



Realm Management Monitor specification

Document number	DEN0137
Document quality	ALP
Document version	1.1-alp9
Document confidentiality	Non-confidential
Document build information	6ae3a6f4 doctool 0.56.1

Copyright © 2022-2024 Arm Limited or its affiliates. All rights reserved.

DRAFT

Realm Management Monitor specification (Draft)

This document is a DRAFT of the Realm Management Monitor specification.
Please send feedback to the Arm CCA architecture team.

Realm Management Monitor specification

DRAFT

Release information

1.1-alp9 (30-09-2024)

New features

None

Clarifications

- Clarify rules regarding execution of RSI_HOST_CALL by Pn
- State the device functionality required by the RMM (FENIMORE-837)
- Clarify transfer of GIC ownership on Plane exit

Defects

- RMI_REC_{CREATE,DESTROY}: add realm.num_recs to footprint (FENIMORE-859)
- RMI_RTT_SET_S2AP: don't return RMI_ERROR_RTT{_AUX} if S2AP overlay index already matches requested value (FENIMORE-869)
- Remove RmiPdevFlags::spdm (FENIMORE-872)
- Simplify PDEV attributes and creation flags (FENIMORE-871)
- RMI_VDEV_ABORT: transition to VDEV_ERROR state (FENIMORE-856)
- RMI_RTT_AUX_{CREATE/DESTROY/FOLD}: return RMI_ERROR_RTT_AUX, not RMI_ERROR_RTT (FENIMORE-863)
- Update interrupt and timer rules to accommodate Pn being GIC owner (FENIMORE-866)
- Fix defects in HIPAS / RIPAS transition diagrams (FENIMORE-857)
 - The only permitted transition to RIPAS=IO is from HIPAS=ASSIGNED_DEV_*, RIPAS=EMPTY.
- RMI_RTT_AUX_UNMAP_{PROTECTED,UNPROTECTED}: return RMI_ERROR_RTT_AUX, not RMI_ERROR_RTT (FENIMORE-873)

Relaxations

None

1.0-rel0 (10-09-2024)

Clarifications

- RMI_RTT_READ_ENTRY: add ripas_prot success condition
- Clarify rules regarding Realm EL1 timer state
- Correct wording in “Initialize memory of New Realm” flow
- RecAuxCount return value is not greater than 16, and constant for a Realm (FENIMORE-796)
- Clarify purpose of CCA platform hash algorithm ID claim (FENIMORE-811)
- Clarify behaviour of RMI_REC_ENTER if RMI_EMULATED_MMIO flag is set following a REC exit not due to emulatable Data Abort
- RMI_RTT_READ_ENTRY: simplify expression of ripas_unprot pre-condition (FENIMORE-847)

Defects

- Correct typo in “REC entry” section [L_{LFYDV}]
- Add rule regarding Realm execution of data cache invalidate by set / way (FENIMORE-734)
- Remove SH from the set of Host-controlled Unprotected RTT attributes (FENIMORE-736)
- If LPA2 is enabled, ensure that PA written to RTTE is less than 2⁴⁸ (FENIMORE-752)
- RMI_RTT_SET_RIPAS: if base address is not aligned with entry at which RTT walk terminates, only fail if RIPAS of that entry does not match the requested value (FENIMORE-765)
- RMI_DATA_DESTROY: if address is mapped as block, level can be either 1 or 2 (FENIMORE-775)
- RMI_RTT_MAP_UNPROTECTED: remove reference to non-existent output value “nl” (FENIMORE-776)
- Make number of GICv3 List Register values discoverable (FENIMORE-779)
- RMI_REC_ENTER: if RMI_INJECT_SEA is set then RMI_EMULATED_MMIO is ignored (FENIMORE-782)
- Impose IMPLEMENTATION_DEFINED limit on maximum number of RECs per Realm (FENIMORE-800)
- Allow Realm to query RIPAS of an IPA range (FENIMORE-802)

- Introduce RIPAS DEV value (FENIMORE-802)
- Add RPV to RsiRealmConfig (FENIMORE-810)
- Expand RmiFeatureRegister0::{NUM_BPS, NUM_WPS} to support up to 64 counters (FENIMORE-759)
- Attestation token: change profile value to be a versioned tag (FENIMORE-809)
- RSI_ATTESTATION_TOKEN_CONTINUE: add RSI_ERROR_UNKNOWN failure condition (FENIMORE-832)
- RmiFeatureRegister0::GICV3_NUM_LRS: report number of available LRs, minus one (FENIMORE-845)
- Simplify definition of NUM_BPS, NUM_WPS fields (FENIMORE-846)
- RMI_RTT_READ_ENTRY: ripas_unprot failure condition: change && to || (FENIMORE-861)
- RMI_RTT_INIT_RIPAS: correct inconsistency between text and command definition (FENIMORE-864)
- Fix defects in HIPAS / RIPAS transition diagrams (FENIMORE-857)
 - For transitions due to execution of RMI_RTT_DESTROY, remove from the diagram and describe in text.

Relaxations

- RMI_RTT_{INIT,SET}_RIPAS: relax “top_rtt_align” failure condition
 - The previous condition caused the command to fail if the “top” address was misaligned
 - This is replaced with “no_progress”, which only fails if the command does not modify any RTT entries

1.1-alp8 (29-08-2024)

New features

- Allow P0 to transfer GIC ownership to Pn on Plane entry (FENIMORE-855)

Clarifications

- REC exit from Pn: amend “HIPAS DESTROYED” to “RIPAS destroyed” (FENIMORE-825)
- Rename RMI_PDEV_SET_KEY -> RMI_PDEV_SET_PUBKEY (FENIMORE-824)
- Amend description of AUX_DESTROYED state (FENIMORE-822)
- Rename RsiIoCoherent -> RsiDevMemCoherent (FENIMORE-829)
- Clarify behaviour of RMI_REC_ENTER if RMI_EMULATED_MMIO flag is set following a REC exit not due to emulatable Data Abort
- RMI_RTT_READ_ENTRY: simplify expression of ripas_unprot pre-condition (FENIMORE-847)
- Reduce usage of term “IO” in DA-related parts of the spec (FENIMORE-854)
 - RmiIoAction -> RmiVdevAction
 - RmiIoData -> RmiDevCommData
 - RmiIoDelegateFlags -> RmiDevDelegateFlags
 - RmiIoEnter -> RmiDevCommEnter
 - RmiIoEnterStatus -> RmiDevCommStatus
 - RmiIoExit -> RmiDevCommExit
 - RmiIoExit -> RmiDevCommExitFlags
 - RmiIoRequestType -> RmiDevCommProtocol
 - RmiIoShared -> RmiDevMemShared
 - RmmIoState -> RmmDevCommState
 - RmmIoShared -> RmmDevMemShared
 - REC exit due to IO -> REC exit due to device communication
 - RmiRecExit::io_vdev -> RmiRecExit::vdev
 - RmiRecExit::io_action -> RmiRecExit::vdev_action
 - RIPAS IO -> RIPAS DEV
 - Granule states IO_{PRIVATE,SHARED}, IO_DELEGATED_{PRIVATE,SHARED}, IO_UNDELEGATED -> DEV_{PRIVATE,SHARED}, DEV_DELEGATED_{PRIVATE,SHARED}, DEV_UNDELEGATED
 - RMI_GRANULE_IO_{DELEGATE,UNDELEGATE} -> RMI_GRANULE_DEV_{DELEGATE,UNDELEGATE}
 - RMI_IO_{CREATE,DESTROY} -> RMI_DEV_MEM_{MAP,UNMAP}
 - RSI_RDEV_VALIDATE_IO -> RSI_RDEV_VALIDATE_MAPPING
- Replace “Plane exit due to IRQ” with “Plane exit due to synchronous exception”
- Update ECDSA reference (FENIMORE-831)

Defects

- RMI_RTT_AUX_DESTROYED: on success, change RTTE state to AUX_DESTROYED (FENIMORE-827)
- RSI_PLANE_ENTER: fail if plane_idx is zero (FENIMORE-826)
- Add RsiRealmConfig::gicv3_vtr (FENIMORE-807)
- On REC entry, go to P0 if Pn has a maintenance interrupt pending (FENIMORE-830)
- RSI_ATTESTATION_TOKEN_CONTINUE: add RSI_ERROR_UNKNOWN failure condition (FENIMORE-832)
- Remove RMI_RTT_AUX_READ_ENTRY (FENIMORE-823)
- RMI_RTT{ _AUX }DESTROY sets S2AP overlay index to 0 (FENIMORE-820)
- RMI_RTT_AUX_MAP_{ PROTECTED, UNPROTECTED }: correct calculation of output address (FENIMORE-828)
- RmiFeatureRegister0::GICV3_NUM_LRS: report number of available LRs, minus one (FENIMORE-845)
- Simplify definition of NUM_BPS, NUM_WPS fields (FENIMORE-846)
- Add RMI_{ PDEV, VDEV }_AUX_COUNT; remove RmiFeatureRegister0::{ PDEV, VDEV }_NUM_AUX (FENIMORE-833)
- Allow RMI_FEATURES to report that auxiliary RTTs are not supported (FENIMORE-842)
- Replace IDE_{ LINK, SEL, LINK_SEL } with “device protection” enum (FENIMORE-848)
- Enforce that vdev.tdi_id is unique within a segment (FENIMORE-849)
- Enforce that vdev.tdi_id is within PDEV RID range (FENIMORE-850)
- Enforce 1:1 mapping from vdev_id to VDEV within a Realm (FENIMORE-851)
- Add RMI_VDEV arguments required for locking discipline (FENIMORE-852)
- Cacheability attributes for device memory are defined by stage 1 (FENIMORE-858)

1.1-alp7 (09-07-2024)

New features

- Memory Encryption Contexts (MEC) (FENIMORE-788)
- Support for coherent devices and platform devices (FENIMORE-814)
 - Introduce concepts of platform-attested and independently-attested devices
 - Describe differences and commonalities in communication flows for devices with / without SPD
 - RMI_PDEV_CREATE: replace device class with flags (SPDM, IDE)
 - PDEV: increase maximum number of auxiliary granules to 32
 - VDEV: introduce auxiliary granules
 - RMI_IO_CREATE
 - * Remove flags input value
 - * Permit MemAttr to specify either a cacheable or a non-cacheable mapping
 - Rename RSI_RDEV_CONFIG -> RSI_RDEV_GET_INFO
 - RSI_RDEV_GET_INFO: permit execution in any RDEV state
 - Introduce REC exit due to VDEV request, RMI_VDEV_COMPLETE
 - RSI_RDEV_VALIDATE_IO: add “coherent” flag

Clarifications

- Fix typo in “HIPAS change while Realm state is ACTIVE” diagram (FENIMORE-794)
- Correct command signatures in “Realm validation of MMIO” flow (FENIMORE-795)
- Rename Firmware Component Measurement Log (FCML) -> Firmware Activity Log (FAL) (FENIMORE-798)
- Correct wording in “Initialize memory of New Realm” flow
- RecAuxCount return value is not greater than 16, and constant for a Realm (FENIMORE-796)
- Clarify behavior of Pn instruction fetch from Unprotected IPA
- Fix pseudocode for S2AP change (FENIMORE-805)
- Clarify purpose of CCA platform hash algorithm ID claim (FENIMORE-811)
- Clarify that RSI_RDEV_VALIDATE_IO does not permit change from RIPAS DESTROYED (FENIMORE-785)

Defects

- Rename RSI_RDEV_GET_DIGESTS -> RSI_RDEV_CONFIG (FENIMORE-751)
- Allow Host to select PDEV hash algorithm (FENIMORE-751)
- Impose IMPLEMENTATION_DEFINED limit on maximum number of RECs per Realm (FENIMORE-800)
- Remove RmiIoRequestType::RMI_DISCOVERY (FENIMORE-772)
- RMI_RTT_SET_S2AP: add rtt_tree output value (FENIMORE-806)
- Allow Realm to query RIPAS of an IPA range (FENIMORE-802)

- Allow RSI_HOST_CALL execution by Pn to directly exit to the Host (FENIMORE-804)
- On REC exit from Pn, timer state reported to Host is the earliest of P0 and Pn timers (FENIMORE-801)
- Move Pn EL1 sysreg access from RSI_PLANE_ENTER to RSI_PLANE_REG_{READ,WRITE} (FENIMORE-773)
- Add RPV to RsiRealmConfig (FENIMORE-810)
- Expand RmiFeatureRegister0::{NUM_BPS, NUM_WPS} to support up to 64 counters (FENIMORE-759)
- Add S2AP to criteria for RTT homogeneity, when HIPAS=ASSIGNED (FENIMORE-818)

Relaxations

None

1.1-alp6 (31-05-2024)

New features

- Introduce Live Firmware Activation policy (FENIMORE-780)
- Add flag in platform token, indicating whether a software component may be live-activated (FENIMORE-781)
- Make number of GICv3 List Register values discoverable (FENIMORE-779)
- Support for countersignatures on software components (FENIMORE-793)

Clarifications

- Fix typo in “HIPAS change while Realm state is NEW” diagram
- Add RTT_AUX to Granule state diagram (FENIMORE-784)
- Planes: clarifications
 - Exception model
 - * Amend initial state in “Between a Plane entry and a Plane exit, a REC exit and REC entry may occur”
 - * Remove statement that “On Plane entry, all RsiPlaneEnter fields are ignored unless specified otherwise”
 - * Regarding “An exception due to any of the following in Pn cause a REC exit to the Host ...”
 - Introduce “REC exit from Pn” heading to provide appropriate context
 - Separate synchronous and asynchronous exceptions
 - * Add rule that REC entry with a Pending virtual interrupt causes return to P0
 - * Reword info statements regarding action taken by P0 in response to a Plane exit, to remove suggestion that this is an exhaustive list
 - * Remove statement that “all other RsiPlaneExit fields are zero” following Plane exit due to Synchronous Exception
 - Memory management
 - * Reword overview to introduce “S2AP base index”
 - * Remove (“The number of S2AP overlay indices is 16”) as this is specified by the Arm ARM
 - * Reword description of execute permission in Unprotected IPA space
 - * RmiRealmFlags: remove comment that rtt_tree_pp is ignored if the Realm has a single Plane
 - * Correct implementation note to say that S2AP cookie should always be zero if the Realm is configured to use a shared RTT tree
 - * Remove incorrect statement that RSI_MEM_SET_PERM_INDEX does not return a “top” address
- DA
 - Fix typo in RDEV state diagram
 - Make explicit the relationship between RDEV state and SMMU enablement (FENIMORE-787)
 - Remove SPDM_CHALLENGE step from PDEV setup flow (FENIMORE-761)

Defects

- RSI_MEM_SET_PERM_VALUE: fail if Plane index is zero (FENIMORE-778)
- RMI_REC_ENTER: if RMI_INJECT_SEA is set then RMI_EMULATED_MMIO is ignored (FENIMORE-782)
- RMI_PDEV_COMMUNICATE: add missing success conditions (FENIMORE-786)
- Separate Realm flags into measured and unmeasured sets (FENIMORE-792)
 - Rename RmiRealmFlags -> RmiRealmFlags0
 - Move rtt_tree_pp from RmiRealmFlags0 to RmiRealmFlags1

Relaxations

None

1.1-alp5 (30-04-2024)

New features

None

Clarifications

None

Defects

- Planes: add rules regarding EL1 timers
- RMI_VDEV_{COMMUNICATE,CREATE}: add pdev.num_vdevs to footprint (FENIMORE-774)
- RMI_DATA_DESTROY: if address is mapped as block, level can be either 1 or 2 (FENIMORE-775)
- RMI_RTT_MAP_UNPROTECTED: remove reference to non-existent output value “nl” (FENIMORE-776)
- Reduce number of available memory permission overlay indices to 15 (FENIMORE-777)

Relaxations

None

1.1-alp4 (03-04-2024)

New features

None

Clarifications

- Add rationale for permitting RSI_RDEV_GET_MEASUREMENTS while the RDEV is in RDEV_NEW state (FENIMORE-745)
- Add rationale for Host storage of device attestation evidence (FENIMORE-741)

Defects

- RmmRdev: add linkage to corresponding RmmVdev (FENIMORE-746)
- RMI_PDEV_SET_PUBKEY: add failure condition for invalid key length (FENIMORE-740)
- RMI_RTT_SET_RIPAS: if base address is not aligned with entry at which RTT walk terminates, only fail if RIPAS of that entry does not match the requested value (FENIMORE-765)
- Specify that maximum number of auxiliary Planes is either 0 or 3 (FENIMORE-769)
- Specify S2AP change flow (FENIMORE-770)
 - Follow a pattern similar to the RIPAS change flow, including allowing the Host to reject an S2AP change request.
- Remove misleading transition to PDEV_ERROR in response to Host reporting RMI_IO_ERROR (FENIMORE-764)
- Rename RMI_SECURE_SPDM to RMI_SECURE_CMA_SPDM (FENIMORE-767)
- RMI_PDEV_ABORT: amend failure / success conditions to match state transition diagram (FENIMORE-768)
- RMI_PDEV_SET_PUBKEY: add success condition which sets io_state to IO_PENDING (FENIMORE-771)

Relaxations

None

1.1-alp3 (31-01-2024)

Clarifications

- Update “RTT walk” section to reflect addition of auxiliary RTTs (FENIMORE-737)
- RmiPdevEventIsValid(): remove pdev argument (FENIMORE-747)
- RMI_RTT_DESTROY: add ordering “rtte_state < rtt_live” (FENIMORE-748)
- Planes: fix typo in description of memory permission overlay index lock bit (FENIMORE-753)
- VDEV teardown flow: remove RMI_IO_DESTROY calls (FENIMORE-756)
- Clarify meaning of “secure SPDM” (FENIMORE-743)
- Address review feedback on Planes chapter
 - Replace use of undefined term “P0 memory” (in reference to RsiPlaneRun object) with “Realm memory”.
 - Clarify that the access permissions which can vary between Planes are at stage 2.

- Clarify the subset of REC exits in which an RTT tree index is returned to the Host.
- Use “rec_exit” and “plane_exit” to avoid ambiguity.
- Append to list of Pn actions which result in REC exit to the Host.
- Each Plane has a unique VMID, regardless of whether there exists an RTT tree per Plane.
- Limit the statement “if an IPA is auxiliary-live then the corresponding entry in the primary RTT is live” to apply only to Protected IPA space.
- Correct statement regarding variance of stage 2 permissions among Planes for an Unprotected IPA.

Defects

- If LPA2 is enabled, ensure that PA written to RTTE is less than 2^{48} (FENIMORE-752)
- On REC exit due to Data Abort / Instruction Abort, provide the index of the RTT tree (FENIMORE-749)
- RSI_MEM_SET_PERM_INDEX: add success condition which sets lock bit (FENIMORE-754)
- On RIPAS change to RAM, reset memory permission overlay index to 0 (FENIMORE-755)
- Introduce RMI_ERROR_RTT_AUX (FENIMORE-757)
- RSI_PLANE_ENTER: add plane_idx; move run_ptr from output to input (FENIMORE-758)
- RSI_PLANE_ENTER: add failure and success conditions
- Add RMI_PDEV_IDE_RESET command (FENIMORE-726)
- Reassign FIDs for commands added in RMM 1.1, to avoid overlap between RMI and RSI

Relaxations

None

1.1-alp2 (11-12-2023)

New features

None

Defects

- RSI_RDEV_VALIDATE_IO: add “private / shared” flag (FENIMORE-732)
- Add rule regarding Realm execution of data cache invalidate by set / way (FENIMORE-734)

Relaxations

None

Defects

- RMI_RTT_READ_ENTRY: add success condition for HIPAS=ASSIGNED_IO_* (FENIMORE-733)
- Remove SH from the set of Host-controlled Unprotected RTT attributes (FENIMORE-736)

1.1-alp1 (06-11-2023)

New features

- Planes (FENIMORE-731)
 - RmiFeatureRegister0: add RTT_TREE_SINGLE, MAX_NUM_AUX_PLANES fields
 - RmiRealmFlags: add RTT_TREE_PP field
 - RmiRealmParams: add members
 - * num_aux_planes
 - * rtt_tree_pp
 - * aux_vmid
 - * aux_rtt_base
 - RmmRealm
 - * Add members
 - rtt_tree_pp
 - num_aux_planes
 - * Change type of members to “array of per-Plane values”
 - vmid
 - rtt_base

- Functions
 - * Add functions to compare arrays of per-Plane values
 - RealmRttBaseEqual()
 - RealmVmidEqual()
 - * Modify signatures to accept array of per-Plane values
 - VmidIsFree() -> VmidsAreFree()
 - VmidIsValid() -> VmidsAreValid()
 - * Modify signature of RttWalk() to take an RTT tree index
- RMI_REALM_CREATE
 - * Add success conditions
 - rtt_tree_pp
 - num_aux_planes
- Add RMI_RTT_AUX_* commands for manipulation of auxiliary RTTs
- RMI_{DATA_CREATE,RTT_SET_RIPAS}: add failure condition if IPA is live in an auxiliary RTT
- Add RSI_PLANE_ENTER and describe Plane exception model
- Add RSI_MEM_* commands and describe memory access permission management
- RsiRealmConfig: add num_aux_planes
- RsiFeatureRegister0: add MRO field

Clarifications

- Realm device assignment
 - Replace references to DOE, ECAM, IDE with reference to PCIe 6.0
 - Amend “Realm device assignment overview” to correctly describe flow for retrieval of device interface report and device measurements
 - Amend “Device requests and responses” to include RSI_RDEV_GET_{INTERFACE_REPORT,MEASUREMENTS}
 - Amend IDE setup flow to include RSI_RDEV_{LOCK,GET_INTERFACE_REPORT,GET_MEASUREMENTS}
 - Clarify that RSI_RDEV_CONTINUE causes a REC exit due to IO, which must be completed by Host execution of RMI_VDEV_COMMUNICATE
 - Amend “Virtual device teardown flow” to show disabling of SMMU on RSI_VDEV_STOP
 - Remove incorrect statement that IPA base is present in the interface report
 - Amend description of RSI_RDEV_STOP: in-flight transactions may still complete, but future device accesses to Realm memory are blocked
 - RMI_PDEV_NOTIFY: add check for validity of event identifier
 - State that VDEV and RDEV are Host and Realm views of the same underlying RMM object
 - Describe relationship between RDEV state and TDI state
- Compress rendering of array struct members
- Consistently postfix name of “before: true” context values with “_pre”
- Replace inline “Realm(rd)” expressions with “realm” context value
- Replace inline “Rec(rec_ptr)” expressions with “rec” context value
- Improve consistency of function names
 - Casting a memory location to an object is suffixed “At”, for example “RecAt”
 - Casting an object to an interface type is suffixed “To”, for example “RttEntryStateToRmi”
 - Rename RttEntryFromDescriptor() to RttDescriptorDecode()
- RMI_RTT_READ_ENTRY: add ripas_prot success condition
- Feature discovery and selection
 - Rename RmiFeatureRegister0 fields
 - * SVE_EN -> SVE
 - * PDEV_AUX -> PDEV_NUM_AUX
 - * PMU_EN -> PMU
 - {RMI,RSI}_FEATURES: add {Rmi,Rsi}FeatureRegisterEncode() functions

- * Formalise the mapping from features supported by the platform to the output value of these commands
- RMI_REALM_CREATE: implement RealmParamsSupported()
- * Formalise the check that input values to this command are compatible with features supported by the platform

Defects

- Realm device assignment
 - RTT folding: for ASSIGNED_IO_*, require that memory attributes (which are Host-controlled) are the same for all entries
 - Permit Host to specify any Device memory attributes for an IO mapping
 - Specify that IO mappings are Outer Shareable
 - Remove RmiIoEnter::req_size, and state that the Host should always provide a 4KB request buffer
 - Add RmiIoExit::req_type, to inform Host via which mailbox to route the device request
 - RSI_RDEV_GET_DIGESTS: do not trigger REC exit due to IO
- Correct typo in “REC entry” section [L_FYDV]

Relaxations

- Realm device assignment
 - RMI_{IO_}_GRANULE_UNDELEGATE: permit new GPT entry to be any value except GPT_REALM
- RMI_RTT_{INIT,SET}_RIPAS: relax “top_rtt_align” failure condition
 - The previous condition caused the command to fail if the “top” address was misaligned
 - This is replaced with “no_progress”, which only fails if the command does not modify any RTT entries

1.1-alp0 (05-10-2023)

New features

- Realm device assignment (FENIMORE-722)
 - New Granule states: PDEV{AUX}, VDEV, IO{UNDELEGATED,DELEGATED_PRIVATE,DELEGATED_SHARED}, IO_{PRIVATE,SHARED}
 - * RMI commands for (un)delegation of IO physical memory
 - New RMM objects: PDEV, VDEV
 - * RMI commands for creation, destruction and lifecycle management
 - Realm-facing abstraction for an assigned device: RDEV
 - RMM-device communication flow (protected by SPDm)
 - * Mediated through “REC exit due to IO”, RMI_{PDEV,VDEV}_COMMUNICATE commands and buffers in NS memory
 - * When triggered by a Realm action, managed via an “interruptible operation” programming interface, similar to the existing one for RSI_ATTESTATION_TOKEN commands
 - Host caching of device attestation evidence (certificate, measurements, interface report), with integrity ensured via RMM-held digests
 - New HIPAS values (ASSIGNED_IO_PRIVATE, ASSIGNED_IO_SHARED) and RIPAS value (IO)
 - * Extension of rules regarding Realm access to Protected IPA space
 - * Extension of rules regarding RTT (un)folding
 - Extension of RIPAS change flow for the purpose of validating mappings against the device interface report, and transition to RIPAS IO

Clarifications

None

Defects

None

Relaxations

None

1.0-eac5 (05-10-2023)

Clarifications

- Fix attestation token flows (FENIMORE-718)
- Clarify behavior on Host rejection of a RIPAS change request (FENIMORE-719)
- Replace Granule::pas attribute with Granule::gpt
 - PAS is an attribute of a memory access, not of a Granule.

