known Issues

This chapter presents known issues found during the prior testing of Intel® RSD PSME software v2.5.

Table 6. Status Definitions

Status	Description
Under investigation	The sighting is being investigated.
Root cause identified	The root cause of the defect is identified.
Workaround available	A temporary solution to the defect is provided until the bug is fixed.
As designed	The issue reported is not a defect, and the behavior will not be modified.
Closed no repro	The situation is not observed anymore, and no further investigation is scheduled.
Fixed	The defect has been fixed.

[Table 7](#) is a detailed description of all known issues. Each issue provides a problem statement, implication, workaround, note, and status for each.

Table 7. Known Issues

Issue	Description
HSD-ES 1807484993	Status 503 instead of 201 when creating NVMe Volume
Problem	Creating big volume takes too much time and service return timeout (503 HTTP response code).
Implication	Volume will be created after few seconds.
Note	Volume will be visible on the API after polling interval specified in configuration file.
Workaround	N/A
Status	Under investigation

Issue	Description
HSD-ES 1807485208	Sporadically POD returns 404 on different resources after restart PSME
Problem	After restart PSME build-up resources structure from scratch. It should not expose resources before discovery process is finished and all resources are enumerated.
Implication	PODM returns 404 on resources that are already in PODM database but have not been discovered by PSME yet.
Note	N/A
Workaround	N/A
Status	Root cause identified.



Issue	Description
HSD-ES 2201225313	Too many tasks will slow down performance of PSME
Problem	Too many tasks can slow down Assemble and Delete Node operations.
Implication	If the user sends Power Off and then Power On command before the first command is applied, the platform will shut down but will not start.
Note	N/A
Workaround	Control number of tasks by performing DELETE action on unnecessary task on Intel® RSD PSME REST API Specification, Table 4.
Status	Root cause identified.

Issue	Description
HSD-ES 1807485184	Node is in failed state after assembling from snapshot
Problem	The PSME creates a lock on the Baseboard Management Controller (BMC) for exclusive operations on the MDR region, and then performs the operations (writing an Internet Small Computer Systems Interfaces (iSCSI) boot attempt in the region). This lock has a timeout. In the defect, the thread performing writes appears to have been pre-empted, because the lock timed-out before the thread could finish writing.
Implication	Setting iSCSI OOB Boot attributes on PSME sporadically fails. This causes the node to enter “Failed” state. The node cannot be recovered from the state and must be deleted.
Note	If it occurs, this issue causes a failure of Node assembly.
Workaround	Delete the failed Node, Allocate, and Assemble it again until successful.
Status	Root cause identified.


Issue	Description
HSD-ES 1807485222	No information about malformed configuration file after setting incorrect value
Problem	PSME will not log if values in configuration files exceed maximum value.
Implication	PSME will use the wrong value.
Note	N/A
Workaround	N/A
Status	Root cause identified



Issue	Description
HSD-ES 1807484921	Aggregated power measurements value from RMM is bigger than aggregated power measurements value from PSME
Problem	Manually reading consumed power measurement provides different values than the aggregated value provided by the Chassis agent to the RMM
Implication	RMM returns invalid consumed power measurement
Note	N/A
Workaround	Use manual readings from BMCs to obtain correct values
Status	Under investigation

Issue	Description
HSD-ES 1807485048	After sending "LoadFactoryDefaults" no changes on RMM REST API immediately after receiving 204.
Problem	Due to incorrect synchronization, changes caused by the action appear a while after it's completed successfully
Implication	There are no changes on the API immediately after "LoadFactoryDefaults" action
Note	N/A
Workaround	Wait until PSME REST server refreshes its state after polling interval specified in configuration file.
Status	Root cause identified.

Issue	Description
HSD-ES 1807484777	RMM Rest API is reading a different properties for this same type of chassis
Problem	Sporadically, RMM is not able to read Serial Number of a drawer
Implication	Without the Serial Number, RMM is not able to assign a stable REST identifier to the drawer.
Note	N/A
Workaround	N/A
Status	Under investigation



Issue	Description
HSD-ES 1807484605	GPT Managers Ethernet interface doesn't have MACAddress
Problem	Reading the MAC Address of the interface used by the PSME Rest server fails on GPT NVMe storage hosts.
Implication	Upper tier services like PODM are unable to match the interface with ports on an Ethernet Switch.
Note	N/A
Workaround	N/A
Status	Under investigation