Defects

- {RMI,RSI}_VERSION: (FENIMORE-724)
 - Clarify rules regarding returned interface version, and provide examples
 - Remove rule that if the return code is SUCCESS, subsequent calls to the interface adhere to the behavior corresponding with the returned interface version
- Specify that SMCCC registers not specified as command input / output values are SBZ and MBZ respectively (FENIMORE-724)
- RSI_ATTESTATION_TOKEN_INIT: return upper bound on token size (FENIMORE-720)
- RMI_DATA_CREATE: move RIPAS=RAM from being a pre-condition to a post-condition (FENIMORE-721)

Relaxations

None

1.0-eac4 (06-09-2023)

Clarifications

- Exclude GIC, timer and PMU values from “On REC exit ... all other REC exit fields are zero” (FENIMORE-712)
- Amend contradictory statement regarding RTT folding to level 1 (FENIMORE-715) [IQWQSB]

Defects

- RMI_RTT_{INIT,SET}_RIPAS: fix “top” alignment check
 - Ensure that “top” is Granule aligned (FENIMORE-710)
 - Ensure that return code is deterministically specified (FENIMORE-711)
 - Prevent RIPAS change from proceeding beyond the “top” address provided by the Realm (FENIMORE-711)
- {RMI,RSI}_VERSION: add handshake (FENIMORE-708)
 - The caller provides a “requested version”
 - The RMM either returns:
 - * A version which it can implement, that is compatible with the requested version (and a SUCCESS return code)
 - * A version which it implements, that is incompatible with the requested version (and an error code)
 - If the return code is SUCCESS, subsequent calls to the interface adhere to the behavior corresponding with the returned interface version
- Increase width of PsciReturnCode to 64 bits (FENIMORE-709)

Relaxations

- RMI_REALM_CREATE: permit number of PMU counters to be less than number supported by the implementation (FENIMORE-716)
- RMI_REALM_CREATE: permit number of breakpoints or watchpoints to be less than number supported by the implementation (FENIMORE-717)

1.0-eac3 (20-07-2023)

Clarifications

- Clarify which bits of command input / output values should / must be zero (FENIMORE-674)
- Explain distinction between concrete and abstract types (FENIMORE-693)
- Clarify return value from RSI_IPA_STATE_SET when stopping at first DESTROYED entry (FENIMORE-699) [IGXDDX]

Defects

- PSCI_SYSTEM_{OFF,RESET}: change Realm state to SYSTEM_OFF (FENIMORE-694)

- RMI_REC_CREATE: update RIM only if runnable flag is set (FENIMORE-697)
- RMI_REALM_CREATE: fix list of measured parameters (FENIMORE-695)
- Remove members from RmmSystemRegisters (FENIMORE-700)
 - State saved / restored depends on architecture features supported by the platform, so defining this type as an empty placeholder
- Avoid use of reserved ASL v1 keyword “entry” in MRS (FENIMORE-702)
 - RmiRecEntry -> RmiRecEnter
 - RmiRecEntryFlags -> RmiRecEnterFlags
 - RmiRecRun::entry -> RmiRecRun::enter
 - RmmRttWalkResult::entry -> RmmRttWalkResult::rtte
- RSI_IPA_STATE_SET: prohibit RSI_DESTROYED input value (FENIMORE-705)
- RMI_PSCI_COMPLETE: PSCI_CPU_ON: fix copy of context_id to target CPU X0 (FENIMORE-703)
- Allow Host to reject request to change RIPAS to RAM (FENIMORE-661)
- Allow Host to reject PSCI_CPU_ON request via RMI_PSCI_COMPLETE (FENIMORE-706)

Relaxations

- Permit folding of level 2 RTT to create level 1 block mapping (FENIMORE-608)
- Remove restriction that attestation token size must not exceed 4KB (FENIMORE-691)

1.0-eac2 (07-06-2023)

Clarifications

- Remove reference to triggering ERROR_INPUT by setting MBZ bit to 1 (FENIMORE-675)
- Clarify constraints on output values in case of command failure [R_TFZMS] (FENIMORE-676)
- Clarify encoding of RmiRealmParams::sve_sz (FENIMORE-684)
- Clarify set of SMCCC interfaces available to a Realm [R_NPLKX] (FENIMORE-685)

Defects

- Replace PMU fields in RmiRecExit with single bit indicating the PMU overflow status [R_WXTZF] (FENIMORE-679)
- RMI_PSCI_COMPLETE: failure condition should compare against MPIDR, not RD address (FENIMORE-681)
- RMI_REC_CREATE: remove params_valid failure condition (FENIMORE-686)
- RMI_RTT_{INIT,SET}_RIPAS: check alignment of “top” input value (FENIMORE-687)
- Reduce coupling between HIPAS and RIPAS (FENIMORE-680)
 - Replace HIPAS=DESTROYED with RIPAS=DESTROYED
 - Remove RmiRttEntryState::RMI_DESTROYED
 - Change encoding of RmiRttEntryState::RMI_TABLE
 - Add RmiRipas::RMI_DESTROYED
 - Add RsiRipas::RSI_DESTROYED
 - RMI_DATA_CREATE_UNKNOWN: remove pre-condition that RIPAS=RAM
 - RMI_DATA_DESTROY:
 - * In all cases, post-condition now states that HIPAS=UNASSIGNED
 - * If pre-condition was RIPAS=RAM, post-condition states that RIPAS=DESTROYED
 - RMI_RTT_DESTROY:
 - * Remove post-condition that HIPAS=DESTROYED
 - * Add post-condition that state of parent RTTE is UNASSIGNED
 - * Add post-condition that RIPAS=DESTROYED
 - RMI_RTT_SET_IPA_STATE: stop at first DESTROYED entry if “destroyed” flag is set
 - RSI_IPA_STATE_SET: add “destroyed” flag
 - Clarify distinction between “RTT folding” [D_QPXCp] and “RTT destruction” [D_VXRZW]
- RMI_RTT_INIT_RIPAS: success conditions should be bounded by walk_top, not top

Relaxations

- RSI_REALM_CONFIG: provide Realm hash algorithm (FENIMORE-678)

1.0-eac1 (31-03-2023)

Clarifications

- Unused bits of RmiRecEntry::gicv3_hcr are SBZ [I_{SMHXB}] (FENIMORE-666)
- RMI_REC_ENTER: all RMI_ERROR_INPUT failure conditions precede all RMI_ERROR_REC failure conditions (FENIMORE-668)
- Avoid use of raw Xn values in command conditions where possible (FENIMORE-671)
- Clarify definition of REC exit due to (Non-)emulatable Data Abort [D_{CYRMT}, D_{MTZMC}] (FENIMORE-673)

Defects

- RMI_RTT_INIT_RIPAS: take account of “top” IPA value when calculating RIM contribution (FENIMORE-662)
- RttSkipEntriesWithRipas: fix inverted logic (FENIMORE-663)
- RMI_RTT_SET_RIPAS: on success, modify IPA range [base, walk_top] (FENIMORE-669)
- RMI_RTT_{INIT,SET}_RIPAS: remove redundant failure conditions (FENIMORE-670)
- Clarify HIPAS=DESTROYED implies RIPAS=UNDEFINED [R_{JYDRL}] (FENIMORE-672)

Relaxations

- RSI_HOST_CALL: relax alignment requirement from 4KB to 256B

1.0-eac0 (31-01-2023)

Clarifications

None

Defects

- RmiRealmParams: reduce width of integer attributes (FENIMORE-647)
- RSI_IPA_STATE_SET: replace (base, size) with (base, top) (FENIMORE-656)
- RMI_RTT_INIT_RIPAS, RMI_RTT_SET_RIPAS: allow single command to modify multiple RTT entries (FENIMORE-656)

Relaxations

- RMI_RTT_SET_RIPAS: remove “ripas” input value (FENIMORE-659)

1.0-bet2 (16-12-2022)

Clarifications

- Flows: update RMI_REC_ENTRY to take a single ‘run’ input value
- Clarify meaning of “TTD” [I_{YMNSR}] (FENIMORE-641)
- Fix typo in reference to “CCA platform token claim map” [I_{FJKFY}] (FENIMORE-647)
- Fix reference to “RME system architecture spec” (FENIMORE-648)
- Flows: remove stale reference to parameters passed to RMI_DATA_CREATE (FENIMORE-649)
- Improve definition and consistency of usage of the term “REC” (FENIMORE-650)
 - Where referring to the RMM data structure “REC object” is now used
- Clarify description of properties of Realm IPA space [I_{TPGKW}] (FENIMORE-639)
 - Replace “permitted, under control of host” with statements which refer to particular HIPAS values.
 - Add “Protected IPA, HIPAS=DESTROYED” row, thereby removing contradictory statements regarding SEA taken to Realm, previously in “Protected IPA, RIPAS=EMPTY”.
- On assertion of an EL1 timer, the RMM guarantees a *REC exit*, not only a *Realm exit* (FENIMORE-651)
- RMI_RTT_FOLD: preserve RIPAS value if IPA is Protected (FENIMORE-638)

Defects

- Attestation: wrap sub-tokens in byte stream (FENIMORE-643)
- RMI_DATA_DESTROY, RMI_RTT_{DESTROY,FOLD}: return PA of destroyed object (FENIMORE-563)
- RMI_REALM_DESTROY, RMI_REC_DESTROY, RMI_REC_ENTER, RMI_RTT_DESTROY, RMI_RTT_FOLD, RMI_RTT_SET_RIPAS: Remove RMI_ERROR_IN_USE (FENIMORE-588)

- RMI_DATA_CREATE, RMI_DATA_CREATE_UNKNOWN, RMI_REC_CREATE, RMI_RTT_CREATE: pass RD pointer in X1 (FENIMORE-655)
- Replace RmiRealmParams::features_0 with discrete fields (FENIMORE-655)
- RMI_DATA_CREATE(_UNKNOWN): require RIPAS=RAM (FENIMORE-645)
- Apply “must / should be zero” consistently (FENIMORE-619)
 - In command inputs, unused bits are SBZ
 - In command outputs, unused bits are MBZ

Relaxations

- RSI_HOST_CALL: expand set of GPRs to X0-X30 (FENIMORE-607)
 - This enables the RMM to support any calling convention.
- RMI_DATA_DESTROY, RMI_RTT_DESTROY, RMI_RTT_UNMAP_UNPROTECTED: return IPA of next live RTT entry (FENIMORE-563)

1.0-beta1 (31-10-2022)

Clarifications

- Rename HIPAS_VALID_NS -> UNASSIGNED (FENIMORE-631)
- SEA injection is independent of whether Host emulates MMIO (FENIMORE-632)
- In RIPAS change flow, permit Host to apply the change to zero or more pages of the target IPA region (FENIMORE-633)
- Flows: replace HVC with Host call (FENIMORE-611)
- Clarify behavior of VmidIsValid() function (FENIMORE-630)
- Qualify “all other exit fields are zero” statements [R_GTJRP, R_LRCP] (FENIMORE-634)
 - GIC, timer and PMU fields are valid on every REC exit.

Defects

- Change size of RsiHostCall type to 256 bytes (FENIMORE-629)
- Correct the set of ESR_EL2 fields which are returned to the Host on REC exit due to Data abort [R_RYVFL]
 - On all Data Aborts, add FnV.
 - On Emulatable Data Aborts, add SF.
 - On Non-emulatable Data Abort at an Unprotected IPA, add IL.

Relaxations

None

Non-Confidential Proprietary Notice

This document is protected by copyright and other related rights and the use or implementation of the information contained in this document may be protected by one or more patents or pending patent applications. No part of this document may be reproduced in any form by any means without the express prior written permission of Arm Limited (“Arm”). **No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document unless specifically stated.**

Your access to the information in this document is conditional upon your acceptance that you will not use or permit others to use the information for the purposes of determining whether the subject matter of this document infringes any third party patents.

The content of this document is informational only. Any solutions presented herein are subject to changing conditions, information, scope, and data. This document was produced using reasonable efforts based on information available as of the date of issue of this document. The scope of information in this document may exceed that which Arm is required to provide, and such additional information is merely intended to further assist the recipient and does not represent Arm’s view of the scope of its obligations. You acknowledge and agree that you possess the necessary expertise in system security and functional safety and that you shall be solely responsible for compliance with all legal, regulatory, safety and security related requirements concerning your products, notwithstanding any information or support that may be provided by Arm herein. In addition, you are responsible for any applications which are used in conjunction with any Arm technology described in this document, and to minimize risks, adequate design and operating safeguards should be provided for by you.

This document may include technical inaccuracies or typographical errors. THIS DOCUMENT IS PROVIDED “AS IS”. ARM PROVIDES NO REPRESENTATIONS AND NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE DOCUMENT. For the avoidance of doubt, Arm makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of, any patents, copyrights, trade secrets, trademarks, or other rights.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF ARM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Reference by Arm to any third party’s products or services within this document is not an express or implied approval or endorsement of the use thereof.

This document consists solely of commercial items. You shall be responsible for ensuring that any permitted use, duplication, or disclosure of this document complies fully with any relevant export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws. Use of the word “partner” in reference to Arm’s customers is not intended to create or refer to any partnership relationship with any other company. Arm may make changes to this document at any time and without notice.

This document may be translated into other languages for convenience, and you agree that if there is any conflict between the English version of this document and any translation, the terms of the English version of this document shall prevail.

The validity, construction and performance of this notice shall be governed by English Law.

The Arm corporate logo and words marked with ® or ™ are registered trademarks or trademarks of Arm Limited (or its affiliates) in the US and/or elsewhere. Please follow Arm’s trademark usage guidelines at <https://www.arm.com/company/policies/trademarks>. All rights reserved. Other brands and names mentioned in this document may be the trademarks of their respective owners.

Copyright © 2022-2024 Arm Limited or its affiliates. All rights reserved.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

PRE-20349

8 March 2024

DRAFT

Contents

Realm Management Monitor specification

Realm Management Monitor specification	ii
Release information	iii
Non-Confidential Proprietary Notice	xv

Preface

Quality level	xxix
Conventions	xxx
Typographical conventions	xxx
Numbers	xxx
Pseudocode descriptions	xxx
Addresses	xxx
Rules-based writing	xxxi
Content item identifiers	xxxi
Content item rendering	xxxi
Content item classes	xxxi
Additional reading	xxxii
Feedback	xxxiv
Feedback on this book	xxxiv
Open issues	xxxv

Part A Architecture

Chapter A1

Overview

A1.1 Confidential computing	37
A1.2 System software components	38
A1.3 Realm Management Monitor	38

Chapter A2

Concepts

A2.1 Realm	41
A2.1.1 Overview	41
A2.1.2 Realm execution environment	41
A2.1.3 Realm attributes	42
A2.1.4 Realm liveness	44
A2.1.5 Realm lifecycle	44
A2.1.6 Realm parameters	45
A2.1.7 Realm Descriptor	45
A2.2 Granule	47
A2.2.1 Granule attributes	47
A2.2.2 Granule lifecycle	48
A2.2.3 Granule ownership	51
A2.2.4 Granule wiping	52
A2.3 Realm Execution Context	53
A2.3.1 Overview	53
A2.3.2 REC attributes	53
A2.3.3 REC index and MPIDR value	55
A2.3.4 REC lifecycle	55

Chapter A3

Feature discovery and configuration

A3.1	Feature discovery and configuration overview	58
A3.2	Realm hash algorithm	58
A3.3	Realm LPA2 and IPA width	58
A3.4	Realm support for Scalable Vector Extension	59
A3.5	Realm support for self-hosted debug	60
A3.6	Realm support for Performance Monitors Extension	60
A3.7	Realm support for Activity Monitors Extension	60
A3.8	Realm support for Statistical Profiling Extension	60
A3.9	Realm support for Trace Buffer Extension	60
A3.10	Support for Realm device assignment	61
A3.11	Support for auxiliary Planes	61
A3.12	Live Firmware Activation	61
A3.13	GICv3 virtualization	62
A3.14	Support for Realm memory encryption	62

Chapter A4

Realm exception model

A4.1	Realm exception model overview	64
A4.2	REC entry	66
A4.2.1	RmiRecEnter object	66
A4.2.2	General purpose registers restored on REC entry	66
A4.2.3	REC entry following REC exit due to Data Abort	67
A4.3	REC exit	68
A4.3.1	RmiRecExit object	68
A4.3.2	Realm exit reason	69
A4.3.3	General purpose registers saved on REC exit	69
A4.3.4	REC exit due to synchronous exception	70
A4.3.5	REC exit due to IRQ	73
A4.3.6	REC exit due to FIQ	73
A4.3.7	REC exit due to PSCI	73
A4.3.8	REC exit due to RIPAS change pending	74
A4.3.9	REC exit due to Host call	75
A4.3.10	REC exit due to SError	75
A4.3.11	REC exit due to device communication	75
A4.3.12	REC exit due to RTT request	76
A4.3.13	REC exit due to S2AP change pending	76
A4.3.14	REC exit due to VDEV request	77
A4.4	Emulated Data Aborts	78
A4.5	Host call	78

Chapter A5

Realm memory management

A5.1	Realm memory management overview	80
A5.2	Realm view of memory management	80
A5.2.1	Realm IPA space	80
A5.2.2	Realm IPA state	80
A5.2.3	Realm access to a Protected IPA	81
A5.2.4	Changes to RIPAS while Realm state is REALM_NEW	81
A5.2.5	Changes to RIPAS while Realm state is REALM_ACTIVE	82
A5.2.6	Realm access to an Unprotected IPA	83
A5.2.7	Synchronous External Aborts	83
A5.2.8	Realm access outside IPA space	84
A5.2.9	Summary of Realm IPA space properties	85
A5.2.10	Cache maintenance operations	85
A5.3	Host view of memory management	86
A5.3.1	Host IPA state	86
A5.3.2	Changes to HIPAS while Realm state is REALM_NEW	87

	A5.3.3	Changes to HIPAS while Realm state is REALM_ACTIVE	88
	A5.3.4	Summary of changes to HIPAS and RIPAS of a Protected IPA	89
	A5.3.5	Dependency of RMI command execution on RIPAS and HIPAS values	91
	A5.3.6	Changes to HIPAS of an Unprotected IPA	92
A5.4		RIPAS change	93
A5.5		Realm Translation Table	95
	A5.5.1	RTT overview	95
	A5.5.2	RTT structure and configuration	95
	A5.5.3	RTT starting level	95
	A5.5.4	RTT entry	96
	A5.5.5	RTT reading	97
	A5.5.6	RTT folding	97
	A5.5.7	RTT unfolding	98
	A5.5.8	RTTE liveness and RTT liveness	99
	A5.5.9	RTT destruction	99
	A5.5.10	RTT walk	100
	A5.5.11	RTT entry attributes	101
Chapter A6		Realm interrupts and timers	
	A6.1	Realm interrupts	105
	A6.2	Realm timers	107
Chapter A7		Realm measurement and attestation	
	A7.1	Realm measurements	110
	A7.1.1	Realm Initial Measurement	110
	A7.1.2	Realm Extensible Measurement	111
	A7.2	Realm attestation	112
	A7.2.1	Attestation token	112
	A7.2.2	Attestation token generation	112
	A7.2.3	Attestation token format	114
Chapter A8		Realm debug and performance monitoring	
	A8.1	Realm PMU	135
Chapter A9		Realm device assignment	
	A9.1	Realm device assignment overview	137
	A9.1.1	Assignment of an independently-attested device	137
	A9.1.2	Assignment of a platform-attested device	139
	A9.2	Communication between RMM and a device	142
	A9.2.1	Device requests and responses	142
	A9.2.2	Mapping from virtual device ID to VDEV object	143
	A9.2.3	Host-side device communication flow	145
	A9.2.4	Host caching of device attestation evidence	146
	A9.2.5	Device communication data structures	147
	A9.3	Physical device object	149
	A9.3.1	Physical device attributes	149
	A9.3.2	Physical device lifecycle	150
	A9.3.3	Physical device flows	154
	A9.4	Virtual device object	160
	A9.4.1	Virtual device attributes	160
	A9.4.2	Virtual device invariants	160
	A9.4.3	Virtual device lifecycle	160
	A9.4.4	Virtual device flows	163
	A9.5	Realm management of an assigned device interface	164
	A9.5.1	Interruptible Realm device operations	164

A9.5.2	Realm retrieval of device attestation evidence	165
A9.5.3	Realm validation of device memory mappings	165
A9.5.4	Realm device attributes	166
A9.5.5	Realm device lifecycle	167
A9.5.6	Realm device flows	172
A9.6	Device access to a Protected IPA	174

Chapter A10

Planes

A10.1	Planes overview	176
A10.2	Planes exception model	177
A10.2.1	Plane exception model overview	177
A10.2.2	Plane entry	177
A10.2.3	Plane exit	178
A10.2.4	REC exit from Pn	179
A10.2.5	Pn execution of HVC and SMC	180
A10.2.6	Pn system registers	180
A10.3	Planes memory management	181
A10.3.1	Auxiliary RTT	181
A10.3.2	Stage 2 access permissions	182
A10.4	Planes interrupts	187
A10.5	Planes timers	188

Chapter A11

Realm memory encryption

Part B Interface

Chapter B1

Commands

B1.1	Overview	195
B1.2	Command definition	196
B1.2.1	Example command	196
B1.3	Command registers	197
B1.4	Command condition expressions	197
B1.5	Command context values	198
B1.6	Command failure conditions	199
B1.7	Command success conditions	200
B1.8	Concrete and abstract types	200
B1.9	Command footprint	200
B1.10	Command testing	201

Chapter B2

Interface versioning

Chapter B3

Command condition functions

B3.1	AddrInRange function	206
B3.2	AddrIsAligned function	206
B3.3	AddrIsAuxLive function	207
B3.4	AddrIsGranuleAligned function	207
B3.5	AddrIsProtected function	207
B3.6	AddrIsRttLevelAligned function	207
B3.7	AddrIsWithin function	207
B3.8	AddrRangelsAuxLive function	208
B3.9	AddrRangelsProtected function	208
B3.10	AddrRangelsWithin function	208
B3.11	AlignDownToRttLevel function	208
B3.12	AlignUpToRttLevel function	208
B3.13	AuxAlias16 function	209

B3.14	AuxAlias32 function	209
B3.15	AuxAligned16 function	209
B3.16	AuxAligned32 function	210
B3.17	AuxEqual16 function	210
B3.18	AuxEqual32 function	210
B3.19	AuxSort function	211
B3.20	AuxStateEqual16 function	211
B3.21	AuxStateEqual32 function	211
B3.22	AuxStates function	211
B3.23	CurrentRealm function	212
B3.24	CurrentRec function	212
B3.25	DeviceCommunicate function	212
B3.26	Equal function	212
B3.27	FeatureToRmi function	214
B3.28	FeatureToRsi function	214
B3.29	Gicv3ConfigIsValid function	215
B3.30	GranuleAccessPermitted function	215
B3.31	GranuleAt function	215
B3.32	ImplFeatures function	215
B3.33	MecMembers function	215
B3.34	MecPolicy function	216
B3.35	MecState function	216
B3.36	MemPermLabelSupported function	216
B3.37	MinAddress function	216
B3.38	MpidrEqual function	216
B3.39	MpidrIsUsed function	217
B3.40	PalsDelegable function	217
B3.41	PdevAt function	217
B3.42	PdevAuxCount function	217
B3.43	PdevFlags function	217
B3.44	PlaneRegIsValid function	218
B3.45	PlaneRegValue function	218
B3.46	PsciReturnCodeEncode function	218
B3.47	PsciReturnCodePermitted function	218
B3.48	RdevFromId function	218
B3.49	RdevFromIds function	219
B3.50	RdevIdsValid function	219
B3.51	RdevIdsAreValid function	219
B3.52	RdevMeasurementParamsValid function	219
B3.53	ReadMemory function	220
B3.54	RealmAt function	220
B3.55	RealmsLive function	220
B3.56	RealmParamsSupported function	220
B3.57	RealmRttBaseEqual function	222
B3.58	RealmVmidEqual function	222
B3.59	RecAt function	222
B3.60	RecAuxCount function	222
B3.61	RecFromMpidr function	223
B3.62	RecIndex function	223
B3.63	RecRipasResponseToRsi function	223
B3.64	RecS2APResponseToRsi function	223
B3.65	RemExtend function	224
B3.66	ResultEqual function	224
B3.67	RimExtendData function	224
B3.68	RimExtendRec function	224

B3.69	RimExtendRipas function	225
B3.70	RimExtendRipasForEntry function	225
B3.71	RimInit function	225
B3.72	RipasToRmi function	225
B3.73	RmiAddressRangesEqual16 function	226
B3.74	RmiAddressRangesEqual4 function	226
B3.75	RmiDevCommDataAt function	226
B3.76	RmiFeatureRegister0Decode function	227
B3.77	RmiFeatureRegisterEncode function	227
B3.78	RmiPdevEventIsValid function	228
B3.79	RmiPdevFlagsDecode function	228
B3.80	RmiPdevParamsAt function	228
B3.81	RmiPdevParamsIsValid function	228
B3.82	RmiRealmParamsAt function	228
B3.83	RmiRealmParamsIsValid function	229
B3.84	RmiRecParamsAt function	229
B3.85	RmiRecRunAt function	229
B3.86	RmiVdevFlagsDecode function	229
B3.87	RmiVdevParamsAt function	229
B3.88	RmiVdevParamsIsValid function	229
B3.89	RsiDeviceInfoAt function	229
B3.90	RsiDeviceMeasParamsAt function	230
B3.91	RsiFeatureRegisterEncode function	230
B3.92	RsiHostCallAt function	230
B3.93	RsiPlaneRunAt function	230
B3.94	RsiRealmConfigAt function	230
B3.95	RttAllEntriesContiguous function	231
B3.96	RttAllEntriesRipas function	231
B3.97	RttAllEntriesState function	231
B3.98	RttAt function	231
B3.99	RttConfigIsValid function	231
B3.100	RttDescriptorDecode function	231
B3.101	RttDescriptorIsValidForUnprotected function	232
B3.102	RttEntriesInRangeRipas function	232
B3.103	RttEntryAt function	232
B3.104	RttEntryIndex function	232
B3.105	RttEntryStateToRmi function	232
B3.106	RttFold function	233
B3.107	RttIsHomogeneous function	233
B3.108	RttIsLive function	233
B3.109	RttLevellsBlockOrPage function	233
B3.110	RttLevellsStarting function	234
B3.111	RttLevellsValid function	234
B3.112	RttLevelSize function	234
B3.113	RttsAllProtectedEntriesRipas function	234
B3.114	RttsAllProtectedEntriesState function	234
B3.115	RttsAllUnprotectedEntriesState function	235
B3.116	RttsGranuleState function	235
B3.117	RttSkipEntriesUnlessRipas function	235
B3.118	RttSkipEntriesUnlessState function	235
B3.119	RttSkipEntriesWithRipas function	235
B3.120	RttSkipNonLiveEntries function	236
B3.121	RttsStateEqual function	237
B3.122	RttWalk function	237
B3.123	RttWalkAnyNotAligned function	237

B3.124	TdildlsFree function	238
B3.125	ToAddress function	238
B3.126	ToBits64 function	238
B3.127	VdevAt function	238
B3.128	VdevAuxCount function	238
B3.129	VmidsAreFree function	238
B3.130	VmidsAreValid function	239

Chapter B4

Realm Management Interface

B4.1	RMI version	241
B4.2	RMI command return codes	241
B4.3	RMI commands	242
B4.3.1	RMI_DATA_CREATE command	244
B4.3.2	RMI_DATA_CREATE_UNKNOWN command	247
B4.3.3	RMI_DATA_DESTROY command	250
B4.3.4	RMI_DEV_MEM_MAP command	253
B4.3.5	RMI_DEV_MEM_UNMAP command	256
B4.3.6	RMI_FEATURES command	259
B4.3.7	RMI_GRANULE_DELEGATE command	260
B4.3.8	RMI_GRANULE_DEV_DELEGATE command	262
B4.3.9	RMI_GRANULE_DEV_UNDELEGATE command	264
B4.3.10	RMI_GRANULE_UNDELEGATE command	266
B4.3.11	RMI_MEC_SET_PRIVATE command	268
B4.3.12	RMI_MEC_SET_SHARED command	269
B4.3.13	RMI_PDEV_ABORT command	270
B4.3.14	RMI_PDEV_AUX_COUNT command	272
B4.3.15	RMI_PDEV_COMMUNICATE command	273
B4.3.16	RMI_PDEV_CREATE command	276
B4.3.17	RMI_PDEV_DESTROY command	279
B4.3.18	RMI_PDEV_GET_STATE command	281
B4.3.19	RMI_PDEV_IDE_RESET command	283
B4.3.20	RMI_PDEV_NOTIFY command	285
B4.3.21	RMI_PDEV_SET_PUBKEY command	287
B4.3.22	RMI_PDEV_STOP command	289
B4.3.23	RMI_PSCI_COMPLETE command	291
B4.3.24	RMI_REALM_ACTIVATE command	295
B4.3.25	RMI_REALM_CREATE command	297
B4.3.26	RMI_REALM_DESTROY command	301
B4.3.27	RMI_REC_AUX_COUNT command	303
B4.3.28	RMI_REC_CREATE command	304
B4.3.29	RMI_REC_DESTROY command	308
B4.3.30	RMI_REC_ENTER command	310
B4.3.31	RMI_RTT_AUX_CREATE command	313
B4.3.32	RMI_RTT_AUX_DESTROY command	316
B4.3.33	RMI_RTT_AUX_FOLD command	319
B4.3.34	RMI_RTT_AUX_MAP_PROTECTED command	322
B4.3.35	RMI_RTT_AUX_MAP_UNPROTECTED command	325
B4.3.36	RMI_RTT_AUX_UNMAP_PROTECTED command	328
B4.3.37	RMI_RTT_AUX_UNMAP_UNPROTECTED command	331
B4.3.38	RMI_RTT_CREATE command	334
B4.3.39	RMI_RTT_DESTROY command	337
B4.3.40	RMI_RTT_FOLD command	340
B4.3.41	RMI_RTT_INIT_RIPAS command	343
B4.3.42	RMI_RTT_MAP_UNPROTECTED command	346
B4.3.43	RMI_RTT_READ_ENTRY command	349

B4.3.44	RMI_RTT_SET_RIPAS command	352
B4.3.45	RMI_RTT_SET_S2AP command	355
B4.3.46	RMI_RTT_UNMAP_UNPROTECTED command	358
B4.3.47	RMI_VDEV_ABORT command	361
B4.3.48	RMI_VDEV_AUX_COUNT command	363
B4.3.49	RMI_VDEV_COMMUNICATE command	364
B4.3.50	RMI_VDEV_COMPLETE command	367
B4.3.51	RMI_VDEV_CREATE command	369
B4.3.52	RMI_VDEV_DESTROY command	373
B4.3.53	RMI_VDEV_GET_STATE command	376
B4.3.54	RMI_VDEV_STOP command	378
B4.3.55	RMI_VERSION command	380
B4.4	RMI types	382
B4.4.1	RmiAddressRange type	382
B4.4.2	RmiBoolean type	382
B4.4.3	RmiCommandReturnCode type	382
B4.4.4	RmiDataFlags type	383
B4.4.5	RmiDataMeasureContent type	383
B4.4.6	RmiDevCommData type	384
B4.4.7	RmiDevCommEnter type	384
B4.4.8	RmiDevCommExit type	385
B4.4.9	RmiDevCommExitFlags type	385
B4.4.10	RmiDevCommProtocol type	386
B4.4.11	RmiDevCommStatus type	387
B4.4.12	RmiDevDelegateFlags type	387
B4.4.13	RmiDevMemShared type	388
B4.4.14	RmiEmulatedMmio type	388
B4.4.15	RmiFeature type	388
B4.4.16	RmiFeatureRegister0 type	389
B4.4.17	RmiFeatureRegister1 type	390
B4.4.18	RmiHashAlgorithm type	391
B4.4.19	RmiInjectSea type	391
B4.4.20	RmiInterfaceVersion type	392
B4.4.21	RmiLfaPolicy type	392
B4.4.22	RmiPdevEvent type	393
B4.4.23	RmiPdevFlags type	393
B4.4.24	RmiPdevParams type	394
B4.4.25	RmiPdevProtConfig type	394
B4.4.26	RmiPdevState type	395
B4.4.27	RmiPlaneRttFeature type	395
B4.4.28	RmiPmuOverflowStatus type	396
B4.4.29	RmiRealmFlags0 type	396
B4.4.30	RmiRealmFlags1 type	397
B4.4.31	RmiRealmParams type	398
B4.4.32	RmiRecCreateFlags type	399
B4.4.33	RmiRecEnter type	400
B4.4.34	RmiRecEnterFlags type	401
B4.4.35	RmiRecExit type	401
B4.4.36	RmiRecExitFlags type	403
B4.4.37	RmiRecExitReason type	404
B4.4.38	RmiRecMpidr type	404
B4.4.39	RmiRecParams type	405
B4.4.40	RmiRecRun type	405
B4.4.41	RmiRecRunnable type	406
B4.4.42	RmiResponse type	406

B4.4.43	RmiRipas type	406
B4.4.44	RmiRttEntryState type	407
B4.4.45	RmiSignatureAlgorithm type	407
B4.4.46	RmiStatusCode type	408
B4.4.47	RmiTrap type	409
B4.4.48	RmiUnprotectedS2AP type	409
B4.4.49	RmiVdevAction type	409
B4.4.50	RmiVdevFlags type	410
B4.4.51	RmiVdevParams type	410
B4.4.52	RmiVdevState type	411

Chapter B5

Realm Services Interface

B5.1	RSI version	413
B5.2	RSI command return codes	413
B5.3	RSI commands	414
B5.3.1	RSI_ATTESTATION_TOKEN_CONTINUE command	415
B5.3.2	RSI_ATTESTATION_TOKEN_INIT command	417
B5.3.3	RSI_FEATURES command	419
B5.3.4	RSI_HOST_CALL command	420
B5.3.5	RSI_IPA_STATE_GET command	422
B5.3.6	RSI_IPA_STATE_SET command	424
B5.3.7	RSI_MEASUREMENT_EXTEND command	426
B5.3.8	RSI_MEASUREMENT_READ command	428
B5.3.9	RSI_MEM_GET_PERM_VALUE command	430
B5.3.10	RSI_MEM_SET_PERM_INDEX command	432
B5.3.11	RSI_MEM_SET_PERM_VALUE command	434
B5.3.12	RSI_PLANE_ENTER command	436
B5.3.13	RSI_PLANE_REG_READ command	438
B5.3.14	RSI_PLANE_REG_WRITE command	440
B5.3.15	RSI_RDEV_CONTINUE command	442
B5.3.16	RSI_RDEV_GET_INFO command	445
B5.3.17	RSI_RDEV_GET_INTERFACE_REPORT command	447
B5.3.18	RSI_RDEV_GET_MEASUREMENTS command	450
B5.3.19	RSI_RDEV_GET_STATE command	453
B5.3.20	RSI_RDEV_LOCK command	455
B5.3.21	RSI_RDEV_START command	457
B5.3.22	RSI_RDEV_STOP command	459
B5.3.23	RSI_RDEV_VALIDATE_MAPPING command	461
B5.3.24	RSI_REALM_CONFIG command	463
B5.3.25	RSI_VERSION command	465
B5.4	RSI types	467
B5.4.1	RsiBoolean type	467
B5.4.2	RsiCommandReturnCode type	467
B5.4.3	RsiDeviceInfo type	468
B5.4.4	RsiDeviceMeasurementsParams type	468
B5.4.5	RsiDeviceState type	468
B5.4.6	RsiDevMemCoherent type	469
B5.4.7	RsiDevMemShared type	469
B5.4.8	RsiFeature type	470
B5.4.9	RsiFeatureRegister0 type	470
B5.4.10	RsiGicOwner type	471
B5.4.11	RsiHashAlgorithm type	471
B5.4.12	RsiHostCall type	472
B5.4.13	RsiInterfaceVersion type	472
B5.4.14	RsiPlaneEnter type	472

B5.4.15	RsiPlaneEnterFlags type	473
B5.4.16	RsiPlaneExit type	474
B5.4.17	RsiPlaneExitReason type	475
B5.4.18	RsiPlaneRun type	475
B5.4.19	RsiRdevValidateloFlags type	475
B5.4.20	RsiRealmConfig type	476
B5.4.21	RsiResponse type	477
B5.4.22	RsiRipas type	477
B5.4.23	RsiRipasChangeDestroyed type	477
B5.4.24	RsiRipasChangeFlags type	478
B5.4.25	RsiTrap type	478

Chapter B6

Power State Control Interface

B6.1	PSCI overview	481
B6.2	PSCI version	481
B6.3	PSCI commands	482
B6.3.1	PSCI_AFFINITY_INFO command	483
B6.3.2	PSCI_CPU_OFF command	485
B6.3.3	PSCI_CPU_ON command	486
B6.3.4	PSCI_CPU_SUSPEND command	488
B6.3.5	PSCI_FEATURES command	489
B6.3.6	PSCI_SYSTEM_OFF command	490
B6.3.7	PSCI_SYSTEM_RESET command	491
B6.3.8	PSCI_VERSION command	492
B6.4	PSCI types	493
B6.4.1	PsciInterfaceVersion type	493
B6.4.2	PsciReturnCode type	493

Part C Constants and types

Chapter C1

RMM constants

C1.1	RMM_GRANULE_SIZE	496
C1.2	RMM_NUM_PERM_OVERLAY_INDICES	496
C1.3	RMM_RTT_BLOCK_LEVEL	496
C1.4	RMM_RTT_PAGE_LEVEL	496
C1.5	RMM_RTT_TREE_PRIMARY	497

Chapter C2

RMM types

C2.1	RmmAddressRange type	498
C2.2	RmmBoolean type	498
C2.3	RmmDataFlags type	499
C2.4	RmmDataMeasureContent type	499
C2.5	RmmDevCommState type	500
C2.6	RmmDevMemShared type	500
C2.7	RmmFeature type	501
C2.8	RmmFeatures type	501
C2.9	RmmGptEntry type	502
C2.10	RmmGranule type	502
C2.11	RmmGranuleState type	502
C2.12	RmmHashAlgorithm type	503
C2.13	RmmHipas type	504
C2.14	RmmLfaPolicy type	504
C2.15	RmmMeasurementDescriptorData type	504
C2.16	RmmMeasurementDescriptorRec type	505

C2.17	RmmMeasurementDescriptorRipas type	505
C2.18	RmmMecPolicy type	506
C2.19	RmmMecState type	506
C2.20	RmmMemPermLocked type	507
C2.21	RmmMemPerms type	507
C2.22	RmmPdev type	507
C2.23	RmmPdevProtConfig type	508
C2.24	RmmPdevState type	509
C2.25	RmmPhysicalAddressSpace type	509
C2.26	RmmPlaneRttFeature type	510
C2.27	RmmRdev type	510
C2.28	RmmRdevOperation type	510
C2.29	RmmRdevState type	511
C2.30	RmmRealm type	511
C2.31	RmmRealmMeasurement type	512
C2.32	RmmRealmState type	513
C2.33	RmmRec type	513
C2.34	RmmRecAttestState type	514
C2.35	RmmRecEmulatableAbort type	514
C2.36	RmmRecFlags type	515
C2.37	RmmRecPending type	515
C2.38	RmmRecResponse type	515
C2.39	RmmRecRunnable type	516
C2.40	RmmRecState type	516
C2.41	RmmRipas type	516
C2.42	RmmRipasChangeDestroyed type	517
C2.43	RmmRtt type	517
C2.44	RmmRttEntry type	517
C2.45	RmmRttEntryState type	518
C2.46	RmmRttWalkNotAligned type	519
C2.47	RmmRttWalkResult type	519
C2.48	RmmSystemRegisters type	519
C2.49	RmmVdev type	520
C2.50	RmmVdevState type	520

Chapter C3

Generic types

C3.1	Address type	521
C3.2	BitsN type	521
C3.3	IntN type	521
C3.4	UIntN type	522

Part D Usage

Chapter D1

Flows

D1.1	Granule delegation flows	525
D1.1.1	Granule delegation flow	525
D1.1.2	Granule undelegation flow	525
D1.2	Realm lifecycle flows	527
D1.2.1	Realm creation flow	527
D1.2.2	Realm Translation Table creation flow	527
D1.2.3	Initialize memory of New Realm flow	528
D1.2.4	REC creation flow	530
D1.2.5	Realm destruction flow	532
D1.3	Realm exception model flows	534

D1.3.1	Realm entry and exit flow	534
D1.3.2	Host call flow	534
D1.3.3	REC exit due to Data Abort fault flow	535
D1.3.4	MMIO emulation flow	536
D1.4	PSCI flows	538
D1.4.1	PSCI_CPU_ON flow	538
D1.5	Realm memory management flows	541
D1.5.1	Add memory to Active Realm flow	541
D1.5.2	NS memory flow	541
D1.5.3	RIPAS change flow	542
D1.5.4	S2AP change flow	543
D1.6	Realm interrupts and timers flows	545
D1.6.1	Interrupt flow	545
D1.6.2	Timer interrupt delivery flow	545
D1.7	Realm attestation flows	547
D1.7.1	Attestation token generation flow	547
D1.7.2	Handling interrupts during attestation token generation flow	547
D1.8	Realm device assignment flows	549

Chapter D2

Realm shared memory protocol

D2.1	Realm shared memory protocol description	551
D2.2	Realm shared memory protocol flow	551

Glossary

Preface

Quality level

This table below summarises the quality level of the features which have been added in version 1.1 of this specification.

Feature	Quality level
Realm device assignment	ALPHA
Planes	BETA
Realm memory encryption	BETA
Live firmware activation	ALPHA
Platform software component countersigners	BETA

Due to the fact that some features are at ALPHA, the overall quality level of this version of the specification is ALPHA.

Conventions

Typographical conventions

The typographical conventions are:

italic

Introduces special terminology, and denotes citations.

`monospace`

Used for pseudocode and source code examples.

Also used in the main text for instruction mnemonics and for references to other items appearing in pseudocode and source code examples.

SMALL CAPITALS

Used for some common terms such as IMPLEMENTATION DEFINED.

Used for a few terms that have specific technical meanings, and are included in the Glossary.

Red text

Indicates an open issue.

Blue text

Indicates a link. This can be

- A cross-reference to another location within the document
- A URL, for example <http://developer.arm.com>

Numbers

Numbers are normally written in decimal. Binary numbers are preceded by 0b, and hexadecimal numbers by 0x. In both cases, the prefix and the associated value are written in a monospace font, for example 0xFFFF0000. To improve readability, long numbers can be written with an underscore separator between every four characters, for example 0xFFFF_0000_0000_0000. Ignore any underscores when interpreting the value of a number.

Pseudocode descriptions

This book uses a form of pseudocode to provide precise descriptions of the specified functionality. This pseudocode is written in a monospace font. The pseudocode language is described in the Arm Architecture Reference Manual.

Addresses

Unless otherwise stated, the term *address* in this specification refers to a physical address.

Rules-based writing

This specification consists of a set of individual *content items*. A content item is classified as one of the following:

- Declaration
- Rule
- Goal
- Information
- Rationale
- Implementation note
- Software usage

Declarations and Rules are normative statements. An implementation that is compliant with this specification must conform to all Declarations and Rules in this specification that apply to that implementation.

Declarations and Rules must not be read in isolation. Where a particular feature is specified by multiple Declarations and Rules, these are generally grouped into sections and subsections that provide context. Where appropriate, these sections begin with a short introduction.

Arm strongly recommends that implementers read *all* chapters and sections of this document to ensure that an implementation is compliant.

Content items other than Declarations and Rules are informative statements. These are provided as an aid to understanding this specification.

Content item identifiers

A content item may have an associated identifier which is unique among content items in this specification.

After this specification reaches beta status, a given content item has the same identifier across subsequent versions of the specification.

Content item rendering

In this document, a content item is rendered with a token of the following format in the left margin: L_{iiii}

- L is a label that indicates the content class of the content item.
- $iiii$ is the identifier of the content item.

Content item classes

Declaration

A Declaration is a statement that does one or more of the following:

- Introduces a concept
- Introduces a term
- Describes the structure of data
- Describes the encoding of data

A Declaration does not describe behaviour.

A Declaration is rendered with the label D .

Rule

A Rule is a statement that describes the behaviour of a compliant implementation.

A Rule explains what happens in a particular situation.

A Rule does not define concepts or terminology.

A Rule is rendered with the label *R*.

Goal

A Goal is a statement about the purpose of a set of rules.

A Goal explains why a particular feature has been included in the specification.

A Goal is comparable to a “business requirement” or an “emergent property.”

A Goal is intended to be upheld by the logical conjunction of a set of rules.

A Goal is rendered with the label *G*.

Information

An Information statement provides information and guidance as an aid to understanding the specification.

An Information statement is rendered with the label *I*.

Rationale

A Rationale statement explains why the specification was specified in the way it was.

A Rationale statement is rendered with the label *X*.

Implementation note

An Implementation note provides guidance on implementation of the specification.

An Implementation note is rendered with the label *U*.

Software usage

A Software usage statement provides guidance on how software can make use of the features defined by the specification.

A Software usage statement is rendered with the label *S*.

Additional reading

This section lists publications by Arm and by third parties.

See Arm Developer (<http://developer.arm.com>) for access to Arm documentation.

- [1] *Introducing Arm CCA*. (ARM DEN 0125) Arm Limited.
- [2] *Arm Architecture Reference Manual Supplement, The Realm Management Extension (RME), for Armv9-A*. (ARM DDI 0615 A.d) Arm Ltd.
- [3] *Arm Architecture Reference Manual for A-Profile architecture*. (ARM DDI 0487 I.a) Arm Ltd.
- [4] *Arm CCA Security model*. (ARM DEN 0096) Arm Limited.
- [5] *Live Firmware Activation SMC Interface*. (ARM DEN 0147) Arm Limited.
- [6] *Arm Generic Interrupt Controller (GIC) Architecture Specification version 3 and version 4*. (ARM IHI 0069 G) Arm Ltd.
- [7] *Concise Binary Object Representation (CBOR)*. See <https://tools.ietf.org/html/rfc7049>
- [8] *CBOR Object Signing and Encryption (COSE)*. See <https://tools.ietf.org/html/rfc8152>
- [9] *Entity Attestation Token (EAT)*. See <https://datatracker.ietf.org/doc/draft-ietf-rats-eat/>

- [10] *Concise Data Definition Language (CDDL)*. See <https://tools.ietf.org/html/rfc8610>
- [11] *IANA Named Information Hash Algorithm Registry*. See <http://www.iana.org/assignments/named-information>
- [12] *SEC 1: Elliptic Curve Cryptography, version 2.0*. See <https://www.secg.org/sec1-v2.pdf>
- [13] *RME system architecture spec.* (ARM DEN 0129) Arm Ltd.
- [14] *PCI Express 6.0 specification*. See <https://pcisig.com/pci-express-6.0-specification>
- [15] *Secured Messages using SPDm Specification version 1.1.0*. See https://www.dmtf.org/sites/default/files/standards/documents/DSP0277_1.1.0.pdf
- [16] *Arm SMC Calling Convention*. (ARM DEN 0028 D) Arm Ltd.
- [17] *Arm Specification Language Reference Manual*. (ARM DDI 0612 00bet7) Arm Ltd.
- [18] *Security Protocol and Data Model (SPDM)*. See <https://www.dmtf.org/dsp/DSP0274>
- [19] *Secure Hash Standard (SHS)*. See <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.180-4.pdf>
- [20] *RSA Cryptography Specifications Version 2.2*. See <https://datatracker.ietf.org/doc/rfc8017/>
- [21] *Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA)*. See <https://datatracker.ietf.org/doc/html/rfc6979>
- [22] *Arm Power State Coordination Interface (PSCI)*. (ARM DEN 0022 D.b) Arm Ltd.

Feedback

Arm welcomes feedback on its documentation.

Feedback on this book

If you have any comments or suggestions for additions and improvements, create a ticket at <https://support.developer.arm.com>.

As part of the ticket, include:

- The title (Realm Management Monitor specification).
- The number (DEN0137 1.1-alp9).
- The section name(s) to which your comments refer.
- The page number(s) to which your comments apply.
- The rule identifier(s) to which your comments apply, if applicable.
- A concise explanation of your comments.

Arm also welcomes general suggestions for additions and improvements.

Note

Arm tests PDFs only in Adobe Acrobat and Acrobat Reader, and cannot guarantee the appearance or behavior of any document when viewed with any other PDF reader.

Open issues

The following table lists known open issues in this version of the document.

Key	Description
-	Consider prefixing states with a namespace, for example changing DEV_PRIVATE to GRANULE_DEV_PRIVATE. This would avoid potential confusion with the RmiDevMemShared and RmmDevMemShared values.
-	Consider how teardown of DRAM mappings (via RMI_DATA_DESTROY) composes with teardown of device memory mappings (via RMI_DEV_MEM_UNMAP). In each case, the command returns the IPA of the next live entry - but it doesn't tell the caller whether this is DRAM or IO. How then can the caller know which of the two commands to call next, while still avoiding a (race-prone) call to RMI_RTT_READ_ENTRY?
-	In RMI_RTT_SET_RIPAS, consider how to combine: <ul style="list-style-type: none">• Modification of a range of RTT entries in a single command, and• Checking of output address and HIPAS values against rec.ripas_dev_pa and rec.ripas_dev_shared respectively.

DRAFT

DRAFT

Part A
Architecture

Chapter A1

Overview

The RMM is a software component which forms part of a system which implements the Arm Confidential Compute Architecture (Arm CCA). Arm CCA is an architecture which provides protected execution environments called *Realms*.

The threat model which Arm CCA is designed to address is described in [Introducing Arm CCA \[1\]](#).

The hardware architecture of Arm CCA is called the Realm Management Extension (RME), and is described in [Arm Architecture Reference Manual Supplement, The Realm Management Extension \(RME\), for Armv9-A \[2\]](#).

A1.1 Confidential computing

The Armv8-A architecture ([Arm Architecture Reference Manual for A-Profile architecture \[3\]](#)) includes mechanisms that establish a privilege hierarchy. Software operating at higher privilege levels is responsible for managing the resources (principally memory and processor cycles) that are used by entities at lower privilege levels.

Prior to Arm CCA, resource management was coupled with a right of access. That is, a resource that is managed by a higher-privileged entity is also accessible by it. A *Realm* is a protected execution environment for which this coupling is broken, so that the right to manage resources is separated from the right to access those resources.

The purpose of a Realm is to provide to the Realm owner an environment for confidential computing, without requiring the Realm owner to trust the software components that manage the resources used by the Realm.

Construction of a Realm, and allocation of resources to a Realm at runtime, are the responsibility of the Virtual Machine Monitor (VMM). In this specification, the term *Host* is used to refer to the VMM.

See also:

- [A2.1 Realm](#)

A1.2 System software components

The system software architecture of Arm CCA is summarised in the following figure.




Figure A1.1: System software architecture

The components shown in the diagram are listed below.

Component	Description
Monitor	The most privileged software component, which is responsible for switching between the Security states used at EL2, EL1 and EL0.
Realm	A protected execution environment.
Realm Management Monitor (RMM)	The software component which is responsible for the management of Realms.
Virtual Machine (VM)	An execution environment within which an operating system can run. Note that a Realm is a VM which executes in the Realm security state.
Hypervisor	The software component which is responsible for the management of VMs.
Secure Partition Manager (SPM)	The software component which is responsible for the management of Secure Partitions.
Trusted OS (TOS)	An operating system which runs in a Secure Partition.
Trusted Application (TA)	An application hosted by a TOS.

A1.3 Realm Management Monitor

The Realm Management Monitor (RMM) is the system component that is responsible for the management of Realms.

The responsibilities of the RMM are to:

- Provide services that allow the Host to create, populate, execute and destroy Realms.
- Provide services that allow the initial configuration and contents of a Realm to be attested.
- Protect the confidentiality and integrity of Realm state during the lifetime of the Realm.
- Protect the confidentiality of Realm state during and following destruction of the Realm.
- Act as the Trusted Security Manager (TSM) in Realm device assignment.

The RMM exposes the following interfaces, which are accessed via SMC instructions, to the Host:

- The *Realm Management Interface* (RMI), which provides services for the creation, population, execution and destruction of Realms.

The RMM exposes the following interfaces, which are accessed via SMC instructions, to Realms:

- The *Realm Services Interface* (RSI), which provides services used to manage resources allocated to the Realm, and to request an attestation report.
- The *Power State Coordination Interface* (PSCI), which provides services used to control power states of VPEs within a Realm. Note that the HVC conduit for PSCI is not supported for Realms.

The RMM operates by manipulating data structures which are stored in memory accessible only to the RMM.

See also:

- [Chapter A9 Realm device assignment](#)
- [Chapter B4 Realm Management Interface](#)
- [Chapter B5 Realm Services Interface](#)
- [Chapter B6 Power State Control Interface](#)

DRAFT

Chapter A2

Concepts

This chapter introduces the following concepts which are central to the RMM architecture:

- [A2.1 Realm](#)
- [A2.2 Granule](#)
- [A2.3 Realm Execution Context](#)

A2.1 Realm

This section describes the concept of a Realm.

A2.1.1 Overview

D_DLRSR A *Realm* is an execution environment which is protected from agents in the Non-secure and Secure Security states, and from other Realms.

A2.1.2 Realm execution environment

I_LQYLY The execution environment of a Realm is an EL0 + EL1 environment, as described in [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#).

A2.1.2.1 Realm registers

R_NJHQK On first entry to a Realm VPE, PE state is initialized according to “PE state on reset to AArch64 state” in [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#), except for GPR and PC values which are specified by the Host during Realm creation.

G_ZFCQX Confidentiality is guaranteed for a Realm VPE’s general purpose and SIMD / floating point registers.

G_QHZCS Confidentiality is guaranteed for other Realm VPE register state (including stack pointer, program counter and EL0 / EL1 system registers).

G_XRMHP Integrity is guaranteed for a Realm VPE’s general purpose and SIMD / floating point registers.

G_YKRWG Integrity is guaranteed for other Realm VPE register state (including stack pointer, program counter and EL0 / EL1 system registers).

I_GPGFB A Realm can use a Host call to pass arguments to the Host and receive results from the Host.

See also:

- [A2.3 Realm Execution Context](#)
- [A4.5 Host call](#)
- [B4.3.25 RMI_REALM_CREATE command](#)

A2.1.2.2 Realm memory

I_TQMMZ A Realm is able to determine whether a given IPA is *protected* or *unprotected*.

G_LQFQH Confidentiality is guaranteed for memory contents accessed via a protected address. Informally, this means that a change to the contents of such a memory location is not observable by any agent outside the *CCA platform*.

G_QMLCJ Integrity is guaranteed for memory contents accessed via a protected address. Informally, this means that the Realm does not observe the contents of the location to change unless the Realm itself has either written a different value to the location, or provided consent to the RMM for integrity of the location to be violated.

See also:

- [A5.2.1 Realm IPA space](#)

A2.1.2.3 Realm processor features

R_JGHYJ The value returned to a Realm from reading a feature register is architecturally valid and describes the set of features which are present in the Realm’s execution environment.

I_KKBDP The RMM may suppress a feature which is supported by the underlying hardware platform, if exposing that feature to a Realm could lead to a security vulnerability.

See also:

- [Chapter A3 Feature discovery and configuration](#)

A2.1.2.4 IMPDEF system registers

R_{FQCKH} A Realm read from or write to an IMPLEMENTATION DEFINED system register causes an Unknown exception taken to the Realm.

A2.1.3 Realm attributes

This section describes the attributes of a Realm.

D_{JSGFY} A *Realm attribute* is a property of a Realm whose value can be observed or modified either by the Host or by the Realm.

I_{TTDVX} An example of a way in which a Realm attribute may be observable is the outcome of an RMM command.

D_{MHJCK} The attributes of a Realm are summarized in the following table.