Issue	Description
HSD-ES 1807485029	Event Subscription on PSME isn't compliance with Specification
Problem	The PSME rest server does not allow for specifying "SubscriptionType" property when creating a new Event Destination
Implication	A specification compliant request is rejected by the service
Note	N/A
Workaround	Send POST request without the property to create a new Event Destination
Status	Root cause identified

Issue	Description
HSD-ES 1807288860	There are differences between PCIDevices-> Function - before and after deploy build
Problem	Sporadically, installing a newer version of the PSME Compute software causes unexpected changes in the list of discovered PCIDevices and PCIeFunctions
Implication	PCIeFunction resources are assigned wrong identifiers on the REST API
Note	N/A
Workaround	N/A
Status	Under investigation



Issue	Description
HSD-ES 1807484565	500 Whitelabel Error Page after sending "GracefullRestart" on drawer
Problem	The PSME REST API for Chassis incorrectly displays the list of available actions. This causes an error if the drawer reset action is attempted through the PODM software.
Implication	The user is not able to reset a drawer using PODM REST API
Note	N/A
Workaround	To reset a drawer, send the request directly to the RMM REST API
Status	Root cause identified



4.0 Fixed Issues

This chapter presents Fixed issues for Intel® RSD PSME software v2.5.

Note: The Fixed issues in this revision were resolved in prior Intel® RSD PSME Software testing.

Table 8. Fixed Issues

Issue	Description
HSD-ES 1807484996	Status 500 instead of 201 when creating NVMe Volume
Problem	Creating new volume in selected storage pool immediately after deleting existing one will cause internal server error.
Implication	Volume will not be created.
Status	Fixed

Issue	Description
HSD-ES 1807485151	Bad links to Storage Service under System's Hosted Storage Services collection
Problem	List of links under /redfish/v1/Systems/1/StorageServices collection point to non-existing resources.
Implication	Navigation using these links will fail with 404 HTTP response code.
Workaround	Use link to hosted system in storage service resource instead.
Status	Fixed

Issue	Description
HSD-ES 1807485057	Storage Pool "ConsumedBytes" property value has not changed after patch capacity bytes of volume
Problem	Storage pool shows incorrect capacity information.
Implication	Next volume creation request can fail due to incorrect capacity availability.
Status	Fixed

Issue	Description
HSD-ES 1807485089	Status 503 during DELETE SPDK volume equal to the size of the entire storage pool
Problem	Delete volume operation takes too much time and service returns timeout (503 HTTP response code).
Implication	Volume will exist on the REST API but not exist in SPDK.
Status	Root cause identified




Issue	Description
HSD-ES 1807485225	PSME FPGAoF incorrectly exposes UUID of the AFU that is programmed onto the FPGA
Problem	PSME FPGAoF incorrectly reads property of the FPGA
Implication	The value on the REST API is inconsistent with reality
Status	Fixed

Issue	Description
HSD-ES 1807485229	Parallel FPGA operations from many initiators may not be handled properly in FPGA target Service
Problem	Parallel operations on FPGAs are not supported
Implication	Parallel workloads may fail
Status	Fixed

Issue	Description
HSD-ES 1807484821	Delete zone task ends before changes are made
Problem	PSME PNC returns information about zone deletion task finished while the changes on PSME REST API appear after some time
Implication	The task result cannot be reliably monitored
Status	Fixed

Issue	Description
HSD-ES 1807484963	Patch on zone returns 200 OK but effect after more time.
Problem	PSME PNC returns information about zone edition task finished while the changes on PSME REST API appear after some time
Implication	The task result cannot be reliably monitored
Status	Fixed



Issue	Description
HSD-ES 1807485046	When the I2C communication on the PCIe switch is corrupted, PSME PNC presents devices of unknown type as processors
Problem	Physical devices are falsely reported due to HW issue
Implication	The value on the REST API is inconsistent with reality
Status	Fixed

Issue	Description
HSD-ES 1807484796	Can't SSH to Node after PATCH and reset
Problem	Very sporadically, it takes over an hour for an Assembled node to be ready for use.
Implication	The assembled node is unworkable.
Status	Root cause identified – environmental issue

Issue	Description
HSD-ES 1807485135	Patch System PerformanceConfiguration returns Exception
Problem	PSME accepts a request with incorrect configuration identifier and attempts to process it
Implication	Processing PerformanceConfiguration with incorrect configuration identifier returns an exception.
Status	Fixed

Issue	Description
HSD-ES 1807485130	All values in Memory/AEP_DIMM/Metrics are "null"
Problem	If the firmware on the Optane* DC Persistent Memory modules is too old, PSME Compute software is unable to read telemetry data from the modules
Implication	The Metric resource contains null values for all readings.
Status	Fixed