Name	Type	Description
feat_lpa2	RmmFeature	Whether LPA2 is enabled for this Realm
ipa_width	UInt8	IPA width in bits
measurements	RmmRealmMeasurement [5]	Realm measurements
hash_algo	RmmHashAlgorithm	Algorithm used to compute Realm measurements
rec_index	UInt64	Index of next REC to be created
rtt_base	Address [4]	Realm Translation Table base addresses If rtt_tree_pp is FEATURE_FALSE then only the first entry is valid. If rtt_tree_pp is FEATURE_TRUE then only the first (num_aux_planes + 1) entries are valid.
rtt_level_start	Int64	RTT starting level
rtt_num_start	UInt64	Number of physically contiguous starting level RTTs
state	RmmRealmState	Lifecycle state
vmid	Bits16 [4]	Virtual Machine Identifiers If rtt_tree_pp is FEATURE_FALSE then only the first entry is valid. If rtt_tree_pp is FEATURE_TRUE then only the first (num_aux_planes + 1) entries are valid.
rpv	Bits512	Realm Personalization Value
feat_da	RmmFeature	Whether Realm device assignment is enabled for this Realm
rtt_tree_pp	RmmFeature	Whether this Realm has an RTT per Plane
num_aux_planes	UInt64	Number of auxiliary Planes
overlay_perms	RmmMemPerms [4]	Memory overlay permissions
overlay_locked	RmmMemPermLocked [16]	Whether memory overlay value is locked
lfa_policy	RmmLfaPolicy	Live Firmware Activation policy for components within the Realm's TCB
mecid	Bits64	Memory Encryption Context Identifier
mec_policy	RmmMecPolicy	MEC policy

Name	Type	Description
num_recs	UInt64	Number of RECs owned by this Realm
num_vdevs	UInt64	Number of VDEVs which have been assigned to this Realm

D _{MGGPT}	A <i>Realm Initial Measurement</i> (RIM) is a measurement of the configuration and contents of a Realm at the time of activation.	
D _{GRFCS}	A <i>Realm Extensible Measurement</i> (REM) is a measurement value which can be extended during the lifetime of a Realm.	
I _{FMPYL}	Attributes of a Realm include an array of measurement values. The first entry in this array is a RIM. The remaining entries in this array are REMs.	
X _{DNDKV}	During Realm creation, the Host provides ipa_width, rtt_level_start and rtt_num_start values as Realm parameters. According to the VMSA, the rtt_num_start value is architecturally defined as a function of the ipa_width and rtt_level_start values. It would therefore have been possible to design the Realm creation interface such that the Host provided only the ipa_width and rtt_level_start values. However, this would potentially allow a Realm to be successfully created, but with a configuration which did not match the Host's intent. For this reason, it was decided that the Host should specify all three values explicitly, and that Realm creation should fail if the values are not consistent. See Arm Architecture Reference Manual for A-Profile architecture [3] for further details.	
I _{QRVTT}	The VMID of a Realm is chosen by the Host. The VMID must be within the range supported by the hardware platform. The RMM ensures that every Realm on the system has a unique VMID.	
D _{FTWBK}	A <i>Realm Personalization Value</i> (RPV) is a provided by the Host, to distinguish between Realms which have the same Realm Initial Measurement, but different behavior.	
S _{FCNBF}	Possible uses of the RPV include: <ul style="list-style-type: none"> • A GUID • Hash of Realm Owner public key • Hash of a “personalisation document” which is provided to the Realm via a side-band (for example, via NS memory) and contains configuration information used by Realm software. 	
I _{ZFSWC}	The RMM treats the RPV as an opaque value.	
I _{BFSRK}	The RPV is included in the Realm attestation report as a separate claim.	
I _{MFRXD}	The RPV is included in the output of the RSI_REALM_CONFIG command.	
I ₀₀₀₁	If Realm device assignment is not enabled for a Realm then all of the following are true: <ul style="list-style-type: none"> • Assignment of a virtual device to the Realm by execution of RMI_VDEV_CREATE fails. • The device assignment feature is reported to the Realm by RSI_FEATURES as not enabled. Consequently, execution of any RSI_RDEV command fails. See also: <ul style="list-style-type: none"> • A2.1.5 Realm lifecycle • A2.3 Realm Execution Context • A3.3 Realm LPA2 and IPA width • A5.2.1 Realm IPA space • A5.5 Realm Translation Table • A7.1 Realm measurements • A7.2.3.1.3 Realm Personalization Value claim • B4.3.51 RMI_VDEV_CREATE command • B5.3.3 RSI_FEATURES command • B5.3.24 RSI_REALM_CONFIG command 	

- [C2.30 RmmRealm type](#)

A2.1.4 Realm liveness

- D_{W_{TX}TJ}** *Realm liveness* is a property which means that there exists one or more Granules, other than the RD and the starting level RTTs, which are owned by the Realm.
- I_{PVPQB}** If a Realm is live, it cannot be destroyed.
- D_{PCKRN}** A Realm is *live* if any of the following is true:
- The number of RECs owned by the Realm is not zero
 - A starting level RTT of the Realm is live
- I_{VKKPJ}** If a Realm owns a non-zero number of Data Granules, this implies that it has a starting level RTT which is live, and therefore that the Realm itself is live.
- See also:
- [A2.1.5 Realm lifecycle](#)
 - [A2.2.2 Granule lifecycle](#)
 - [A2.2.3 Granule ownership](#)
 - [A2.3 Realm Execution Context](#)
 - [A5.5.8 RTTE liveness and RTT liveness](#)
 - [B3.55 RealmIsLive function](#)
 - [B4.3.26 RMI_REALM_DESTROY command](#)

A2.1.5 Realm lifecycle

See also:

- [Chapter A3 Feature discovery and configuration](#)
- [D1.2 Realm lifecycle flows](#)

A2.1.5.1 States

D_{GDQPJ} The states of a Realm are listed below.

State	Description
REALM_NEW	Under construction. Not eligible for execution.
REALM_ACTIVE	Eligible for execution.
REALM_SYSTEM_OFF	System has been turned off. Not eligible for execution.

A2.1.5.2 State transitions

I_{RRHFG} Permitted Realm state transitions are shown in the following table. The rightmost column lists the events which can cause the corresponding state transition.

A transition from the pseudo-state *NULL* represents creation of a Realm object. A transition to the pseudo-state *NULL* represents destruction of a Realm object.

From state	To state	Events
<i>NULL</i>	REALM_NEW	RMI_REALM_CREATE
REALM_NEW	<i>NULL</i>	RMI_REALM_DESTROY

From state	To state	Events
REALM_ACTIVE	<i>NULL</i>	RMI_REALM_DESTROY
REALM_SYSTEM_OFF	<i>NULL</i>	RMI_REALM_DESTROY
REALM_NEW	REALM_ACTIVE	RMI_REALM_ACTIVATE
REALM_ACTIVE	REALM_SYSTEM_OFF	PSCI_SYSTEM_OFF PSCI_SYSTEM_RESET

I_{YCPWW}

Permitted Realm state transitions are shown in the following figure. Each arc is labeled with the events which can cause the corresponding state transition.

A transition from the pseudo-state *NULL* represents creation of an RD. A transition to the pseudo-state *NULL* represents destruction of an RD.

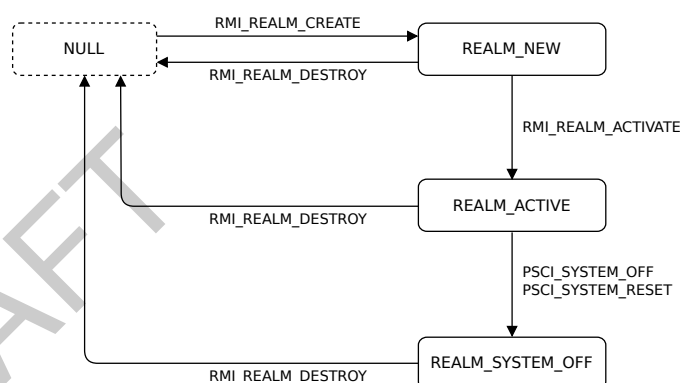


Figure A2.1: Realm state transitions

See also:

- [B6.3.6 PSCI_SYSTEM_OFF command](#)
- [B6.3.7 PSCI_SYSTEM_RESET command](#)
- [B4.3.24 RMI_REALM_ACTIVATE command](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [B4.3.26 RMI_REALM_DESTROY command](#)

A2.1.6 Realm parameters

D_{TGMVZ}

A *Realm parameter* is a value which is provided by the Host during Realm creation.

See also:

- [A2.1.3 Realm attributes](#)
- [Chapter A3 Feature discovery and configuration](#)
- [B3.82 RmiRealmParamsAt function](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [B4.4.31 RmiRealmParams type](#)

A2.1.7 Realm Descriptor

D_{TNSBY}

A *Realm Descriptor* (RD) is an RMM data structure which stores attributes of a Realm.

D_{GGKWX}

The size of an RD is one Granule.

See also:

- [A2.1.3 Realm attributes](#)
- [A2.2.2 Granule lifecycle](#)

DRAFT

A2.2 Granule

This section describes the concept of a Granule.

D_{NBXXX} A *Granule* is a unit of physical memory whose size is 4KB.

I_{DJGZW} A Granule may be used for one of the following purposes:

- To store code or data used by the Host
- To store code or data used by software in the Secure Security state
- To store code or data used by a Realm
- To access device memory, such as a memory-mapped register interface
- To store data used by the RMM to manage a Realm

The use of a Granule is reflected in its lifecycle state.

D_{ZVRXC} A Granule is *delegable* if it can be delegated by the Host for use by the RMM or by a Realm.

U_{KHKLP} In a typical implementation, delegable memory includes memory which is presented to the Host as RAM.

Examples of non-delegable memory may include the following:

- Memory which is carved out for use by the Root world, the RMM or the Secure world

I₀₀₀₂ In a system which supports Realm device assignment, delegable memory includes device memory (that is, regions of the system physical address map which are reserved for use as PCIe device BARs).

I₀₀₀₃ The initial state of delegable memory reflects whether it is:

- Delegable DRAM (initial state is UNDELEGATED), or
- Delegable device memory (initial state is DEV_UNDELEGATED)

See also:

- [A2.2.1 Granule attributes](#)
- [A2.2.2 Granule lifecycle](#)
- [Chapter A9 Realm device assignment](#)

A2.2.1 Granule attributes

This section describes the attributes of a Granule.

D_{JPBBC} A *Granule attribute* is a property of a Granule whose value can be observed or modified either by the Host or by a Realm.

I_{WVXGK} Examples of ways in which a Granule attribute may be observable include the outcome of an RMM command, and whether a memory access generates a fault.

D_{DVMRF} The attributes of a Granule are summarized in the following table.

Name	Type	Description
gpt	RmmGptEntry	GPT entry
state	RmmGranuleState	Lifecycle state

See also:

- [A2.1 Realm](#)
- [A2.1.7 Realm Descriptor](#)
- [A2.2.2 Granule lifecycle](#)
- [B3.30 GranuleAccessPermitted function](#)

- C2.10 RmmGranule type

A2.2.2 Granule lifecycle

A2.2.2.1 States

D_{MPLGT}

The states of a Granule are listed below.

For each state, the corresponding GPT entry value is shown.

Granule state	Description	GPT entry
UNDELEGATED	Not delegated for use by the RMM.	Not GPT_REALM
DELEGATED	Delegated for use by the RMM.	GPT_REALM
RD	Realm Descriptor.	GPT_REALM
REC	Realm Execution Context.	GPT_REALM
REC_AUX	Realm Execution Context auxiliary Granule.	GPT_REALM
DATA	Realm code or data.	GPT_REALM
RTT	Realm Translation Table.	GPT_REALM
PDEV	Physical device.	GPT_REALM
PDEV_AUX	Physical device auxiliary Granule.	GPT_REALM
VDEV	Virtual device.	GPT_REALM
DEV_UNDELEGATED	Device memory, not delegated for use by the RMM.	Not GPT_REALM
DEV_DELEGATED_PRIVATE	Device memory, delegated to the RMM and accessible via Realm PAS only.	GPT_REALM
DEV_DELEGATED_SHARED	Device memory, delegated to the RMM and accessible via any PAS.	GPT_AAP
DEV_PRIVATE	Device memory, mapped into a Realm and inaccessible by other requestors.	GPT_REALM
DEV_SHARED	Device memory, mapped into a Realm and also accessible by other requestors.	GPT_AAP

Issue Consider prefixing states with a namespace, for example changing DEV_PRIVATE to GRANULE_DEV_PRIVATE. This would avoid potential confusion with the RmiDevMemShared and RmmDevMemShared values.

I_{MPLGT}

If the state of a Granule is UNDELEGATED or DEV_UNDELEGATED then the RMM does not prevent the GPT entry of the Granule from being changed by another agent to any value except GPT_REALM.

D_{VRSKZ}

An NS Granule is a Granule whose GPT entry is GPT_NS.

A2.2.2.2 State transitions

I_{ZJBT}

Permitted Granule state transitions are shown in the following table. The rightmost column lists the events which can cause the corresponding state transition.

From state	To state	Events
UNDELEGATED	DELEGATED	RMI_GRANULE_DELEGATE
DELEGATED	UNDELEGATED	RMI_GRANULE_UNDELEGATE
DELEGATED	RD	RMI_REALM_CREATE
RD	DELEGATED	RMI_REALM_DESTROY
DELEGATED	DATA	RMI_DATA_CREATE RMI_DATA_CREATE_UNKNOWN
DATA	DELEGATED	RMI_DATA_DESTROY
DELEGATED	REC	RMI_REC_CREATE
REC	DELEGATED	RMI_REC_DESTROY
DELEGATED	REC_AUX	RMI_REC_CREATE
REC_AUX	DELEGATED	RMI_REC_DESTROY
DELEGATED	RTT	RMI_REALM_CREATE RMI_RTT_CREATE
RTT	DELEGATED	RMI_REALM_DESTROY RMI_RTT_DESTROY
DELEGATED	PDEV	RMI_PDEV_CREATE
PDEV	DELEGATED	RMI_PDEV_DESTROY
DELEGATED	PDEV_AUX	RMI_PDEV_CREATE
PDEV_AUX	DELEGATED	RMI_PDEV_DESTROY
DELEGATED	VDEV	RMI_VDEV_CREATE
VDEV	DELEGATED	RMI_VDEV_DESTROY
DEV_UNDELEGATED	DEV_DELEGATED_PRIVATE	RMI_GRANULE_DEV_DELEGATE
DEV_DELEGATED_PRIVATE	DEV_UNDELEGATED	RMI_GRANULE_DEV_UNDELEGATE
DEV_UNDELEGATED	DEV_DELEGATED_SHARED	RMI_GRANULE_DEV_DELEGATE
DEV_DELEGATED_SHARED	DEV_UNDELEGATED	RMI_GRANULE_DEV_UNDELEGATE
DEV_DELEGATED_PRIVATE	DEV_PRIVATE	RMI_DEV_MEM_MAP
DEV_PRIVATE	DEV_DELEGATED_PRIVATE	RMI_DEV_MEM_UNMAP
DEV_DELEGATED_SHARED	DEV_SHARED	RMI_DEV_MEM_MAP
DEV_SHARED	DEV_DELEGATED_SHARED	RMI_DEV_MEM_UNMAP

I_VVGVM

Permitted Granule state transitions are shown in the following figures. Each arc is labeled with the events which can cause the corresponding state transition.

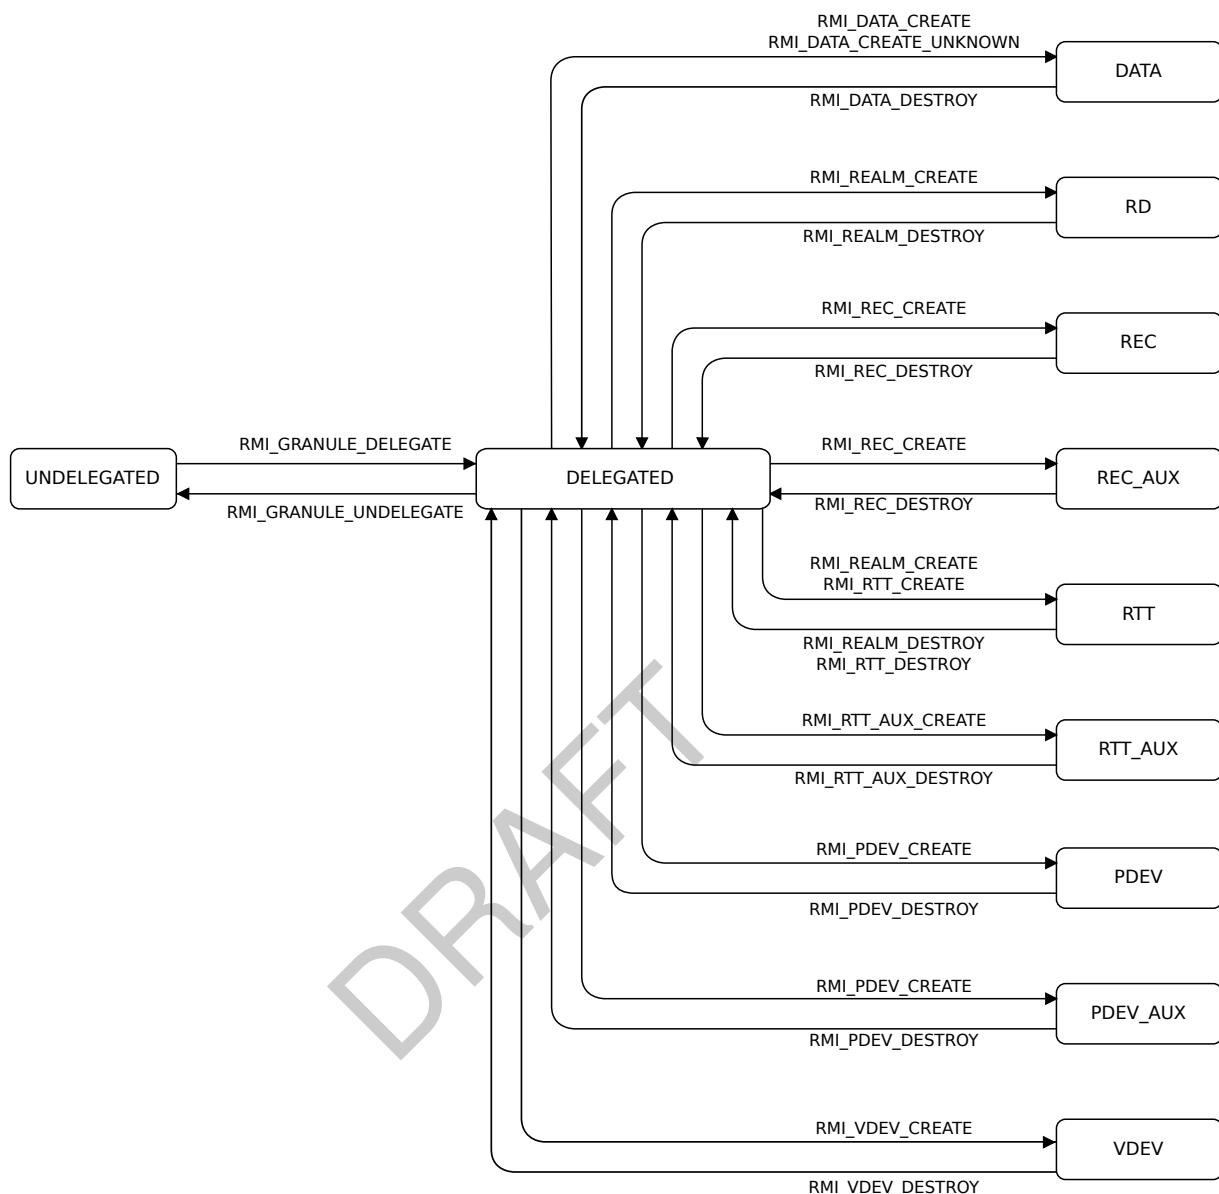


Figure A2.2: Granule state transitions for non-device memory

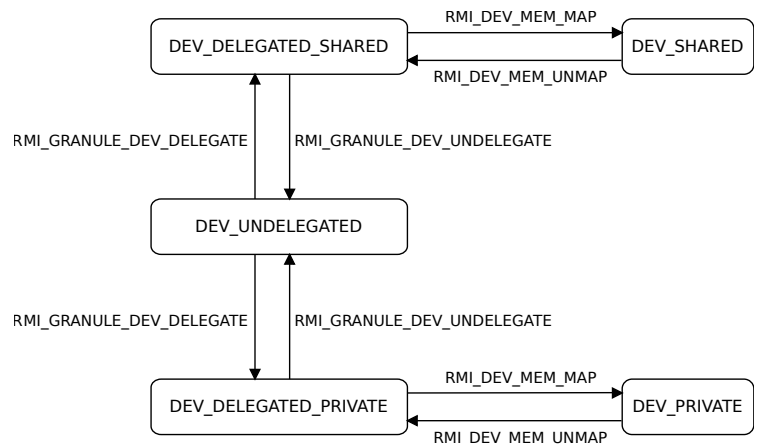


Figure A2.3: Granule state transitions for device memory

See also:

- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.4 RMI_DEV_MEM_MAP command](#)
- [B4.3.5 RMI_DEV_MEM_UNMAP command](#)
- [B4.3.7 RMI_GRANULE_DELEGATE command](#)
- [B4.3.8 RMI_GRANULE_DEV_DELEGATE command](#)
- [B4.3.9 RMI_GRANULE_DEV_UNDELEGATE command](#)
- [B4.3.10 RMI_GRANULE_UNDELEGATE command](#)
- [B4.3.16 RMI_PDEV_CREATE command](#)
- [B4.3.17 RMI_PDEV_DESTROY command](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [B4.3.26 RMI_REALM_DESTROY command](#)
- [B4.3.28 RMI_REC_CREATE command](#)
- [B4.3.29 RMI_REC_DESTROY command](#)
- [B4.3.38 RMI_RTT_CREATE command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)
- [B4.3.51 RMI_VDEV_CREATE command](#)
- [B4.3.52 RMI_VDEV_DESTROY command](#)

A2.2.3 Granule ownership

$\mathcal{I}_{\text{DMVQM}}$

A Granule whose state is none of the following is owned by a Realm:

- UNDELEGATED
- DELEGATED
- PDEV
- DEV_UNDELEGATED
- DEV_DELEGATED_PRIVATE
- DEV_DELEGATED_SHARED

$\mathcal{I}_{\text{PRNTM}}$

The owner of a Granule is identified by the address of a Realm Descriptor (RD).

$\mathcal{I}_{\text{ZXBZM}}$

For a Granule whose state is RD, the ownership relation is recursive: the owning Realm is identified by the address of the RD itself.

$\mathcal{I}_{\text{TYHTD}}$

A Granule whose state is RTT is one of the following:

- A starting level RTT. The address of this RTT is stored in the RD of the owning Realm.

- A non-starting level RTT. The address of this RTT is stored in its parent RTT, in an RTT entry whose state is TABLE. Recursively following the parent relationship leads to the RD of the owning Realm.

I _{QCNRM}	A Granule whose state is DATA is mapped at a Protected IPA, in an RTT entry whose state is ASSIGNED. The Realm which owns the RTT is the owner of the DATA Granule.
I _{HHPVB}	A REC has an “owner” attribute which points to the RD of the owning Realm.
X _{NDNHG}	A REC is not mapped at a Protected IPA. Its ownership therefore needs to be recorded explicitly.
I ₀₀₀₄	A VDEV has an “owner” attribute which points to the RD of the owning Realm.
X ₀₀₀₅	A VDEV is not mapped at a Protected IPA. Its ownership therefore needs to be recorded explicitly.

See also:

- [A2.1 Realm](#)
- [A2.1.7 Realm Descriptor](#)
- [A2.3 Realm Execution Context](#)
- [A5.2.1 Realm IPA space](#)
- [A5.5 Realm Translation Table](#)
- [Chapter A9 Realm device assignment](#)
- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.28 RMI_REC_CREATE command](#)
- [B4.3.38 RMI_RTT_CREATE command](#)

A2.2.4 Granule wiping

R _{TMGSL}	When the state of a Granule has transitioned from <i>P</i> to DELEGATED and then to any other state, any content associated with <i>P</i> has been <i>wiped</i> .
X _{CTGQZ}	Any sequence of Granule state transitions which passes through the DELEGATED state causes the Granule contents to be wiped. This is necessary to ensure that information does not leak from one Realm to another, or from a Realm to the Host. Note that no agent can observe the contents of a Granule while its state is DELEGATED.
R ₀₀₀₆	When the state of a Granule has transitioned from <i>P</i> to DEV_DELEGATED_PRIVATE and then to any other state, any content associated with <i>P</i> has been <i>wiped</i> .
X ₀₀₀₇	Any sequence of Granule state transitions which passes through the DEV_DELEGATED_PRIVATE state causes the Granule contents to be wiped. This is necessary to ensure that information does not leak from one Realm to another, or from a Realm to the Host. Note that no agent can observe the contents of a Granule while its state is DEV_DELEGATED_PRIVATE.
D _{WTWJR}	<i>Wiping</i> is an operation which changes the observable value of a memory location from <i>X</i> to <i>Y</i> , such that the value <i>X</i> cannot be determined from the value <i>Y</i> .
R _{BSXXV}	Wiping of a memory location does not reveal, directly or indirectly, any confidential Realm data.
I _{MRPCQ}	Wiping is not guaranteed to be implemented as zero filling.
S _{VJWYH}	Realm software should not assume that the initial contents of uninitialized memory (that is, Realm IPA space which is backed by DATA Granules created using RMI_DATA_CREATE_UNKNOWN) are zero.

See also:

- [Arm CCA Security model \[4\]](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)

A2.3 Realm Execution Context

This section describes the concept of a Realm Execution Context (REC).

A2.3.1 Overview

D_{LRFCP}

A *Realm Execution Context* (REC) is an R-EL0&1 execution context which is associated with a Realm VPE.

A *REC object* is an RMM data structure which is used to store the register state of a REC.

See also:

- [A2.1.2 Realm execution environment](#)
- [Chapter A4 Realm exception model](#)

A2.3.2 REC attributes

This section describes the attributes of a REC.

D_{ZLGLT}

A *REC attribute* is a property of a REC whose value can be observed or modified either by the Host or by the Realm which owns the REC.

I_{CSGGT}

Examples of ways in which a REC attribute may be observable include the outcome of an RMM command, and the PE state following Realm entry.

D_{LQSFT}

The attributes of a REC are summarized in the following table.

Name	Type	Description
attest_state	RmmRecAttestState	Attestation token generation state
attest_challenge	Bits512	Challenge for under-construction attestation token
aux	Address[16]	Addresses of auxiliary Granules
emulatable_abort	RmmRecEmulatableAbort	Whether the most recent exit from this REC was due to an Emulatable Data Abort
flags	RmmRecFlags	Flags which control REC behavior
gprs	Bits64[32]	General-purpose register values
mpidr	Bits64	MPIDR value
owner	Address	PA of RD of Realm which owns this REC
pc	Bits64	Program counter value
pending	RmmRecPending	Whether a REC operation is pending
vdev_id	Bits64	Virtual device ID
inst_id	UInt64	Device instance ID
inst_id_valid	RmmBoolean	Whether device instance ID is valid
state	RmmRecState	Lifecycle state
sysregs	RmmSystemRegisters	EL1 and EL0 system register values
ripas_addr	Address	Next address to be processed in RIPAS change
ripas_top	Address	Top address of pending RIPAS change
ripas_value	RmmRipas	RIPAS value of pending RIPAS change

Name	Type	Description
ripas_destroyed	RmmRipasChangeDestroyed	Whether a RIPAS change from DESTROYED should be permitted
ripas_response	RmmRecResponse	Host response to RIPAS change request
ripas_dev_pa	Address	Base PA of device memory region, if RIPAS change is pending due to execution of RSI_RDEV_VALIDATE_MAPPING
ripas_dev_shared	RmmDevMemShared	Value of shared bit, if RIPAS change is pending due to execution of RSI_RDEV_VALIDATE_MAPPING
s2ap_addr	Address	Next address to be processed in S2AP change
s2ap_top	Address	Top address of pending S2AP change
s2ap_overlay	UInt3	Overlay index of pending S2AP change
s2ap_response	RmmRecResponse	Host response to S2AP change request
gic_owner	UInt64	Index of Plane which is the GIC owner

I _{PVMTY}	The <i>aux</i> attribute of a REC is a list of <i>auxiliary Granules</i> .
I _{RWFZF}	The number of auxiliary Granules required for a REC is returned by the RMI_REC_AUX_COUNT command.
X _{LRWHB}	Depending on the configuration of the CCA platform and of the Realm, the amount of storage space required for a REC may exceed a single Granule.
I _{TGLBK}	The number of auxiliary Granules required for a REC can vary between Realms on a CCA platform.
R _{MMBNR}	The number of auxiliary Granules required for a REC is a constant for the lifetime of a given Realm.
I _{BGVRT}	The <i>gprs</i> attribute of a REC is the set of general-purpose register values which are saved by the RMM on exit from the REC and restored by the RMM on entry to the REC.
I _{FPJDL}	The <i>mpidr</i> attribute of a REC is a value which can be used to identify the VPE associated with the REC.
I _{BLVKZ}	The <i>pc</i> attribute of a REC is the program counter which is saved by the RMM on exit from the REC and restored by the RMM on entry to the REC.
I _{GHFNQ}	The <i>runnable</i> flag of a REC determines whether the REC is eligible for execution. The RMI_REC_ENTER command results in a REC entry only if the value of the flag is RUNNABLE.
I _{SCCMH}	The runnable flag of a REC is controlled by the Realm. Its initial value is reflected in the Realm Initial Measurement, and during Realm execution its value can be changed by execution of the PSCI_CPU_ON and PSCI_CPU_OFF commands.
I _{PMYBG}	The <i>state</i> attribute of a REC is controlled by the Host, by execution of the RMI_REC_ENTER command.
D _{CDXDZ}	The <i>sysregs</i> attribute of a REC is the set of system register values which are saved by the RMM on exit from the REC and restored by the RMM on entry to the REC.
D ₀₀₀₈	The <i>gic_owner</i> attribute of a REC is the index of the Plane which is the GIC owner for the REC.

See also:

- [A2.3.3 REC index and MPIDR value](#)
- [A2.3.4 REC lifecycle](#)
- [A4.3.4.3 REC exit due to Data Abort](#)
- [A10.4 Planes interrupts](#)
- [B4.3.30 RMI_REC_ENTER command](#)

- [B6.3.2 PSCI_CPU_OFF command](#)
- [B6.3.3 PSCI_CPU_ON command](#)
- [C2.33 RmmRec type](#)

A2.3.3 REC index and MPIDR value

D_{KQVHN}

The *REC index* is the unsigned integer value generated by concatenation of MPIDR fields:

$\text{index} = \text{Aff3}:\text{Aff2}:\text{Aff1}:\text{Aff0}[3:0]$

This is illustrated by the following table.

REC index	Aff3	Aff2	Aff1	Aff0[3:0]
0	0	0	0	0
1	0	0	0	1
...
16	0	0	1	0
...
4096	0	1	0	0
...
1048576	1	0	0	0
...

I_{PVLZY}

The $\text{Aff0}[7:4]$ field of a REC MPIDR value is RES0 for compatibility with GICv3.

I_{TTWVM}

When creating the n th REC in a Realm, the Host is required to use the MPIDR corresponding to REC index n .

See also:

- [B3.62 RecIndex function](#)
- [B4.3.28 RMI_REC_CREATE command](#)
- [B4.4.38 RmiRecMpidr type](#)

A2.3.4 REC lifecycle

A2.3.4.1 States

D_{HTXQY}

The states of a REC are listed below.

State	Description
REC_READY	REC is not currently running.
REC_RUNNING	REC is currently running.

A2.3.4.2 State transitions

I_{PHMWT}

Permitted REC state transitions are shown in the following table. The rightmost column lists the events which can cause the corresponding state transition.

A transition from the pseudo-state *NULL* represents creation of a REC object. A transition to the pseudo-state *NULL* represents destruction of a REC object.

From state	To state	Events
<i>NULL</i>	REC_READY	RMI_REC_CREATE
REC_READY	<i>NULL</i>	RMI_REC_DESTROY
REC_READY	REC_RUNNING	RMI_REC_ENTER
REC_RUNNING	REC_READY	Return from RMI_REC_ENTER

I_{FNSTJ}

Permitted REC state transitions are shown in the following figure. Each arc is labeled with the events which can cause the corresponding state transition.

A transition from the pseudo-state *NULL* represents creation of a REC. A transition to the pseudo-state *NULL* represents destruction of a REC.




Figure A2.4: REC state transitions

I_{LYXCN}

The maximum number of RECs per Realm is an IMPLEMENTATION DEFINED value which is discoverable via [RMI_FEATURES](#).

See also:

- [B4.3.28 RMI_REC_CREATE command](#)
- [B4.3.29 RMI_REC_DESTROY command](#)
- [B4.3.30 RMI_REC_ENTER command](#)

See also:

- [B4.3.6 RMI_FEATURES command](#)

Chapter A3

Feature discovery and configuration

This section describes how the Host discovers features which are supported by the RMM implementation, and how the Host configures the features which are used by or available to a Realm.

A3.1 Feature discovery and configuration overview

I _{GJSMC}	RMM implementations across different CCA platforms may support disparate features and may offer disparate configuration options for Realms.
I _{YRSDX}	The features supported by an RMI implementation are discovered by reading feature pseudo-register values using the RMI_FEATURES command.
X _{WPHWG}	The term <i>pseudo-register</i> is used because, although these values are stored in memory, their usage model is similar to feature registers specified in the Arm A-profile architecture.
I _{QNJTQ}	On Realm creation, the Host provides a desired configuration in a Realm parameters structure to the RMI_REALM_CREATE command. The RMM checks that the configuration provided by the Host is supported by the implementation.
I _{RRHJJ}	Aspects of the Realm configuration which affect the security posture of the Realm are included in the Realm Initial Measurement.
I _{ZHXGX}	The features supported by an RSI implementation are discovered by reading feature pseudo-register values using the RSI_FEATURES command.

See also:

- [A2.1.6 Realm parameters](#)
- [A7.1.1 Realm Initial Measurement](#)
- [B3.32 ImplFeatures function](#)
- [B3.56 RealmParamsSupported function](#)
- [B4.3.6 RMI_FEATURES command](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [B5.3.3 RSI_FEATURES command](#)

A3.2 Realm hash algorithm

I _{WMKGX}	The set of hash algorithms supported by the implementation is reported by the RMI_FEATURES command in RmiFeatureRegister0.
I ₀₀₀₉	The hash algorithm used by a Realm is provided by the Host when calling RMI_REALM_CREATE.
R _{KPBQM}	Providing an unsupported hash algorithm causes execution of RMI_REALM_CREATE to fail.

See also:

- [A7.1 Realm measurements](#)
- [B3.56 RealmParamsSupported function](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [B4.4.16 RmiFeatureRegister0 type](#)

A3.3 Realm LPA2 and IPA width

I _{GVJMZ}	Support by the implementation for LPA2 is reported by the RMI_FEATURES command in RmiFeatureRegister0.
I _{NKLXQ}	Usage of LPA2 for Realm Translation Tables is configured by the Host when calling RMI_REALM_CREATE.
I _{LKJGN}	Realm IPA width is provided by the Host when calling RMI_REALM_CREATE.
R _{SZVDK}	Providing an unsupported IPA width (for example, smaller than the minimum supported, or larger than the maximum supported) causes execution of RMI_REALM_CREATE to fail.
I _{GKCCS}	The Host can choose a smaller IPA width than the maximum supported IPA width reported by RMI_FEATURES. This is true regardless of whether LPA2 is enabled for the Realm.

X _{FTVXQ}	<p>The Host may want to enable LPA2 for a Realm due to either or both of the following reasons:</p> <ul style="list-style-type: none"> • to allow the Realm to be configured with a larger IPA width • to allow access from mappings in the Realm's stage 2 translation to a larger PA space
I _{XDBQB}	A Realm can query its IPA width using the RSI_REALM_CONFIG command.
I _{FSNMG}	<p>If LPA2 is not enabled for a Realm then passing a PA greater than or equal to 2^{48} to any of the following commands causes an error to be returned:</p> <ul style="list-style-type: none"> • RMI_DATA_CREATE • RMI_DATA_CREATE_UNKNOWN • RMI_RTT_CREATE • RMI_RTT_AUX_CREATE • RMI_RTT_MAP_UNPROTECTED <p>See also:</p> <ul style="list-style-type: none"> • A5.2.1 Realm IPA space • B3.56 RealmParamsSupported function • B4.3.25 RMI_REALM_CREATE command • B4.4.16 RmiFeatureRegister0 type • B5.3.24 RSI_REALM_CONFIG command

A3.4 Realm support for Scalable Vector Extension

I _{KJVLJ}	Support by the implementation for the Scalable Vector Extension (FEAT_SVE) is reported by the RMI_FEATURES command in RmiFeatureRegister0.
I _{ZJSMJ}	Availability of SVE to a Realm is configured by the Host when calling RMI_REALM_CREATE.
I _{VNLNH}	SVE vector length for a Realm is provided by the Host when calling RMI_REALM_CREATE.
R _{FZZDS}	Providing a larger-than-supported SVE vector length causes execution of RMI_REALM_CREATE to fail. This is different from the behaviour of the hardware architecture, in which a larger-than-supported SVE vector length value is silently truncated.
X _{YGWTK}	The RMI ABI provides a natural mechanism to signal an invalid feature selection, via the return code of RMI_REALM_CREATE. The analog in the hardware architecture would be to generate an illegal exception return, which would cause undesirable coupling between two disparate parts of the architecture, namely the exception model and the SVE feature.
X _{CWNQC}	Providing a larger-than-supported SVE vector length causes execution of RMI_REALM_CREATE to fail prepares the architecture for addition of Realm live migration support in future. Assuming that the live migration flow starts with creation of an empty destination Realm, configured identically to the source Realm, this provides a point where the necessary feature support can be checked on the destination platform.
R _{NBYKC}	If SVE is supported by the platform but is disabled for the Realm via the RMI_REALM_CREATE command then a read of ID_AA64PFR0_EL1.SVE indicates that SVE is not supported.
U _{ZRJXL}	The RMM should trap and emulate reads of ID_AA64PFR0_EL1.SVE.
S _{VXRNN}	A Realm should discover SVE support by reading ID_AA64PFR0_EL1.SVE rather than based on the platform identity read from MIDR_EL1.

See also:

- [B3.56 RealmParamsSupported function](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [B4.4.16 RmiFeatureRegister0 type](#)

A3.5 Realm support for self-hosted debug

I _{SSTJD}	Self-hosted debug is always available in Armv8-A.
I _{LVMFG}	The number of breakpoints and watchpoints are provided by the Host when calling RMI_REALM_CREATE.
R _{CJQTB}	Providing a number of breakpoints which is larger than the number of breakpoints available causes execution of RMI_REALM_CREATE to fail.
R _{PLMDH}	Providing a number of watchpoints which is larger than the number of watchpoints available causes execution of RMI_REALM_CREATE to fail.
X _{TPBHD}	Specifying that a larger-than-supported number of breakpoints or watchpoints causes execution of RMI_REALM_CREATE to fail prepares the architecture for addition of Realm live migration support in future. Assuming that the live migration flow starts with creation of an empty destination Realm, configured identically to the source Realm, this provides a point where the necessary feature support can be checked on the destination platform.

See also:

- [B3.56 RealmParamsSupported function](#)
- [B4.3.25 RMI_REALM_CREATE command](#)

A3.6 Realm support for Performance Monitors Extension

I _{RVCQD}	Support by the implementation for the Performance Monitors Extension (FEAT_PMU) is reported by the RMI_FEATURES command in RmiFeatureRegister0.
I _{NHCFD}	Availability of PMU to a Realm is configured by the Host when calling RMI_REALM_CREATE.
I _{XZMKC}	The number of PMU counters available to a Realm is provided by the Host when calling RMI_REALM_CREATE.
R _{XVRGD}	Providing a number of PMU counters which is larger than the number of PMU counters available causes RMI_REALM_CREATE to fail.
X _{NTWKF}	Specifying that a larger-than-supported number of PMU counters causes RMI_REALM_CREATE to fail prepares the architecture for addition of Realm live migration support in future. Assuming that the live migration flow starts with creation of an empty destination Realm, configured identically to the source Realm, this provides a point where the necessary feature support can be checked on the destination platform.

See also:

- [A8.1 Realm PMU](#)
- [B3.56 RealmParamsSupported function](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [B4.4.16 RmiFeatureRegister0 type](#)

A3.7 Realm support for Activity Monitors Extension

R _{JJVZS}	The Activity Monitors Extension (FEAT_AMUv1) is not available to a Realm.
--------------------	---------------------------------------------------------------------------

A3.8 Realm support for Statistical Profiling Extension

R _{DCBNL}	The Statistical Profiling Extension (FEAT_SPE) is not available to a Realm.
--------------------	-----------------------------------------------------------------------------

A3.9 Realm support for Trace Buffer Extension

R_{NXD}XG The Trace Buffer Extension (FEAT_TRBE) is not available to a Realm.

A3.10 Support for Realm device assignment

- I₀₀₁₀ Support by the implementation for Realm device assignment is reported by the RMI_FEATURES command in RmiFeatureRegister0.
- I₀₀₁₁ Availability of Realm device assignment for a Realm is configured by the Host when calling RMI_REALM_CREATE.
See also:
- [Chapter A9 Realm device assignment](#)
 - [B3.56 RealmParamsSupported function](#)
 - [B4.3.25 RMI_REALM_CREATE command](#)
 - [B4.4.16 RmiFeatureRegister0 type](#)

A3.11 Support for auxiliary Planes

- I₀₀₁₂ The maximum number of auxiliary Planes supported by the implementation is reported by the RMI_FEATURES command in the NUM_AUX_PLANES field of RmiFeatureRegister0.
- R₀₀₁₃ The maximum number of auxiliary Planes supported by the implementation is either 0 or 3.
- I₀₀₁₄ The number of auxiliary Planes for a Realm is provided by the Host when calling RMI_REALM_CREATE.
- R₀₀₁₅ Providing a number of auxiliary Planes which is larger than the maximum number of auxiliary Planes causes RMI_REALM_CREATE to fail.
- I₀₀₁₆ For a Realm with a non-zero number of auxiliary Planes, the PLANE_RTT field of RmiFeatureRegister0 indicates which one of the following configurations is supported by the implementation:
- The Realm has an RTT tree per Plane
 - The Realm has a single RTT tree
 - The Realm can be configured to either have an RTT tree per Plane, or a single RTT tree.
- I₀₀₁₇ Whether a Realm has an RTT tree per Plane is configured by the Host when calling RMI_REALM_CREATE.
See also:
- [Chapter A10 Planes](#)
 - [B3.32 ImplFeatures function](#)
 - [B3.56 RealmParamsSupported function](#)
 - [B4.3.25 RMI_REALM_CREATE command](#)
 - [B4.4.16 RmiFeatureRegister0 type](#)

A3.12 Live Firmware Activation

- I₀₀₁₈ Live Firmware Activation (LFA) allows an update to a platform firmware component to be activated without rebooting the system. This potentially includes components which are within the TCB of a Realm.
- I₀₀₁₉ A Realm has an LFA policy which is provided by the Host when calling RMI_REALM_CREATE.
- R₀₀₂₀ If the LFA policy of a Realm is LFA_DISALLOW then all firmware components within the Realm's TCB are guaranteed not to be live activated during the lifetime of the Realm.
- I₀₀₂₁ In order to apply LFA to any firmware component (including the RMM) which is within the TCB of a Realm whose LFA policy is LFA_DISALLOW, the Host must first destroy the Realm.
- I₀₀₂₂ The mechanism via which the LFA implementation determines whether any Realm with an LFA policy of LFA_DISALLOW currently exists on the system is IMPLEMENTATION DEFINED.

I₀₀₂₃ If the LFA policy of a Realm is LFA_DISALLOW then the contents of the CCA platform software components claim reflect the state of all firmware components within the Realm's TCB, throughout the lifetime of the Realm.

I₀₀₂₄ The LFA policy of a Realm is reflected in the Realm attestation token.

See also:

- [Live Firmware Activation SMC Interface \[5\]](#)
- [A7.2.3.2.7 CCA platform software components claim](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [B4.4.21 RmiLfaPolicy type](#)

A3.13 GICv3 virtualization

I₀₀₂₅ The number of GICv3 List Registers which can be provided by the Host via the RMI_REC_ENTER command is reported by the RMI_FEATURES command in RmiFeatureRegister0.

X₀₀₂₆ Making the number of GICv3 List Registers discoverable via RMI allows the RMM to reserve List Registers for its own usage.

See also:

- [A6.1 Realm interrupts](#)
- [B4.3.30 RMI_REC_ENTER command](#)
- [B4.4.16 RmiFeatureRegister0 type](#)

A3.14 Support for Realm memory encryption

I₀₀₂₇ A Realm is configured on creation with a *MEC policy*.

A MEC policy describes whether the Realm's memory encryption context is:

- Shared with other Realms
- Private to the Realm

See also:

- [A7.2.3.1 Realm claims](#)
- [Chapter A11 Realm memory encryption](#)

Chapter A4

Realm exception model

This section describes how Realms are executed, and how exceptions which cause exit from a Realm are handled.

See also:

- [A2.1.2 Realm execution environment](#)

A4.1 Realm exception model overview

D_{HCGWL}	A <i>Realm entry</i> is a transfer of control to a Realm.
D_{RMGWJ}	A <i>Realm exit</i> is a transition of control from a Realm.
I_{SMPWB}	When executing in a Realm, an exception taken to R-EL2 or EL3 results in a Realm exit.
D_{XSNZP}	A <i>REC entry</i> is a Realm entry due to execution of RMI_REC_ENTER.
I_{FQZJG}	The Host provides the address of a REC as an input to the RMI_REC_ENTER command.
I_{MDQWG}	In this chapter, both <code>rec</code> and “the target REC” refer to the REC object which is provided to the RMI_REC_ENTER command.
D_{BLJGY}	A <i>RecRun object</i> is a data structure used to pass values between the RMM and the Host on REC entry and on REC exit.
I_{VCCFV}	A RecRun object is stored in Non-secure memory.
I_{WBHYZ}	The Host provides the address of a RecRun object as an input to the RMI_REC_ENTER command.
I_{HMSQM}	An implementation is permitted to return RMI_SUCCESS from RMI_REC_ENTER without performing a REC entry. For example, on observing a pending interrupt, the implementation can generate a REC exit due to IRQ without entering the target REC.
D_{TJCGH}	A <i>REC exit</i> is return from an execution of RMI_REC_ENTER which caused a REC entry.
I_{HPWVY}	The following diagram summarises the possible control flows that result from a Realm exit.




Figure A4.1: Realm exit paths

- The exception is taken to EL3. The Monitor handles the exception and returns control to the Realm.
- The exception is taken to EL3. The Monitor pre-empt's Realm Security state and passes control to the Secure Security state. This may be for example due to an FIQ.
- The exception is taken to EL2. The RMM decides to perform a REC exit. The RMM executes an SMC instruction, requesting the Monitor to pass control to the Non-secure Security state.
- The exception is taken to EL2. The RMM executes an SMC instruction, requesting the Monitor to perform an operation, then returns control to the Realm.
- The exception is taken to EL2. The RMM handles the exception and returns control to the Realm.

See also:

- [A4.2 REC entry](#)
- [A4.3 REC exit](#)
- [A10.2 Planes exception model](#)
- [B4.3.30 RMI_REC_ENTER command](#)
- [B4.4.40 RmiRecRun type](#)

DRAFT

A4.2 REC entry

This section describes REC entry.

See also:

- [A4.3 REC exit](#)
- [B4.3.30 RMI_REC_ENTER command](#)

A4.2.1 RmiRecEnter object

D _{SVSJM}	An <i>RmiRecEnter</i> object is a data structure used to pass values from the Host to the RMM on REC entry.
I _{YSKDN}	An <i>RmiRecEnter</i> object is stored in the <i>RecRun</i> object which is passed by the Host as an input to the <i>RMI_REC_ENTER</i> command.
I _{TRKKX}	On REC entry, execution state is restored from the REC object and from the <i>RmiRecEnter</i> object to the PE.
I _{GHDLM}	An <i>RmiRecEnter</i> object contains attributes which are used to manage Realm virtual interrupts.
D _{CLNLW}	The attributes of an <i>RmiRecEnter</i> object are summarized in the following table.

Name	Byte offset	Type	Description
flags	0x0	RmiRecEnterFlags	Flags
gprs[31]	0x200	Bits64	Registers
gicv3_hcr	0x300	Bits64	GICv3 Hypervisor Control Register value
gicv3_lrs[16]	0x308	Bits64	GICv3 List Register values

I_{ZWRQP} In this chapter, both *rec_enter* and “the *RmiRecEnter* object” refer to the *RmiRecEnter* object which is provided to the *RMI_REC_ENTER* command.

I_{LFYDV} On REC entry, all *rec_enter* fields are ignored unless specified otherwise.

See also:

- [A2.3 Realm Execution Context](#)
- [A4.3.1 RmiRecExit object](#)
- [Chapter A6 Realm interrupts and timers](#)
- [B4.4.33 RmiRecEnter type](#)

A4.2.2 General purpose registers restored on REC entry

R _{NMSFT}	On REC entry, if the most recent exit from the target REC was a REC exit due to PSCI, then all of the following occur: <ul style="list-style-type: none">• X0 to X6 contain the PSCI return code and PSCI output values.• GPR values X7 to X30 are restored from the REC object to the PE.
R _{RZRRM}	On REC entry, if either this is the first entry to this REC, or the most recent exit from the target REC was not a REC exit due to PSCI, then GPR values X0 to X30 are restored from the REC object to the PE.
R _{YSNYQ}	On REC entry, if <i>rec.pending</i> is <i>REC_PENDING_HOST_CALL</i> , then GPR values X0 to X30 are copied from <i>rec_enter.gprs[0..30]</i> to the <i>RsiHostCall</i> data structure.
R _{YWHKC}	On REC entry, if writing to the <i>RsiHostCall</i> data structure fails due to the target IPA not being mapped then a REC exit to Data Abort results.
R _{TZVNK}	On REC entry, if writing to the <i>RsiHostCall</i> data structure succeeds then <i>rec.pending</i> is <i>REC_PENDING_NONE</i> .

R_{NLVXB} On REC entry, if RMM access to `rec_enter` causes a GPF then the `RMI_REC_ENTER` command fails with `RMI_ERROR_INPUT`.

See also:

- [A4.3.3 General purpose registers saved on REC exit](#)
- [A4.3.4.3 REC exit due to Data Abort](#)
- [A4.3.7 REC exit due to PSCI](#)
- [A4.3.9 REC exit due to Host call](#)
- [A4.5 Host call](#)

A4.2.3 REC entry following REC exit due to Data Abort

R_{TWMDB} On REC entry, if `rec_enter.flags.inject_sea == RMI_INJECT_SEA` then the value of `rec_enter.flags.emul_mmio` is ignored.

R_{BWZKH} On REC entry, if the most recent exit from the target REC was a REC exit due to Emulatable Data Abort and `rec_enter.flags.emul_mmio == RMI_EMULATED_MMIO`, then the return address is the next instruction following the faulting instruction.

R_{SCJWG} On REC entry, if the most recent exit from the target REC was a REC exit due to Emulatable Data Abort and the Realm memory access was a read and `rec_enter.flags.emul_mmio == RMI_EMULATED_MMIO`, then the register indicated by `ESR_EL2.ISS.SRT` is set to `rec_enter.gprs[0]`.

I_{KNFDT} On execution of `RMI_REC_ENTER`, if the most recent exit from the target REC was not a REC exit due to Emulatable Data Abort and `rec_enter.flags.emul_mmio == RMI_EMULATED_MMIO`, then the `RMI_REC_ENTER` command fails.

R_{LJWRK} On REC entry, if the most recent exit from the target REC was a REC exit due to Data Abort at an Unprotected IPA and `rec_enter.flags.inject_sea == RMI_INJECT_SEA`, then a Synchronous External Abort is taken to the Realm.

See also:

- [A4.3.4.3 REC exit due to Data Abort](#)
- [A4.4 Emulated Data Aborts](#)
- [A5.2.6 Realm access to an Unprotected IPA](#)
- [A5.2.7 Synchronous External Aborts](#)

A4.3 REC exit

This section describes REC exit.

See also:

- [A4.2 REC entry](#)
- [B4.3.30 RMI_REC_ENTER command](#)

A4.3.1 RmiRecExit object

D _{PBDCB}	An <i>RmiRecExit object</i> is a data structure used to pass values from the RMM to the Host on REC exit.
I _{VHJTL}	An <i>RmiRecExit object</i> is stored in the <i>RecRun object</i> which is passed by the Host as an input to the <i>RMI_REC_ENTER</i> command.
I _{JKWPB}	On REC exit, execution state is saved from the PE to the REC object and to the <i>RmiRecExit object</i> .
I _{ZSCNM}	An <i>RmiRecExit object</i> contains attributes which are used to manage Realm virtual interrupts and Realm timers.
D _{FFCMN}	The attributes of an <i>RmiRecExit object</i> are summarized in the following table.

Name	Byte offset	Type	Description
exit_reason	0x0	RmiRecExitReason	Exit reason
flags	0x8	RmiRecExitFlags	Flags
esr	0x100	Bits64	Exception Syndrome Register
far	0x108	Bits64	Fault Address Register
hpfar	0x110	Bits64	Hypervisor IPA Fault Address register
rtt_tree	0x118	UInt64	Index of RTT tree active at time of the exit
rtt_level	0x120	Int64	Level of requested RTT
gprs[31]	0x200	Bits64	Registers
gicv3_hcr	0x300	Bits64	GICv3 Hypervisor Control Register value
gicv3_lrs[16]	0x308	Bits64	GICv3 List Register values
gicv3_misr	0x388	Bits64	GICv3 Maintenance Interrupt State Register value
gicv3_vmcr	0x390	Bits64	GICv3 Virtual Machine Control Register value
cntp_ctl	0x400	Bits64	Counter-timer Physical Timer Control Register value
cntp_cval	0x408	Bits64	Counter-timer Physical Timer CompareValue Register value
cntv_ctl	0x410	Bits64	Counter-timer Virtual Timer Control Register value
cntv_cval	0x418	Bits64	Counter-timer Virtual Timer CompareValue Register value
ripas_base	0x500	Bits64	Base address of target region for pending RIPAS change

Name	Byte offset	Type	Description
ripas_top	0x508	Bits64	Top address of target region for pending RIPAS change
ripas_value	0x510	RmiRipas	RIPAS value of pending RIPAS change
ripas_dev_pa	0x518	Address	Base PA of device memory region, if RIPAS change is pending due to execution of RSI_RDEV_VALIDATE_MAPPING
s2ap_base	0x520	Bits64	Base address of target region for pending S2AP change
s2ap_top	0x528	Bits64	Top address of target region for pending S2AP change
vdev_id	0x530	Bits64	Virtual device ID
imm	0x600	Bits16	Host call immediate value
plane	0x608	UInt64	Plane index
vdev	0x610	Address	VDEV which triggered REC exit due to device communication
vdev_action	0x618	RmiVdevAction	Action which triggered REC exit due to device communication
pmu_ovf_status	0x700	RmiPmuOverflowStatus	PMU overflow status

I_FQZXZ In this chapter, both `rec_exit` and “the `RmiRecExit` object” refer to the `RmiRecExit` object which is provided to the `RMI_REC_ENTER` command.

R_PNWZV On REC exit, all `rec_exit` fields are zero unless specified otherwise.

See also:

- [A2.3 Realm Execution Context](#)
- [A4.2.1 RmiRecEnter object](#)
- [A4.5 Host call](#)
- [Chapter A6 Realm interrupts and timers](#)
- [Chapter A8 Realm debug and performance monitoring](#)
- [B4.4.35 RmiRecExit type](#)

A4.3.2 Realm exit reason

I_DYWHJ On return from the `RMI_REC_ENTER` command, the reason for the REC exit is indicated by `rec_exit.exit_reason` and `rec_exit.esr`.

See also:

- [B4.4.37 RmiRecExitReason type](#)

A4.3.3 General purpose registers saved on REC exit

R_PBKVB On REC exit due to PSCI, all of the following are true:

- `rec_exit.gprs[0]` contains the PSCI FID.
- `rec_exit.gprs[1..3]` contain the corresponding PSCI arguments. If the PSCI command has fewer than 3 arguments, the remaining values contain zero.

- GPR values X7 to X30 are saved from the PE to the REC object.

R_{FNZKM} On REC exit for any reason which is not REC exit due to PSCI, GPR values X0 to X30 are saved from the PE to the REC.

R_{MZGPT} On REC exit for any reason which is neither REC exit due to Host call nor REC exit due to PSCI, `rec_exit.gprs` is zero.

R_{FRGVT} On REC exit, if RMM access to `rec_exit` causes a GPF then the RMI_REC_ENTER command fails with RMI_ERROR_INPUT.

See also:

- [A4.2.2 General purpose registers restored on REC entry](#)
- [A4.3.7 REC exit due to PSCI](#)
- [A4.3.9 REC exit due to Host call](#)

A4.3.4 REC exit due to synchronous exception

I_{SNDHF} A synchronous exception taken to R-EL2 can cause a REC exit.

I_{RPSNC} The following table summarises the behavior of synchronous exceptions taken to R-EL2.

Exception class	Behavior
Trapped WFI or WFE instruction execution	REC exit due to WFI or WFE
HVC instruction execution in AArch64 state	Unknown exception taken to Realm
SMC instruction execution in AArch64 state	One of: <ul style="list-style-type: none"> • REC exit due to PSCI • RSI command handled by RMM, followed by return to Realm
Trapped MSR, MRS or System instruction execution in AArch64 state	Emulated by RMM, followed by return to Realm
Instruction Abort from a lower Exception level	REC exit due to Instruction Abort
Data Abort from a lower Exception level	REC exit due to Data Abort

R_{YLFMD} Realm execution of an SMC which is not part of one of the following ABIs results in a return value of SMCCC_NOT_SUPPORTED:

- PSCI
- RSI

See also:

- [A4.5 Host call](#)
- [Chapter B5 Realm Services Interface](#)
- [Chapter B6 Power State Control Interface](#)

A4.3.4.1 REC exit due to WFI or WFE

D_{GLHPX} A *REC exit due to WFI or WFE* is a REC exit due to WFI, WFIT, WFE or WFET instruction execution in a Realm.

R_{VTJQF} On WFI or WFIT instruction execution in a Realm, a REC exit due to WFI or WFE is caused if `rec_enter.trap_wfi` is RMI_TRAP.

R_{GBNGW} On WFE or WFET instruction execution in a Realm, a REC exit due to WFI or WFE is caused if `rec_enter.trap_wfe` is RMI_TRAP.

R_{YQWST}

On REC exit due to WFI or WFE, all of the following are true:

- `rec_exit.exit_reason` is `RMI_EXIT_SYNC`.
- `rec_exit.esr.EC` contains the value of `ESR_EL2.EC` at the time of the Realm exit.
- `rec_exit.esr.ISS.TI` contains the value of `ESR_EL2.ISS.TI` at the time of the Realm exit.
- All other `rec_exit` fields except for `rec_exit.givc3_*`, `rec_exit_cnt*` and `rec_exit.pmu_ovf_status` are zero.

R_{BPYBC}

On REC exit due to WFI or WFE, if the exit was caused by WFET or WFIT instruction execution then `rec_exit.gpr[0]` contains the timeout value.

See also:

- [A6.1 Realm interrupts](#)
- [A6.2 Realm timers](#)
- [A8.1 Realm PMU](#)

A4.3.4.2 REC exit due to Instruction Abort

D_{GYQXK}

A *REC exit due to Instruction Abort* is a REC exit due to a Realm instruction fetch from a Protected IPA for which either of the following is true:

- HIPAS is UNASSIGNED and RIPAS is RAM
- RIPAS is DESTROYED

R_{MGWRC}

On REC exit due to Instruction Abort, all of the following are true:

- `rec_exit.exit_reason` is `RMI_EXIT_SYNC`.
- `rec_exit.esr.EC` contains the value of `ESR_EL2.EC` at the time of the Realm exit.
- `rec_exit.esr.ISS.SET` contains the value of `ESR_EL2.ISS.SET` at the time of the Realm exit.
- `rec_exit.esr.ISS.EA` contains the value of `ESR_EL2.ISS.EA` at the time of the Realm exit.
- `rec_exit.esr.ISS.IFSC` contains the value of `ESR_EL2.ISS.IFSC` at the time of the Realm exit.
- `rec_exit.hpfar` contains the value of `HPFAR_EL2` at the time of the Realm exit.
- `rec_exit.rtt_tree` contains the index of the RTT tree within which the contents of an RTTE cause the Realm to exit.
- All other `rec_exit` fields except for `rec_exit.givc3_*`, `rec_exit_cnt*` and `rec_exit.pmu_ovf_status` are zero.

I_{HMFEM}

`HPFAR_EL2.FIPA` does not include the lowest 12 bits of the faulting IPA. `rec_exit.hpfar` therefore only reveals the Realm's access patterns at a granularity of 4KB. If support was added to this specification for Granule sizes larger than 4KB, `rec_exit.hpfar` would need to be masked accordingly.

See also:

- [A5.2.2 Realm IPA state](#)
- [A5.2.3 Realm access to a Protected IPA](#)
- [A6.1 Realm interrupts](#)
- [A6.2 Realm timers](#)
- [A8.1 Realm PMU](#)

A4.3.4.3 REC exit due to Data Abort

D_{CYRMT}

A *REC exit due to Emulatable Data Abort* is a REC exit due to a Realm data access to one of the following:

- an Unprotected IPA whose HIPAS is UNASSIGNED_NS, where the access caused `ESR_EL2.ISS.ISV` to be set to '1'
- an Unprotected IPA whose HIPAS is ASSIGNED_NS, where the access caused a stage 2 permission fault and caused `ESR_EL2.ISS.ISV` to be set to '1'

D_{MTZMC}

A *REC exit due to Non-emulatable Data Abort* is a REC exit due to a Realm data access to one of the following:

- an Unprotected IPA whose HIPAS is UNASSIGNED_NS, where the access caused `ESR_EL2.ISS.ISV` to be set to '0'

- an Unprotected IPA whose HIPAS is ASSIGNED_NS, where the access caused a stage 2 permission fault and caused `ESR_EL2.ISS.ISV` to be set to '0'
- a Protected IPA whose HIPAS is UNASSIGNED and whose RIPAS is RAM
- a Protected IPA whose RIPAS is DESTROYED.

`R_RYVFL`

On REC exit due to Data Abort, all of the following are true:

- `rec_exit.exit_reason` is `RMI_EXIT_SYNC`.
- `rec_exit.esr.EC` contains the value of `ESR_EL2.EC` at the time of the Realm exit.
- `rec_exit.esr.ISS.SET` contains the value of `ESR_EL2.ISS.SET` at the time of the Realm exit.
- `rec_exit.esr.ISS.FnV` contains the value of `ESR_EL2.ISS.FnV` at the time of the Realm exit.
- `rec_exit.esr.ISS.EA` contains the value of `ESR_EL2.ISS.EA` at the time of the Realm exit.
- `rec_exit.esr.ISS.DFSC` contains the value of `ESR_EL2.ISS.DFSC` at the time of the Realm exit.
- `rec_exit.hpfar` contains the value of `HPFAR_EL2` at the time of the Realm exit.
- `rec_exit.rtt_tree` contains the index of the RTT tree within which the contents of an RTTE cause the Realm to exit.

On REC exit due to Emulatable Data Abort, all of the following are true:

- `rec.emulatable_abort` is `EMULATABLE_ABORT`.
- `rec_exit.esr.ISS.ISV` contains the value of `ESR_EL2.ISS.ISV` at the time of the Realm exit.
- `rec_exit.esr.ISS.SAS` contains the value of `ESR_EL2.ISS.SAS` at the time of the Realm exit.
- `rec_exit.esr.ISS.SF` contains the value of `ESR_EL2.ISS.SF` at the time of the Realm exit.
- `rec_exit.esr.ISS.WnR` contains the value of `ESR_EL2.ISS.WnR` at the time of the Realm exit.
- `rec_exit.far` contains the value of `FAR_EL2` at the time of the Realm exit, with bits more significant than the size of a Granule masked to zero.

On REC exit due to Non-emulatable Data Abort at an Unprotected IPA, all of the following are true:

- `rec_exit.esr.IL` contains the value of `ESR_EL2.IL` at the time of the Realm exit.

On REC exit due to Data Abort, all other `rec_exit` fields except for `rec_exit.givc3_*`, `rec_exit.cnt*` and `rec_exit.pmu_ovf_status` are zero.

`X_XHXJC`

On REC exit due to Emulatable Data Abort, `ESR_EL2.ISS.SSE` is not propagated to the Host. This is because this field is used to emulate sign extension on loads, which must be performed by the RMM so that the Realm can rely on architecturally correct behavior of the virtual execution environment.

`X_HSWFR`

On REC exit due to Emulatable Data Abort, the Host can calculate the faulting IPA from the `rec_exit.hpfar` and `rec_exit.far` values.

`I_WCYNY`

`HPFAR_EL2.FIPA` does not include the lowest 12 bits of the faulting IPA. `rec_exit.hpfar` therefore only reveals the Realm's access patterns at a granularity of 4KB. If support was added to this specification for Granule sizes larger than 4KB, `rec_exit.hpfar` would need to be masked accordingly.

`R_FFNHW`

On REC exit due to Emulatable Data Abort, if the Realm memory access was a write, `rec_exit.gprs[0]` contains the value of the register indicated by `ESR_EL2.ISS.SRT` at the time of the Realm exit.

`R_QBTFR`

On REC exit not due to Emulatable Data Abort, `rec.emulatable_abort` is `NOT_EMULATABLE_ABORT`.

See also:

- [A4.2.3 REC entry following REC exit due to Data Abort](#)
- [A4.4 Emulated Data Aborts](#)
- [A5.2.1 Realm IPA space](#)
- [A5.2.3 Realm access to a Protected IPA](#)
- [A5.2.6 Realm access to an Unprotected IPA](#)
- [A6.1 Realm interrupts](#)
- [A6.2 Realm timers](#)
- [A8.1 Realm PMU](#)

A4.3.5 REC exit due to IRQ

D_{YLWXX} A *REC exit due to IRQ* is a REC exit due to an IRQ exception which should be handled by the Host.

R_{TYJSX} On REC exit due to IRQ, `rec_exit.exit_reason` is `RMI_EXIT_IRQ`.

R_{CSQXV} On REC exit due to IRQ, `rec_exit.esr` is zero.

See also:

- [Chapter A6 Realm interrupts and timers](#)

A4.3.6 REC exit due to FIQ

D_{ZTYMM} A *REC exit due to FIQ* is a REC exit due to an FIQ exception which should be handled by the Host.

R_{PDSBD} On REC exit due to FIQ, `rec_exit.exit_reason` is `RMI_EXIT_FIQ`.

R_{GXZRF} On REC exit due to FIQ, `rec_exit.esr` is zero.

See also:

- [Chapter A6 Realm interrupts and timers](#)

A4.3.7 REC exit due to PSCI

I_{ZSGFP} A PSCI function executed by a Realm is either:

- handled by the RMM, returning to the Realm, or
- forwarded by the RMM to the Host via a *REC exit due to PSCI*.

D_{RFTQD} A *REC exit due to PSCI* is a REC exit due to Realm PSCI function execution by SMC instruction which was forwarded by the RMM to the Host.

I_{VBJXY} The following table summarises the behavior of PSCI function execution by a Realm.

PSCI functions not listed in this table are not supported. Calling a non-supported PSCI function results in a return value of `PSCI_NOT_SUPPORTED`.

PSCI function	Can result in REC exit due to PSCI	Requires Host to call <code>RMI_PSCI_COMPLETE</code>
PSCI_VERSION	No	-
PSCI_FEATURES	No	-
PSCI_CPU_SUSPEND	Yes	No
PSCI_CPU_OFF	Yes	No
PSCI_CPU_ON	Yes	Yes
PSCI_AFFINITY_INFO	Yes	Yes
PSCI_SYSTEM_OFF	Yes	No
PSCI_SYSTEM_RESET	Yes	No

R_{NTZNJ} On REC exit due to PSCI, `rec_exit.exit_reason` is `RMI_EXIT_PSCI`.

R_{SXGJK} On REC exit due to PSCI, `rec_exit.gpr`s contains sanitised parameters from the PSCI call.

R_{YTDGT} On REC exit due to PSCI, if the command arguments include an MPIDR value, `rec.pending` is set to `REC_PENDING_PSCI`. Otherwise, `rec.pending` is set to `REC_PENDING_NONE`.

I_KKFMQ

Following REC exit due to PSCI, if `rec.pending` is `REC_PENDING_PSCI`, the Host must complete the request by calling the `RMI_PSCI_COMPLETE` command, prior to re-entering the REC.

In the call to `RMI_PSCI_COMPLETE`, the Host provides the target REC, which corresponds to the MPIDR value provided by the Realm. This is necessary because the RMM does not maintain a mapping from MPIDR values to REC addresses. The RMM validates that the REC provided by the Host matches the MPIDR value.

In the call to `RMI_PSCI_COMPLETE`, the Host provides a PSCI status value, which the RMM handles as follows:

- If the Host provides `PSCI_SUCCESS`, the RMM performs the PSCI operation requested by the Realm. The result of the PSCI operation is recorded in the REC and returned to the Realm on the next entry to the calling REC.
- If the Host provides a status value other than `PSCI_SUCCESS`, the RMM validates that the status code is permitted for the PSCI operation requested by the Realm. If the status code is permitted, it is recorded in the REC and returned to the Realm on the next entry to the calling REC.

See also:

- [A4.3.3 General purpose registers saved on REC exit](#)
- [B3.47 PsciReturnCodePermitted function](#)
- [B4.3.23 RMI_PSCI_COMPLETE command](#)
- [Chapter B6 Power State Control Interface](#)
- [D1.4 PSCI flows](#)

A4.3.8 REC exit due to RIPAS change pending

D_JGCVY

A REC exit due to RIPAS change pending is a REC exit due to the Realm issuing a RIPAS change request.

R_QSSKK

On REC exit due to RIPAS change pending, all of the following are true:

- `rec_exit.exit_reason` is `RMI_EXIT_RIPAS_CHANGE`.
- `rec_exit.ripas_base` is the base address of the region on which a RIPAS change is pending.
- `rec_exit.ripas_top` is the top address of the region on which a RIPAS change is pending.
- `rec_exit.ripas_value` is the requested RIPAS value.
- `rec.ripas_addr` is the base address of the region on which a RIPAS change is pending.
- `rec.ripas_top` is the top address of the region on which a RIPAS change is pending.
- `rec.ripas_value` is the requested RIPAS value.

On REC exit due to RIPAS change pending, if the exit was triggered by `RSI_RDEV_VALIDATE_MAPPING` then all of the following are true:

- `rec_exit.ripas_dev_pa` is the base physical address of the device memory region.
- `rec_exit.flags.ripas_dev_shared` is the shared state of the device memory region.
- `rec.ripas_dev_pa` is the base physical address of the device memory region.
- `rec.ripas_dev_shared` is the shared state of the device memory region.

I_MCKKH

On REC exit due to RIPAS change pending:

- `rec_exit` holds the base address and the size of the region on which a RIPAS change is pending. These values inform the Host of the bounds of the RIPAS change request.
- `rec` holds the next address to be processed in a RIPAS change, and the top of the requested RIPAS change region. These values are used by the RMM to enforce that the `RMI_RTT_SET_RIPAS` command can only apply RIPAS change within the bounds of the RIPAS change request, and to report the progress of the RIPAS change to the Realm on the next REC entry.

On REC exit due to RIPAS change pending, if the exit was triggered by `RSI_RDEV_VALIDATE_MAPPING` then:

- `rec_exit` holds the base physical address of the device memory region and the shared state of the device memory region. These values allow the Host to create device memory mappings (by calling `RMI_DEV_MEM_MAP`) on demand.

R_{QRMN}

On REC exit not due to RIPAS change pending, all of the following are true:

- `rec.ripas_addr` is 0
- `rec.ripas_top` is 0

See also:

- [A2.3.2 REC attributes](#)
- [A5.4 RIPAS change](#)
- [A9.5.3 Realm validation of device memory mappings](#)

A4.3.9 REC exit due to Host call

D_{WFZK}

A REC exit due to Host call is a REC exit due to RSI_HOST_CALL execution in a Realm.

R_{GTJRP}

On REC exit due to Host call, all of the following are true:

- `rec.pending` is `REC_PENDING_HOST_CALL`.
- `rec_exit.exit_reason` is `RMI_EXIT_HOST_CALL`.
- `rec_exit.imm` contains the immediate value passed to the RSI_HOST_CALL command.
- `rec_exit.plane` contains the index of the Plane which executed the RSI_HOST_CALL command.
- `rec_exit.gprs[0..30]` contain the register values passed to the RSI_HOST_CALL command.
- All other `rec_exit` fields except for `rec_exit.givc3_*`, `rec_exit_cnt*` and `rec_exit.pmu_ovf_status` are zero.

See also:

- [A4.5 Host call](#)
- [A6.1 Realm interrupts](#)
- [A6.2 Realm timers](#)
- [A8.1 Realm PMU](#)
- [B5.3.4 RSI_HOST_CALL command](#)

A4.3.10 REC exit due to SError

D_{PGMHP}

A REC exit due to SError is a REC exit due to an SError interrupt during Realm execution.

R_{LRCFP}

On REC exit due to SError, all of the following occur:

- `rec_exit.exit_reason` is `RMI_EXIT_SERROR`.
- `rec_exit.esr.EC` contains the value of `ESR_EL2.EC` at the time of the Realm exit.
- `rec_exit.esr.ISS.IDS` contains the value of `ESR_EL2.ISS.IDS` at the time of the Realm exit.
- `rec_exit.esr.ISS.AET` contains the value of `ESR_EL2.ISS.AET` at the time of the Realm exit.
- `rec_exit.esr.ISS.EA` contains the value of `ESR_EL2.ISS.EA` at the time of the Realm exit.
- `rec_exit.esr.ISS.DFSC` contains the value of `ESR_EL2.ISS.DFSC` at the time of the Realm exit.
- All other `rec_exit` fields except for `rec_exit.givc3_*`, `rec_exit_cnt*` and `rec_exit.pmu_ovf_status` are zero.

See also:

- [A6.1 Realm interrupts](#)
- [A6.2 Realm timers](#)
- [A8.1 Realm PMU](#)

A4.3.11 REC exit due to device communication

D₀₀₂₈

A REC exit due to device communication is a REC exit due to RSI_RDEV_CONTINUE execution in a Realm.

R₀₀₂₉

On REC exit due to device communication, `rec_exit.vdev` identifies the VDEV which triggered the REC exit.

R₀₀₃₀ On REC exit due to device communication, the state of the VDEV which triggered the REC exit becomes VDEV_COMMUNICATING.

R₀₀₃₁ On REC exit due to device communication, the communication status of the VDEV which triggered the REC exit becomes DEV_COMM_PENDING.

See also:

- [A9.2 Communication between RMM and a device](#)
- [A9.5 Realm management of an assigned device interface](#)
- [B5.3.15 RSI_RDEV_CONTINUE command](#)

A4.3.12 REC exit due to RTT request

D₀₀₃₂ A REC exit due to RTT request is a REC exit due to the RMM requiring an RTT to be created in order to proceed with an operation.

R₀₀₃₃ On REC exit due to RTT request, `rec_exit.rtt_tree` contains the index of the RTT tree within which the contents of an RTTE cause the Realm to exit.

R₀₀₃₄ On REC exit due to RTT request, `rec_exit.rtt_level` identifies the level of the requested RTT.

See also:

- [A10.3.2.4 Stage 2 access permissions change](#)
- [B4.3.31 RMI_RTT_AUX_CREATE command](#)
- [B4.3.38 RMI_RTT_CREATE command](#)

A4.3.13 REC exit due to S2AP change pending

D₀₀₃₅ A REC exit due to S2AP change pending is a REC exit due to the Realm issuing an S2AP change request.

R₀₀₃₆ On REC exit due to S2AP change pending, all of the following are true:

- `rec_exit.exit_reason` is RMI_EXIT_S2AP_CHANGE.
- `rec_exit.s2ap_base` is the base address of the region on which an S2AP change is pending.
- `rec_exit.s2ap_top` is the top address of the region on which an S2AP change is pending.
- `rec.s2ap_addr` is the base address of the region on which an S2AP change is pending.
- `rec.s2ap_top` is the top address of the region on which an S2AP change is pending.
- `rec.s2ap_value` is the requested S2AP value.

I₀₀₃₇ On REC exit due to RIPAS change pending:

- `rec_exit` holds the base address and the size of the region on which an S2AP change is pending. These values inform the Host of the bounds of the RIPAS change request.
- `rec` holds the next address to be processed in an S2AP change, and the top of the requested S2AP change region. These values are used by the RMM to enforce that the RMI_RTT_SET_S2AP command can only apply S2AP change within the bounds of the S2AP change request, and to report the progress of the S2AP change to the Realm on the next REC entry.

R₀₀₃₈ On REC exit not due to S2AP change pending, all of the following are true:

- `rec.s2ap_addr` is 0
- `rec.s2ap_top` is 0

See also:

- [A2.3.2 REC attributes](#)
- [A10.3.2.4 Stage 2 access permissions change](#)

A4.3.14 REC exit due to VDEV request

- D₀₀₃₉ A *REC exit due to VDEV request* is a REC exit due to the RMM requiring the Host to provide the VDEV object which matches a specified virtual device ID.
- R₀₀₄₀ On REC exit due to VDEV request, `rec_exit.exit_reason` is `RMI_EXIT_VDEV_REQUEST`.
- R₀₀₄₁ On REC exit due to VDEV request, `rec_exit.vdev_id` contains the requested virtual device ID.
- R₀₀₄₂ On REC exit due to VDEV request, `rec.vdev_pending` is set to `REC_PENDING_VDEV_REQUEST`.
- I₀₀₄₃ Following REC exit due to VDEV request, the Host must complete the request by calling the `RMI_VDEV_COMPLETE` command, prior to re-entering the REC.

In the call to `RMI_VDEV_COMPLETE`, the Host provides the target VDEV, which corresponds to the virtual device ID value provided by the Realm. This is necessary because the RMM does not maintain a mapping from virtual device IDs to VDEV objects. The RMM validates that the VDEV provided by the Host matches the virtual device ID value.

See also:

- [A9.2.2 Mapping from virtual device ID to VDEV object](#)
- [B4.3.50 RMI_VDEV_COMPLETE command](#)

DRAFT

A4.4 Emulated Data Aborts

I_{SVYDC}

On REC exit due to Emulatable Data Abort, sufficient information is provided to the Host to enable it to emulate the access, for example to emulate a virtual peripheral.

On taking the REC exit, the Host can either

- Establish a mapping in the RTT, in which case it would want the Realm to re-attempt the access. In this case, on the next REC entry the Host sets `enter.flags.emul_mmio = RMI_NOT_EMULATED_MMIO`, which indicates that instruction emulation was not performed. This causes the return address to be the faulting instruction.
- Emulate the access. For an emulated write, the data is provided in `exit.gprs[0]`. For an emulated read, the data is provided in `enter.gprs[0]`. In this case, on the next REC entry the Host sets `enter.flags.emul_mmio = RMI_EMULATED_MMIO`, which indicates that the instruction was emulated. This causes the return address to be the address of the instruction which generated the Data Abort plus 4 bytes.

See also:

- [A4.2.3 REC entry following REC exit due to Data Abort](#)
- [A4.3.4.3 REC exit due to Data Abort](#)
- [A5.2.1 Realm IPA space](#)

A4.5 Host call

This section describes the programming model for Realm communication with the Host.

D_{YDJWT}

A *Host call* is a call made by the Realm to the Host, by execution of the `RSI_HOST_CALL` command.

I_{XNFKZ}

A Host call can be used by a Realm to make a hypercall.

R_{DNBQF}

On Realm execution of HVC, an Unknown exception is taken to the Realm.

See also:

- [A4.2.2 General purpose registers restored on REC entry](#)
- [A4.3.9 REC exit due to Host call](#)
- [B5.3.4 RSI_HOST_CALL command](#)
- [D1.3.2 Host call flow](#)

Chapter A5

Realm memory management

This section describes how Realm memory is managed. This includes:

- How the translation tables which describe the Realm's address space are managed by the Host.
- Properties of the Realm's address space, and of the memory which can be mapped into it.
- How faults caused by Realm memory accesses are handled.

See also:

- [A2.1.2 Realm execution environment](#)
- [D1.5 Realm memory management flows](#)
- [Chapter D2 Realm shared memory protocol](#)

A5.1 Realm memory management overview

Realm memory management can be viewed from one of two standpoints: the Realm and the Host.

From the Realm's point of view, the RMM provides security guarantees regarding the IPA space of the Realm and the memory which is mapped into it. These security guarantees are upheld via RSI commands which the Realm can execute in order to query the initial configuration and contents of its address space, and to modify properties of the address space at runtime.

From the Host's point of view, Realm memory management involves manipulating the stage 2 translation tables which describe the Realm's address space, and handling faults which are caused by Realm memory accesses. These operations are similar to those involved in managing the memory of a normal VM, but in the case of a Realm they are performed via execution of RMI commands.

See also:

- [A5.2 Realm view of memory management](#)
- [A5.3 Host view of memory management](#)

A5.2 Realm view of memory management

This section describes memory management from the Realm's point of view.

A5.2.1 Realm IPA space

I_{DLRZF}

The IPA space of a Realm is divided into two halves: Protected IPA space and Unprotected IPA space.

S_{LZHXC}

Software in a Realm should treat the most significant bit of an IPA as a protection attribute.

D_{KXGDV}

A *Protected IPA* is an address in the lower half of a Realm's IPA space. The most significant bit of a Protected IPA is 0.

D_{MRWGM}

An *Unprotected IPA* is an address in the upper half of a Realm's IPA space. The most significant bit of an Unprotected IPA is 1.

See also:

- [A2.1.3 Realm attributes](#)
- [A3.3 Realm LPA2 and IPA width](#)

A5.2.2 Realm IPA state

D_{WWCBD}

A Protected IPA has an associated *Realm IPA state* (RIPAS).

The RIPAS values are shown in the following table.

Name	Description
DESTROYED	Address which is inaccessible to the Realm due to an action taken by the Host.
DEV	Address where memory of an assigned Realm device is mapped.
EMPTY	Address where no Realm resources are mapped.
RAM	Address where private code or data owned by the Realm is mapped.

I_{VZCZV}

RIPAS values are stored in an RTT.

I_ZPNZT The Realm can query the RIPAS of an IPA range by executing RSI_IPA_STATE_GET.

See also:

- [A5.5 Realm Translation Table](#)
- [Chapter A9 Realm device assignment](#)
- [B5.3.5 RSI_IPA_STATE_GET command](#)

A5.2.3 Realm access to a Protected IPA

R_JVQQR Realm data access to a Protected IPA whose RIPAS is EMPTY causes a Synchronous External Abort taken to the Realm.

R_MKLSQ Realm instruction fetch from a Protected IPA whose RIPAS is EMPTY causes a Synchronous External Abort taken to the Realm.

R_QSQLF Realm data access to a Protected IPA whose RIPAS is RAM does not cause a Synchronous External Abort taken to the Realm.

I_PGHBZ Realm data access to a Protected IPA whose RIPAS is RAM can cause an REC exit due to Data Abort.

R_FCJCP Realm instruction fetch from a Protected IPA whose RIPAS is RAM does not cause a Synchronous External Abort taken to the Realm.

I_XHKQY Realm instruction fetch from a Protected IPA whose RIPAS is RAM can cause a REC exit due to Instruction Abort.

R_CLVKF Realm data access to a Protected IPA whose RIPAS is DESTROYED causes a REC exit due to Data Abort.

R_MZYQT Realm instruction fetch from a Protected IPA whose RIPAS is DESTROYED causes a REC exit due to Instruction Abort.

R_0044 Realm data access to a Protected IPA whose RIPAS is DEV does not cause a Synchronous External Abort taken to the Realm.

I_0045 Realm data access to a Protected IPA whose RIPAS is DEV can cause an REC exit due to Data Abort.

R_0046 Realm instruction fetch from a Protected IPA whose RIPAS is DEV causes a Synchronous External Abort taken to the Realm.

See also:

- [A4.3.4.2 REC exit due to Instruction Abort](#)
- [A4.3.4.3 REC exit due to Data Abort](#)
- [A5.2.7 Synchronous External Aborts](#)
- [A9.6 Device access to a Protected IPA](#)

A5.2.4 Changes to RIPAS while Realm state is REALM_NEW

This section describes how the RIPAS of a Protected IPA can change while the Realm state is REALM_NEW.

I_BSBHN For a Realm in the REALM_NEW state, the RIPAS of a Protected IPA can change to RAM due to Host execution of RMI_DATA_CREATE or RMI_RTT_INIT_RIPAS.

I_BSGSW For a Realm in the REALM_NEW state, changing the RIPAS of a Protected IPA to RAM causes the RIM to be updated.

I_YCPNY For a Realm in the REALM_NEW state, the RIPAS of a Protected IPA can change to DESTROYED due to Host execution of RMI_DATA_DESTROY or RMI_RTT_DESTROY.

I_YXLCP For a Realm in the REALM_NEW state, changing the RIPAS of a Protected IPA to DESTROYED does not cause the RIM to be updated.

See also:

- [A5.4 RIPAS change](#)

- [A7.1.1 Realm Initial Measurement](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)
- [B4.3.41 RMI_RTT_INIT_RIPAS command](#)

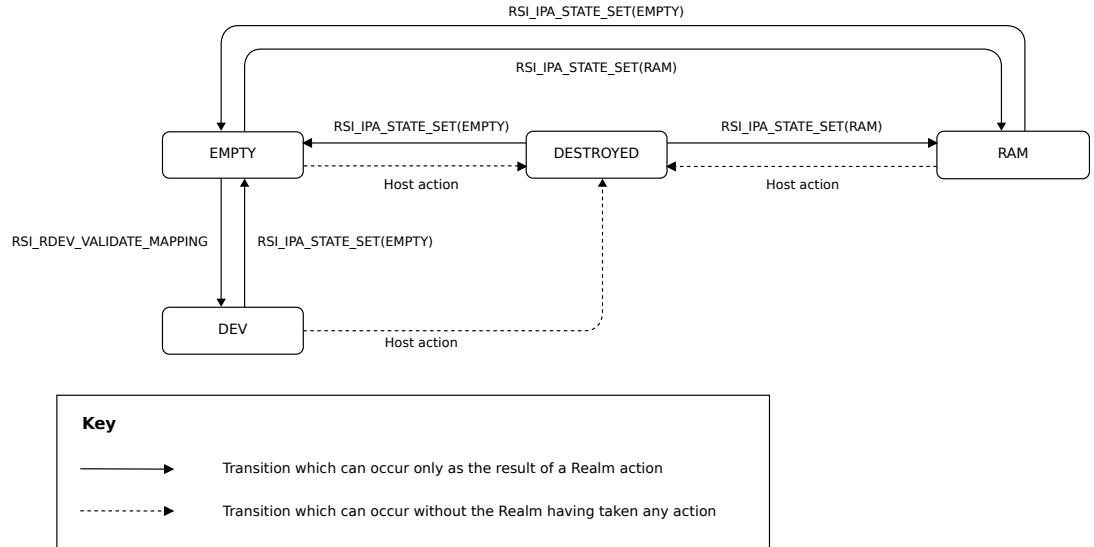
A5.2.5 Changes to RIPAS while Realm state is REALM_ACTIVE

This section describes how the RIPAS of a Protected IPA can change while the Realm state is REALM_ACTIVE.

I _{NZXP} G	A Realm in the REALM_ACTIVE state can request the RIPAS of a region of Protected IPA space to be changed to EMPTY, RAM or DEV.
I _{RHXF}	A Realm in the REALM_ACTIVE state cannot request the RIPAS of a region of Protected IPA space to be changed to DESTROYED.
I _{FRJJH}	For a Realm in the REALM_ACTIVE state, the RIPAS of a Protected IPA can change to EMPTY only in response to Realm execution of RSI_IPA_STATE_SET.
X _{HQLVY}	The fact that the Host cannot change the RIPAS of a Protected IPA to EMPTY without the Realm having consented to this change prevents the Host from injecting an SEA at a Protected IPA which has been configured to have a RIPAS of RAM, which could potentially trigger unexpected behavior in the Realm.
I _{HNFYR}	For a Realm in the REALM_ACTIVE state, the RIPAS of a Protected IPA can change to RAM only in response to Realm execution of RSI_IPA_STATE_SET.
I _{VVFMX}	On execution of RSI_IPA_STATE_SET, a Realm can optionally specify that the RIPAS change should only succeed if the current RIPAS is not DESTROYED.
X _{VXHBV}	<p>An expected pattern for Realm creation is as follows:</p> <ol style="list-style-type: none"> 1. Host populates an “initial image” range of Realm IPA space with measured content: Host executes RMI_DATA_CREATE, establishing a mapping to physical memory, changing RIPAS to RAM and updating the RIM. 2. Host informs the Realm of the range of IPA space which should be considered by the Realm as DRAM. This is a superset of the IPA range populated in step 1. For unpopulated parts of this IPA range, the RIPAS is EMPTY. 3. Realm executes RSI_IPA_STATE_SET(ripas=RAM) for the DRAM IPA range described to it in step 2. Following this command, the desired state is: <ol style="list-style-type: none"> a. For the initial image IPA range, the contents match those described by the RIM. b. For the entire DRAM IPA range, RIPAS is RAM. <p>If at step 2, the Host were to execute RMI_DATA_DESTROY on a page within the initial image IPA range, its RIPAS would change to DESTROYED. The Host could then execute RMI_DATA_CREATE_UNKNOWN, with the result that contents of the initial image IPA range no longer match those described by the RIM.</p> <p>By specifying at step 3 that the RIPAS change should only succeed if the current RIPAS is not DESTROYED, the Realm is able to prevent loss of integrity within the initial image IPA range.</p>
I ₀₀₄₇	For a Realm in the REALM_ACTIVE state, the RIPAS of a Protected IPA can change to DEV only in response to Realm execution of RSI_RDEV_VALIDATE_MAPPING.
I _{KZVDC}	For a Realm in the REALM_ACTIVE state, the RIPAS of a Protected IPA can change to DESTROYED due to Host execution of RMI_DATA_DESTROY or RMI_RTT_DESTROY.
X _{JJPHJ}	The result of changing the RIPAS of a Protected IPA to DESTROYED is that subsequent Realm accesses to that address do not make forward progress. This is consistent with the principle that the RMM does not provide an availability guarantee to a Realm.

I_{NMMSG}

The following diagram summarizes RIPAS changes which can occur when the Realm state is REALM_ACTIVE.



See also:

- [A5.4 RIPAS change](#)
- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)
- [B4.3.41 RMI_RTT_INIT_RIPAS command](#)
- [B5.3.6 RSI_IPA_STATE_SET command](#)
- [B5.3.23 RSI_RDEV_VALIDATE_MAPPING command](#)

A5.2.6 Realm access to an Unprotected IPA

I_{KQJML}

An access by a Realm to an Unprotected IPA can result in a *Granule Protection Fault* (GPF).

The RMM does not ensure that the GPT entry of a Granule mapped at an Unprotected IPA permits access via Non-secure PAS.

S_{ZZBQF}

Realm software must be able to handle taking a GPF during access to an Unprotected IPA.

I_{WCVBZ}

Realm data access to an Unprotected IPA can cause a REC exit due to Data Abort.

I_{RNDTJ}

On taking a REC exit due to Data Abort at an Unprotected IPA, the Host can inject a Synchronous External Abort to the Realm.

X_{MGBDH}

The Host can inject an SEA in response to an unexpected Realm data access to an Unprotected IPA.

I_{FVYCM}

Realm data access to an Unprotected IPA which caused `ESR_EL2.ISS.ISV` to be set to '1' can be emulated by the Host.

R_{XLSKP}

Realm instruction fetch from an Unprotected IPA causes a Synchronous External Abort taken to the Realm.

See also:

- [A4.2.3 REC entry following REC exit due to Data Abort](#)
- [A4.3.4.3 REC exit due to Data Abort](#)
- [A4.4 Emulated Data Aborts](#)
- [A5.2.7 Synchronous External Aborts](#)

A5.2.7 Synchronous External Aborts

R_{VKNJW} When a Synchronous External Abort is taken to a Realm, $ESR_{EL1}.EA == '1'$.

A5.2.8 Realm access outside IPA space

R_{GYVZQ} If stage 1 translation is enabled, Realm access to an IPA which is greater than the IPA space of the Realm causes a stage 1 Address Size Fault taken to the Realm, with the fault status code indicating the level at which the fault occurred.

R_{LSJJR} If stage 1 translation is disabled, Realm access to an IPA which is greater than the IPA space of the Realm causes a stage 1 level 0 Address Size Fault taken to the Realm.

DRAFT

A5.2.9 Summary of Realm IPA space properties

TPGKW

The following table summarizes the properties of Realm IPA space.

Realm IPA	Data access causes abort to Realm?	Data access causes REC exit due to Data Abort?	Instruction fetch causes abort to Realm?	Instruction fetch causes REC exit due to Instruction Abort?
Protected, RIPAS=EMPTY	Always (SEA)	Never	Always (SEA)	Never
Protected, RIPAS=RAM	Never	When HIPAS=UNASSIGNED	Never	When HIPAS=UNASSIGNED
Protected, RIPAS=DEV	Never	When HIPAS=UNASSIGNED	Always (SEA)	Never
Protected, RIPAS=DESTROYED	Never	Always	Never	Always
Unprotected	Host can inject SEA following REC exit due to Data Abort	When HIPAS=UNASSIGNED_NS	Always (SEA)	Never
Outside Realm IPA space	Always (Address Size Fault)	Never	Always (Address Size Fault)	Never

See also:

- [A4.2.3 REC entry following REC exit due to Data Abort](#)

A5.2.10 Cache maintenance operations

R_{TZQDY}

A data cache invalidate by set / way instruction executed by a Realm either has no effect, or performs a data cache clean and invalidate.

X_{XZRDW}

This is to ensure that a Realm cannot invalidate a cache line owned by another Realm.

U_{VQMTB}

Arm expects that the RMM will set HCR_EL2.VM == '1', which causes a data cache invalidate instruction executed at EL1 to perform a data cache clean and invalidate.

A5.3 Host view of memory management

This section describes memory management from the Host's point of view.

A5.3.1 Host IPA state

D_YZT_ZJ A Realm IPA has an associated *Host IPA state* (HIPAS).

The HIPAS values are shown in the following table.

Name	Description
HIPAS_ASSIGNED	Protected IPA which is associated with a DATA Granule.
HIPAS_ASSIGNED_DEV_PRIVATE	Protected IPA which is associated with a DEV_PRIVATE Granule.
HIPAS_ASSIGNED_DEV_SHARED	Protected IPA which is associated with a DEV_SHARED Granule.
HIPAS_ASSIGNED_NS	Unprotected IPA which is associated with an NS Granule.
HIPAS_UNASSIGNED	Protected IPA which is not associated with any Granule.
HIPAS_UNASSIGNED_NS	Unprotected IPA which is not associated with any Granule.

I_TRSK_J HIPAS values are stored in a Realm Translation Table (RTT).

I_GZMK_Q HIPAS transitions are caused by execution of RMI commands.

I_NQCG_S A mapping at a Protected IPA is valid if the HIPAS is ASSIGNED and the RIPAS is RAM.

I₀₀₄₈ A mapping at a Protected IPA is valid if the HIPAS is ASSIGNED_DEV_PRIVATE or ASSIGNED_DEV_SHARED and the RIPAS is DEV.

I_{YMNSR}

The following table summarizes, for each combination of RIPAS and HIPAS for a Protected IPA:

- the translation table entry attributes, and
- the behavior which results from Realm access to that IPA.

Each TTD.X column refers to the value of the corresponding “X” field in the architecturally-defined Stage 2 translation table descriptor which is written by the RMM.

RIPAS	HIPAS	TTD.ADDR	TTD.NS	TTD.VALID	Data access	Instruction fetch
EMPTY	UNASSIGNED			0	SEA to Realm	SEA to Realm
EMPTY	ASSIGNED	DATA		0	SEA to Realm	SEA to Realm
EMPTY	ASSIGNED_DEV_*	DEV		0	SEA to Realm	SEA to Realm
RAM	UNASSIGNED			0	REC exit due to Data Abort	REC exit due to Instruction Abort
RAM	ASSIGNED	DATA	0	1	Data access	Instruction fetch
RAM	ASSIGNED_DEV_*	DEV		0	REC exit due to Data Abort	REC exit due to Instruction Abort
DESTROYED	UNASSIGNED			0	REC exit due to Data Abort	REC exit due to Instruction Abort
DESTROYED	ASSIGNED	DATA		0	REC exit due to Data Abort	REC exit due to Instruction Abort
DESTROYED	ASSIGNED_DEV_*	DEV		0	REC exit due to Data Abort	REC exit due to Instruction Abort
DEV	UNASSIGNED			0	REC exit due to Data Abort	SEA to Realm
DEV	ASSIGNED	DATA		0	REC exit due to Data Abort	SEA to Realm
DEV	ASSIGNED_DEV_*	DEV		1	Device access	SEA to Realm

See also:

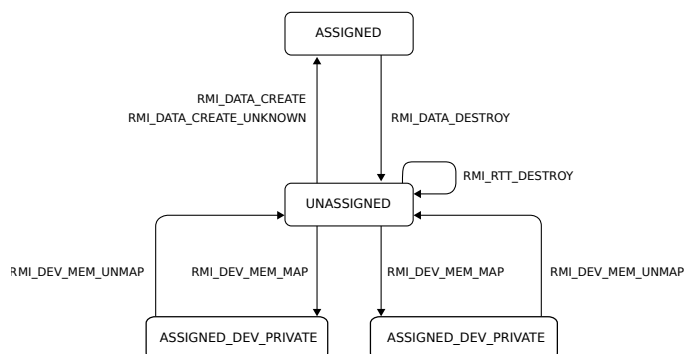
- [A5.5 Realm Translation Table](#)
- [Chapter A9 Realm device assignment](#)

A5.3.2 Changes to HIPAS while Realm state is REALM_NEW

This section describes how the HIPAS of a Protected IPA can change while the Realm state is REALM_NEW.

I_{YNFGD}

The following diagram summarizes HIPAS changes at a Protected IPA which can occur when the Realm state is REALM_NEW.



See also:

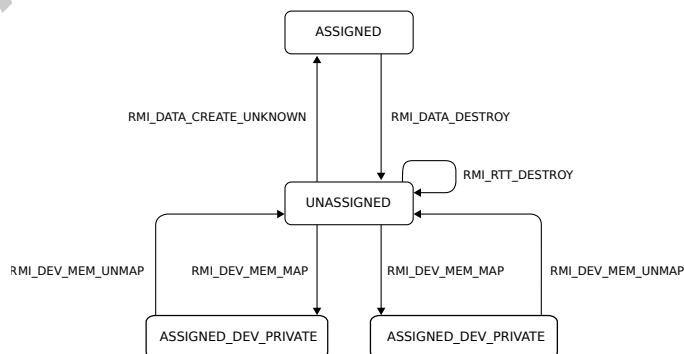
- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.4 RMI_DEV_MEM_MAP command](#)
- [B4.3.5 RMI_DEV_MEM_UNMAP command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)

A5.3.3 Changes to HIPAS while Realm state is REALM_ACTIVE

This section describes how the HIPAS of a Protected IPA can change while the Realm state is REALM_ACTIVE.

I_{WKZXY}

The following diagram summarizes HIPAS changes at a Protected IPA which can occur when the Realm state is REALM_ACTIVE.



See also:

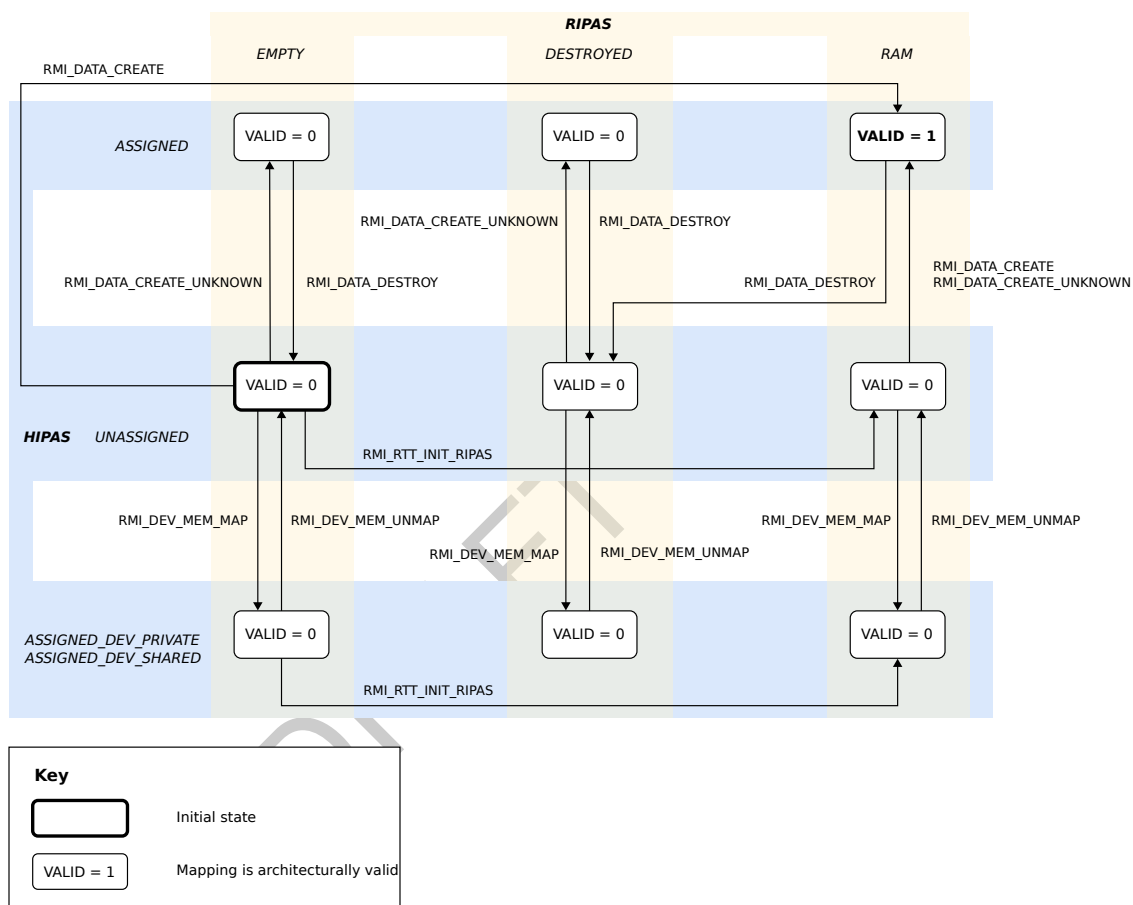
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.4 RMI_DEV_MEM_MAP command](#)
- [B4.3.5 RMI_DEV_MEM_UNMAP command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)

A5.3.4 Summary of changes to HIPAS and RIPAS of a Protected IPA

└_{TJMC}P

The following diagram summarizes HIPAS and RIPAS changes at a Protected IPA which can occur when the Realm state is NEW.

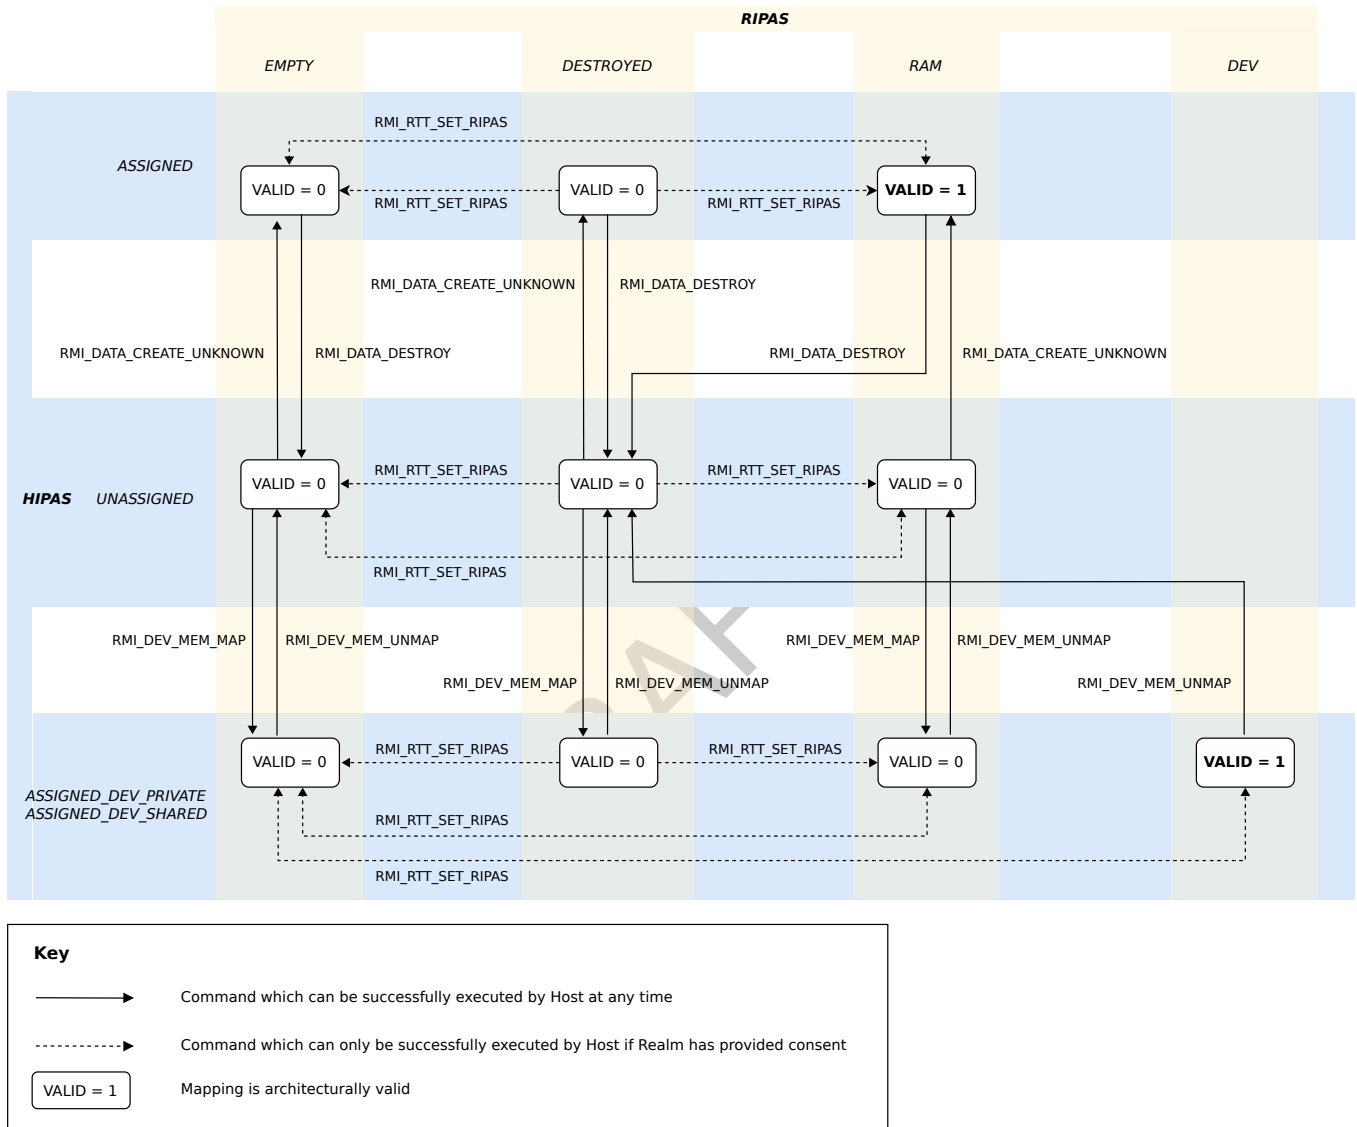
Transitions due to execution of RMI_RTT_DESTROY are omitted from the diagram. Execution of this command results in a transition to HIPAS=UNASSIGNED, RIPAS=DESTROYED.



I_VGKNJ

The following diagram summarizes HIPAS and RIPAS changes at a Protected IPA which can occur when the Realm state is REALM_ACTIVE.

Transitions due to execution of RMI_RTT_DESTROY are omitted from the diagram. Execution of this command results in a transition to HIPAS=UNASSIGNED, RIPAS=DESTROYED.



See also:

- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.4 RMI_DEV_MEM_MAP command](#)
- [B4.3.5 RMI_DEV_MEM_UNMAP command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)
- [B4.3.41 RMI_RTT_INIT_RIPAS command](#)
- [B4.3.44 RMI_RTT_SET_RIPAS command](#)

A5.3.5 Dependency of RMI command execution on RIPAS and HIPAS values

I_{HLHZS} The following table summarizes dependencies on RMI command execution on the current Protected IPA.

Command	Dependency on RIPAS	Dependency on HIPAS	New RIPAS	New HIPAS
RMI_DATA_CREATE	None	HIPAS is UNASSIGNED	RAM	ASSIGNED
RMI_DATA_CREATE_UNKNOWN	None	HIPAS is UNASSIGNED	Unchanged	ASSIGNED
RMI_DATA_DESTROY	If RIPAS is not RAM	HIPAS is ASSIGNED	Unchanged	UNASSIGNED
RMI_DATA_DESTROY	If RIPAS is RAM	HIPAS is ASSIGNED	DESTROYED	UNASSIGNED
RMI_RTT_CREATE	None	None	Unchanged	Unchanged
RMI_RTT_DESTROY	None	HIPAS of all entries is UNASSIGNED	DESTROYED	Unchanged
RMI_RTT_FOLD	RIPAS of all entries is identical	HIPAS of all entries is identical	Unchanged	Unchanged
RMI_RTT_INIT_RIPAS	None	HIPAS is UNASSIGNED	RAM	Unchanged
RMI_RTT_SET_RIPAS	Optionally, Realm may specify that RIPAS is not DESTROYED	None	As specified by Realm	Unchanged
RMI_DEV_MEM_MAP	None	HIPAS is UNASSIGNED	Unchanged	ASSIGNED_DEV_*
RMI_DEV_MEM_UNMAP	If RIPAS is not DEV	HIPAS is ASSIGNED_DEV_*	Unchanged	UNASSIGNED
RMI_DEV_MEM_UNMAP	If RIPAS is DEV	HIPAS is ASSIGNED_DEV_*	DESTROYED	UNASSIGNED

I_{WBRCN} Successful execution of RMI_DATA_CREATE_UNKNOWN does not depend on the RIPAS value of the target IPA.

I_{LCSVH} Successful execution of RMI_DATA_DESTROY does not depend on the RIPAS value of the target IPA.

I_{MMSBL} Successful execution of RMI_RTT_DESTROY does not depend on the RIPAS values of entries in the target RTT.

I_{TJCGT} Successful execution of RMI_RTT_FOLD does depend on the RIPAS values of entries in the target RTT.

I₀₀₄₉ Successful execution of RMI_DEV_MEM_UNMAP does not depend on the RIPAS value of the target IPA.


See also:

- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.38 RMI_RTT_CREATE command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)
- [B4.3.40 RMI_RTT_FOLD command](#)
- [B4.3.41 RMI_RTT_INIT_RIPAS command](#)
- [B4.3.44 RMI_RTT_SET_RIPAS command](#)

A5.3.6 Changes to HIPAS of an Unprotected IPA

└─YNYBY

The following diagram summarises HIPAS transitions for an Unprotected IPA.



See also:

- [A5.4 RIPAS change](#)
- [A5.5 Realm Translation Table](#)
- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)
- [B4.3.41 RMI_RTT_INIT_RIPAS command](#)
- [B4.3.44 RMI_RTT_SET_RIPAS command](#)
- [B5.3.6 RSI_IPA_STATE_SET command](#)

DRAFT

A5.4 RIPAS change

D_{BTSQY} A *RIPAS change* is a process via which the RIPAS of a region of Protected IPA space is changed, for a Realm whose state is `REALM_ACTIVE`.

I_{KXXBV} A RIPAS change consists of actions taken first by the Realm, and then by the Host:

- The Realm issues a *RIPAS change request* by executing `RSI_IPA_STATE_SET` or `RSI_RDEV_VALIDATE_MAPPING`
 - The input values to `RSI_IPA_STATE_SET` include:
 - The requested IPA range: `[base, top)`
 - The requested RIPAS value (either `EMPTY` or `RAM`)
 - A flag which indicates whether a change from `DESTROYED` should be permitted
 - For `RSI_RDEV_VALIDATE_MAPPING`:
 - The input values include the requested IPA range: `[base, top)`
 - The requested RIPAS value is implicitly `DEV`
 - A change from `DESTROYED` is implicitly not permitted
 - The RMM records these values in the REC, and then performs a REC exit due to RIPAS change pending.
- In response, the Host executes zero or more `RMI_RTT_SET_RIPAS` commands.
- If the requested RIPAS value was not `EMPTY` then at the next `RMI_REC_ENTER` the Host can optionally indicate that it rejects the RIPAS change request.

Output values from `RSI_IPA_STATE_SET` or `RSI_RDEV_VALIDATE_MAPPING` indicate:

- The top of the IPA range which has been modified by the command (`new_base`).
- If the requested RIPAS value was not `EMPTY`, whether the Host rejected the Realm request.

S_{CTTQV} Output values from `RSI_IPA_STATE_SET` or `RSI_RDEV_VALIDATE_MAPPING` are expected to be handled by the Realm as follows:

new_base	response	Meaning	Expected Realm action
<code>new_base == base</code>	<code>RSI_ACCEPT</code>	RIPAS change incomplete.	Call the command again, with <code>base = new_base</code> .
<code>base < new_base < top</code>	<code>RSI_ACCEPT</code>	RIPAS change incomplete.	Call the command again, with <code>base = new_base</code> .
<code>new_base == top</code>	<code>RSI_ACCEPT</code>	RIPAS change complete.	No further Realm action required.
<code>new_base == base</code>	<code>RSI_REJECT</code>	RIPAS change request rejected.	Depends on protocol agreed between Realm and Host, out of scope of this specification.
<code>base < new_base < top</code>	<code>RSI_REJECT</code>	RIPAS change to partial region <code>[base, new_base)</code> . Host rejected request to change RIPAS for region <code>[new_base, top)</code> .	Depends on protocol agreed between Realm and Host, out of scope of this specification.

I_{RFVTG} The RIPAS change process, together with the Realm Initial Measurement ensures that a Realm can always reliably determine the RIPAS of any Protected IPA.

I_{LPZWK} A RIPAS change is applied by one or more calls to the `RMI_RTT_SET_RIPAS` command.

I_{MMHMZ} Successful execution of `RMI_RTT_SET_RIPAS` targets an RTTE at address `rec.ripas_addr`.

I_{JHJGZ} On successful execution of `RMI_RTT_SET_RIPAS`, both of the following are set to the address of the next page whose RIPAS is to be modified:

- `rec.ripas_addr`
- The command output value

`I_GXDDX`

If both of the following are true on successful execution of `RMI_RTT_SET_RIPAS`

- The RIPAS change request indicated that a change from DESTROYED should not be permitted
- A page *P* within the target IPA range has RIPAS value DESTROYED

then `rec.ripas_addr` and the command output value are both set to *P*.

`I_HXKPB`

On REC entry following a REC exit due to RIPAS change, GPR values are updated to indicate for how much of the target IPA range the RIPAS change has been applied.

`S_TZYZV`

To complete a RIPAS change for a given target IPA range, a Realm should execute `RSI_IPA_STATE_SET` or `RSI_RDEV_VALIDATE_MAPPING` in a loop, until the value of X1 reaches the top of the target IPA range.

`R_LDMLC`

On REC entry following a REC exit due to RIPAS change, `rec.ripas_response` is set to the value of `enter.flags.ripas_response`.

`I_DRPPK`

If all of the following are true then the output value of `RSI_IPA_STATE_SET` or `RSI_RDEV_VALIDATE_MAPPING` indicates “Host rejected the request”:

- `rec.ripas_value` is RAM.
- `rec.ripas_addr` is not equal to `rec.ripas_top`.
- `rec.ripas_response` is REJECT.

Otherwise, the output value of `RSI_IPA_STATE_SET` or `RSI_RDEV_VALIDATE_MAPPING` indicates “Host accepted the request”.

`S_BZWWC`

Receipt of a rejection for a RIPAS change request whose parameters were valid is expected to be fatal for the Realm.

See also:

- [A2.3.2 REC attributes](#)
- [A4.2 REC entry](#)
- [A4.3.8 REC exit due to RIPAS change pending](#)
- [A5.2.2 Realm IPA state](#)
- [A7.1.1 Realm Initial Measurement](#)
- [Chapter A9 Realm device assignment](#)
- [B3.63 RecRipasResponseToRsi function](#)
- [B4.3.30 RMI_REC_ENTER command](#)
- [B4.3.44 RMI_RTT_SET_RIPAS command](#)
- [B5.3.6 RSI_IPA_STATE_SET command](#)
- [B5.3.23 RSI_RDEV_VALIDATE_MAPPING command](#)
- [D1.5.3 RIPAS change flow](#)

A5.5 Realm Translation Table

This section introduces the stage 2 translation table used by a Realm.

A5.5.1 RTT overview

D _{FRNCX}	A <i>Realm Translation Table</i> (RTT) is an abstraction over an Armv8-A stage 2 translation table used by a Realm.
I _{MBCVZ}	The attributes and format of an Armv8-A stage 2 translation table are defined by the Armv8-A Virtual Memory System Architecture (VMSA) Arm Architecture Reference Manual for A-Profile architecture [3] .
R _{PXNHQ}	The translation granule size of an RTT is 4KB.
I _{TQVTP}	The RMM architecture can only be deployed on a hardware platform which implements a translation granule size of 4KB.
I _{PHGQQ}	The contents of an RTT are not directly accessible to the Host.
I _{FPLRL}	The contents of an RTT are manipulated using RMM commands. These commands allow the Host to manipulate the contents of the RTT used by a Realm, subject to constraints imposed by the RMM.
D _{QTZDW}	An <i>RTT entry</i> (RTTE) is an abstraction over an Armv8-A stage 2 translation table descriptor.
I _{VYLT}	An RTTE contains an output address which can point to one of the following: <ul style="list-style-type: none">• Another RTT• A DATA Granule which is owned by the Realm• Non-secure memory which is accessible to both the Realm and the Host

A5.5.2 RTT structure and configuration

D _{VHLWF}	An <i>RTT tree</i> is a hierarchical data structure composed of RTTs, connected via Table Descriptors.
I _{KNPNX}	An RTT contains an array of RTTEs.
D _{HYTCJ}	An <i>RTT level</i> is the depth of an RTT within an RTT tree.
I _{KKMSX}	An RTT does not have an intrinsic “level” attribute. The level of an RTT is determined by its position within an RTT tree.
D _{QSYBS}	The RTT level of the root of an RTT tree is called the <i>starting level</i> .
I _{SSDBT}	The maximum depth of an RTT tree depends on all of the following: <ul style="list-style-type: none">• whether LPA2 is selected when the Realm is created• the <code>rtt_level_start</code> attribute of the Realm• the <code>ipa_width</code> attribute of the Realm.

See also:

- [A2.1.3 Realm attributes](#)
- [A3.3 Realm LPA2 and IPA width](#)

A5.5.3 RTT starting level

I _{FDWZF}	The RTT starting level is set when a Realm is created.
I _{YCPMF}	The number of starting level RTTs is architecturally defined as a function of the Realm IPA width and the RTT starting level. See Arm Architecture Reference Manual for A-Profile architecture [3] for further details.
I _{RYNXB}	The address of the first starting level RTT is stored in the RTT base attribute of the owning Realm.

I_{XXWQW} The RTT base attribute is set when a Realm is created.
See also:

- [A2.1.3 Realm attributes](#)

A5.5.4 RTT entry

I_{ZBGGZ} An RTT entry (RTTE) is an abstraction over an Armv8-A stage 2 translation table descriptor. The attributes and format of an Armv8-A stage 2 translation table descriptor are defined by the Armv8-A Virtual Memory System Architecture (VMSA) [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#).

D_{BNHQQ} An RTTE has a *state*.

The RTTE state values are shown in the following table.

Name	Description
ASSIGNED	This RTTE is identified by a Protected IPA. The output address of this RTTE points to a DATA Granule.
ASSIGNED_DEV_PRIVATE	This RTTE is identified by a Protected IPA. The output address of this RTTE points to a DEV_PRIVATE Granule.
ASSIGNED_DEV_SHARED	This RTTE is identified by a Protected IPA. The output address of this RTTE points to a DEV_SHARED Granule.
ASSIGNED_NS	This RTTE is identified by an Unprotected IPA. The output address of this RTTE points to an NS Granule.
AUX_DESTROYED	An auxiliary RTT was destroyed while a corresponding primary RTT entry was live.
TABLE	The output address of this RTTE points to the next-level RTT.
UNASSIGNED	This RTTE is identified by a Protected IPA. This RTTE is not associated with any Granule.
UNASSIGNED_NS	This RTTE is identified by an Unprotected IPA. This RTTE is not associated with any Granule.

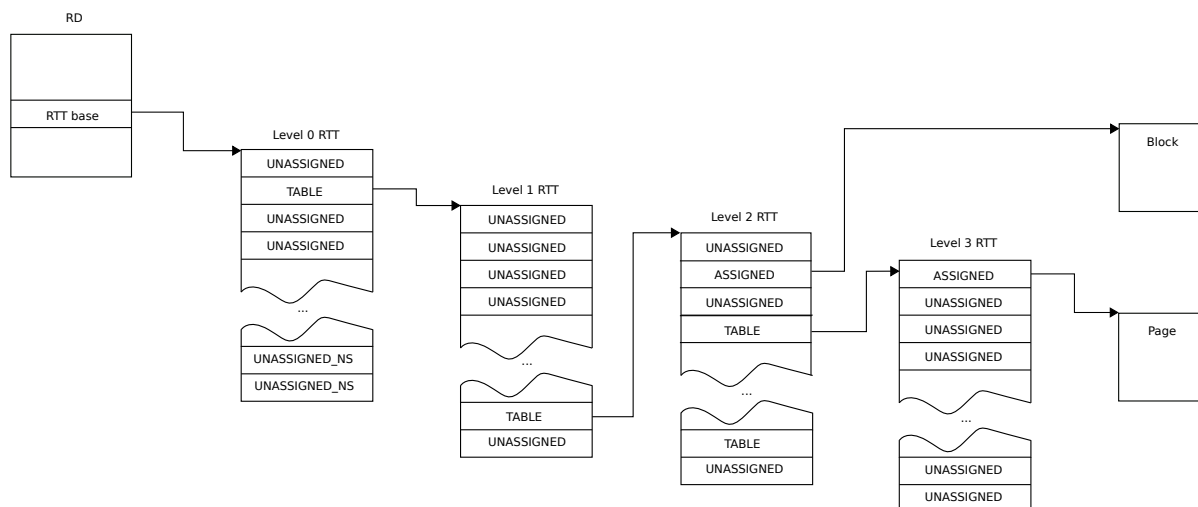
I_{QWQSB} The state of an RTTE in a RTT which is not level 1 or level 2 or level 3 is UNASSIGNED, UNASSIGNED_NS or TABLE.

D_{NSHSL} The output address of an RTTE whose state is TABLE and which is in a level n RTT is the physical address of a level $n+1$ RTT.

I_{DJZTM} An RTT whose level n is not the starting RTT level is pointed-to by exactly one TABLE RTTE in a level $n-1$ RTT.

I_{DXQWZ}

The following diagram shows an example RTT tree, annotated with RTTE states.



I_{FGWQS}

The function `AddrIsRttLevelAligned()` is used to evaluate whether an address is aligned to the address range described by an RTTE at a specified RTT level.

See also:

- [A5.3.1 Host IPA state](#)
- [B1.4 Command condition expressions](#)

A5.5.5 RTT reading

I_{KJWKQ}

Attributes of an RTTE, including the RTTE state, can be read by calling the `RMI_RTT_READ_ENTRY` command. The set of RTTE attributes which are returned depends on the state of the RTTE.

See also:

- [B4.3.43 RMI_RTT_READ_ENTRY command](#)

A5.5.6 RTT folding

D_{RMCLC}

An RTT is *homogeneous* if its entries satisfy one of the conditions in the following table. If an RTT is homogeneous, the following table specifies the state to which the parent RTTE is set.

Conditions on child RTT contents	Parent RTTE state
All of the following are true: <ul style="list-style-type: none"> • State of all entries is UNASSIGNED • RIPAS of all entries is the same 	UNASSIGNED
State of all entries is UNASSIGNED_NS	UNASSIGNED_NS
All of the following are true: <ul style="list-style-type: none"> • Level is 2 or 3 • State of all entries is ASSIGNED • Output address of first entry is aligned to size of the address range described by an entry in the parent RTT • Output addresses of all entries are contiguous • RIPAS of all entries is the same • S2AP of all entries is the same 	ASSIGNED

Conditions on child RTT contents	Parent RTTE state
<p>All of the following are true:</p> <ul style="list-style-type: none"> • Level is 2 or 3 • State of all entries is ASSIGNED_NS • Output address of first entry is aligned to size of the address range described by an entry in the parent RTT • Output addresses of all entries are contiguous • Attributes of all entries are identical 	ASSIGNED_NS
<p>All of the following are true:</p> <ul style="list-style-type: none"> • Level is 2 or 3 • State of all entries is ASSIGNED_DEV_PRIVATE • Output address of first entry is aligned to size of the address range described by an entry in the parent RTT • Output addresses of all entries are contiguous • RIPAS of all entries is the same • Memory attributes of all entries are the same 	ASSIGNED_DEV_PRIVATE
<p>All of the following are true:</p> <ul style="list-style-type: none"> • Level is 2 or 3 • State of all entries is ASSIGNED_DEV_SHARED • Output address of first entry is aligned to size of the address range described by an entry in the parent RTT • Output addresses of all entries are contiguous • RIPAS of all entries is the same • Memory attributes of all entries are the same 	ASSIGNED_DEV_SHARED

I_KDXLT	The function <code>RttIsHomogeneous()</code> is used to evaluate whether an RTT is homogeneous.
D_QPXP	<i>RTT folding</i> is the operation of destroying a homogeneous child RTT, and moving information which was stored in the child RTT into the parent RTTE.
I_QMGWK	On RTT folding, the state of the parent RTTE is determined from the contents of the child RTTEs.
I_LLWGH	The function <code>RttFold()</code> is used to evaluate the parent RTTE state which results from an RTT folding operation.
I_TPMGT	<p>On RTT folding, if the state of the parent RTTE is any of the following then the attributes of the parent RTTE are copied from the child RTTEs:</p> <ul style="list-style-type: none"> • ASSIGNED • ASSIGNED_NS • ASSIGNED_DEV_PRIVATE • ASSIGNED_DEV_SHARED <p>See also:</p> <ul style="list-style-type: none"> • A5.5.9 RTT destruction • A10.3.2 Stage 2 access permissions • B3.106 RttFold function • B3.107 RttIsHomogeneous function • B4.3.40 RMI_RTT_FOLD command

A5.5.7 RTT unfolding

D_HQQMG	<i>RTT unfolding</i> is the operation of creating a child RTT, and populating it based on the contents of the parent RTTE.
I_KWZXX	On RTT unfolding, the state of all RTTEs in the child RTT are set to the state of the parent RTTE.

I_{HMYSW} On RTT unfolding, if the state of the parent RTTE is any of the following then the output addresses of RTTEs in the child RTT are set to a contiguous range which starts from the address of the parent RTTE:

- ASSIGNED
- ASSIGNED_NS
- ASSIGNED_DEV_PRIVATE
- ASSIGNED_DEV_SHARED

See also:

- [B4.3.38 RMI_RTT_CREATE command](#)

A5.5.8 RTTE liveness and RTT liveness

D_{KCMLN} *RTTE liveness* is a property which means that a physical address is stored in the RTTE.

D_{HGYJZ} An RTTE is *live* if the RTTE state is any of the following:

- ASSIGNED
- ASSIGNED_NS
- ASSIGNED_DEV_PRIVATE
- ASSIGNED_DEV_SHARED
- TABLE

I_{RHLYZ} The function `RttSkipNonLiveEntries()` is used to scan an RTT to find the next live RTTE. The resulting IPA is returned to the Host from commands whose successful execution causes a live RTTE to become non-live.

X_{GQPSF} Identifying the next live RTTE allows the Host to avoid calls to `RMI_RTT_READ_ENTRY` when unmapping ranges of a Realm's IPA space, for example during Realm destruction.

D_{MPWLR} *RTT liveness* is a property which means that there exists another RMM data structure which is referenced by the RTT.

D_{YPSLW} An RTT is *live* if, for any of its entries, the RTTE state is any of the following:

- ASSIGNED
- ASSIGNED_DEV_PRIVATE
- ASSIGNED_DEV_SHARED
- TABLE

I_{MXJNY} Note that an RTT can be non-live, even if one of its entries is live. This would be the case for example if the RTT corresponds to an Unprotected IPA range and the state of one of its entries is `ASSIGNED_NS`.

I_{YPLKM} The function `RttIsLive()` is used to evaluate whether an RTT is live.

See also:

- [A5.5.9 RTT destruction](#)
- [B3.108 RttIsLive function](#)
- [B3.120 RttSkipNonLiveEntries function](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)
- [B4.3.46 RMI_RTT_UNMAP_UNPROTECTED command](#)

A5.5.9 RTT destruction

D_{VXRZW} *RTT destruction* is the operation of destroying a child RTT, and discarding information which was stored in the child RTT.

I_{PRMFR} An RTT cannot be destroyed if it is live.

I_{MDFQN} An RTT can be destroyed regardless of whether it is homogeneous.

I_{MCKSK}

Following RTT destruction, all of the following are true for the parent RTTE:

- RIPAS is DESTROYED
- RTTE state is UNASSIGNED

See also:

- [A5.2 Realm view of memory management](#)
- [A5.5.6 RTT folding](#)
- [A5.5.8 RTTE liveness and RTT liveness](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)

A5.5.10 RTT walk

I_{CBWSX}

An IPA is translated to a PA by walking an RTT tree, starting at the RTT base.

I_{FDWYV}

The behaviour of an RTT walk is defined by the Armv8-A Virtual Memory System Architecture (VMSA) [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#).

I_{TVGQD}

The inputs to an RTT walk are:

- a Realm Descriptor, which contains the address of the initial RTT
- an RTT tree index
- a target IPA
- a target RTT level.

The RTT walk terminates when either:

- it reaches the target RTT level, or
- it reaches an RTTE whose state is not TABLE.

D_{RBHVQ}

The result of an RTT walk performed by the RMM is a data structure of type `RmmRttWalkResult`.

The attributes of an `RmmRttWalkResult` are summarized in the following table.

Name	Type	Description
level	Int8	RTT level reached by the walk
rtt_addr	Address	Address of RTT reached by the walk
rtte	RmmRttEntry	RTTE reached by the walk

I_{ZSRCD}

The function `RmmRttWalkResult RttWalk(rd, addr, level)` is used to represent an RTT walk.

I_{FBZPQ}

The input address to an RTT walk is always less than 2^w , where w is the IPA width of the target Realm.

See also:

- [A2.1.3 Realm attributes](#)
- [A10.3.1 Auxiliary RTT](#)
- [B1.4 Command condition expressions](#)
- [B3.122 RttWalk function](#)
- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.38 RMI_RTT_CREATE command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)
- [B4.3.42 RMI_RTT_MAP_UNPROTECTED command](#)
- [B4.3.46 RMI_RTT_UNMAP_UNPROTECTED command](#)
- [C2.47 RmmRttWalkResult type](#)

A5.5.11 RTT entry attributes

A5.5.11.1 RTT entry attributes for ASSIGNED mappings

R_{KCFCT}	The cacheability attributes of an RTT entry whose state is ASSIGNED are independent of any stage 1 descriptors and of the state of the stage 1 MMU.
U_{NPVGN}	The RMM uses FEAT_S2FWB to ensure that the cacheability attributes of an RTT entry whose state is ASSIGNED are independent of stage 1 translation.
R_{JZKMH}	The attributes of an RTT entry whose state is ASSIGNED include the following: <ul style="list-style-type: none"> • Normal memory • Inner Write-Back Cacheable • Inner Shareable

See also:

- [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#)

A5.5.11.2 RTT entry attributes for ASSIGNED_DEV mappings

R_{0050}	The cacheability attributes of an RTT entry whose state is ASSIGNED_DEV_PRIVATE or ASSIGNED_DEV_SHARED are determined by the stage 1 descriptors.
U_{0051}	In an RTT entry whose state is ASSIGNED_DEV_PRIVATE or ASSIGNED_DEV_SHARED, <code>MemAttr</code> is set to 0b111.
R_{0052}	In an RTT entry whose state is ASSIGNED_DEV_PRIVATE or ASSIGNED_DEV_SHARED, <code>S2AP</code> is set to RW.

See also:

- [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#)
- [A9.5.3 Realm validation of device memory mappings](#)

A5.5.11.3 RTT entry attributes for ASSIGNED_NS mappings

D_{FJTMF}	The following attributes of an RTT entry whose state is ASSIGNED_NS are <i>Host-controlled Unprotected RTT attributes</i> : <ul style="list-style-type: none"> • ADDR • <code>MemAttr[2:0]</code> • S2AP
X_{QHLKB}	In an RTT entry whose state is ASSIGNED_NS, <code>MemAttr[3]</code> is RES0 because the RMM uses FEAT_S2FWB.
R_{QFLWD}	In an RTT entry whose state is ASSIGNED_NS, the shareability attributes are as follows: <ul style="list-style-type: none"> • Inner Shareable if the mapping is cacheable. • Outer Shareable if the mapping is non-cacheable.
U_{MCCRT}	The shareability attributes of an RTT entry which corresponds to an Unprotected IPA are expected to be controlled by the RMM as follows: <ul style="list-style-type: none"> • If LPA2 is enabled at stage 2 then the RMM is expected to set <code>VTOR_EL2.DS == '1'</code>. • If LPA2 is not enabled at stage 2 then the RMM is expected to set the value of the <code>SH</code> field in the translation table descriptor based on the value of the <code>MemAttr</code> field.

See also:

- [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#)
- [B3.101 RttDescriptorIsValidForUnprotected function](#)
- [B4.3.42 RMI_RTT_MAP_UNPROTECTED command](#)

A5.5.11.4 Hardware access flag and dirty bit management

R_{JRZTL}

Hardware access flag and dirty bit management is disabled for the stage 2 translation used by a Realm.

I_{QFGJC}

Hardware access flag and dirty bit management may be enabled by software executing within the Realm, for its own stage 1 translation.

DRAFT

Chapter A6

Realm interrupts and timers

This specification requires that a virtual Generic Interrupt Controller (vGIC) is presented to a Realm. This vGIC should be architecturally compliant with respect to GICv3 with no legacy operation.

The Host is able to inject virtual interrupts using the GIC virtual CPU interface.

The vGIC presented to a Realm is expected to be implemented via a combination of Host emulation and RMM mediation, as follows:

- Management of Non-secure physical interrupts is performed by the Host, via the GIC Interrupt Routing Infrastructure (IRI).
- The Host is responsible for emulating a GICv3 distributor MMIO interface.
- The Host is responsible for emulating a GICv3 redistributor MMIO interface for each REC.
- The GIC MMIO interfaces emulated by the Host must be presented to the Realm via its Unprotected IPA space.
- The Host may optionally provide a virtual Interrupt Translation Service (ITS). The Realm must allocate ITS tables within its Unprotected IPA space.
- The RMM allows the Host to control some of the GIC virtual CPU interface state which is observed by the Realm. This state is designed to be the minimum required to allow the Host to correctly manage interrupts for the Realm, with integrity guaranteed by the RMM for the remainder of the GIC CPU interface state.
- On REC exit, the RMM exposes some of the GIC virtual CPU interface state to the Host. This state is designed to be the minimum required to allow the Host to correctly manage interrupts for the Realm, with confidentiality guaranteed by the RMM for the remainder of the GIC virtual CPU interface state.

On every REC exit, the EL1 timer state is exposed to the Host. The RMM guarantees that a REC exit occurs whenever a Realm EL1 timer asserts or de-asserts its output.

See also:

- [Arm Generic Interrupt Controller \(GIC\) Architecture Specification version 3 and version 4 \[6\]](#)
- [A5.2.1 Realm IPA space](#)
- [D1.6 Realm interrupts and timers flows](#)

DRAFT

A6.1 Realm interrupts

This section describes the programming model for a REC's GIC CPU interface.

D _{XZVGB}	<p>The value of <code>enter.gicv3_lrs[n]</code> is valid if all of the following are true:</p> <ul style="list-style-type: none"> The value is an architecturally valid encoding of <code>ICH_LR<n>_EL2</code> according to Arm Generic Interrupt Controller (GIC) Architecture Specification version 3 and version 4 [6]. <code>HW == '0'</code>.
X _{DMSDZ}	<p>The GICv3 architecture states that, if <code>HW == '1'</code> then the virtual interrupt must be linked to a physical interrupt whose state is Active, otherwise behavior is undefined. The RMM is unable to validate that invariant, so it imposes the constraint that <code>HW == '0'</code>.</p>
D _{CPLDX}	<p>The value of <code>enter.gicv3_hcr</code> is valid if the value is an architecturally valid encoding of <code>ICH_HCR_EL2</code> according to Arm Generic Interrupt Controller (GIC) Architecture Specification version 3 and version 4 [6].</p>
R _{HLEFRY}	<p>REC entry fails if the value of any <code>enter.gicv3_*</code> attribute is invalid.</p>
R _{WNFRW}	<p>On REC entry, <code>ICH_LR<n>_EL2</code> is set to <code>enter.gicv3_lrs[n]</code>, for all values of <code>n</code> supported by the PE.</p>
R _{WVGJFJ}	<p>On REC entry, the following fields in <code>ICH_HCR_EL2</code> are set to the corresponding values in <code>enter.gicv3_hcr</code>:</p> <ul style="list-style-type: none"> UIE LRENPIE NPIE VGrp0EIE VGrp0DIE VGrp1EIE VGrp1DIE TDIR
I _{SMHXB}	<p>On REC entry, fields in <code>enter.gicv3_hcr</code> must be set to '0' except for the following:</p> <ul style="list-style-type: none"> UIE LRENPIE NPIE VGrp0EIE VGrp0DIE VGrp1EIE VGrp1DIE TDIR <p>If any other field in <code>enter.gicv3_hcr</code> is set to '1', then <code>RMI_REC_ENTER</code> fails.</p>
X _{LMXCX}	<p>The RMM provides access to the GIC virtual CPU interface to the Realm and therefore controls the enable bit and most trap bits in <code>ICH_HCR_EL2</code>. The maintenance interrupt control bits are controlled by the Host, because the maintenance interrupts are provided as hints to the hypervisor to allocate List Registers optimally and to correctly emulate GICv3 behavior. The <code>TDIR</code> bit is also controlled by the Host because it is used when supporting <code>EOImode == '1'</code> in the Realm. This mode is used to allow deactivation of virtual interrupts across RECs. This deactivation must be handled by the Host because the RMM can only operate on a single REC during execution of <code>RMI_REC_ENTER</code>.</p>
R _{LNQRL}	<p>A REC exit due to <code>IRQ</code> is not generated for an interrupt which is masked by the value of <code>ICC_PMR_EL1</code> at the time of REC entry.</p>
U _{GXCHC}	<p>The RMM should preserve the value of <code>ICC_PMR_EL1</code> during REC entry.</p>
R _{NKPNC}	<p>On REC exit, <code>exit.gicv3_vmcr</code> contains the value of <code>ICH_VMCR_EL2</code> at the time of the Realm exit.</p>
R _{SKQNF}	<p>On REC exit, <code>exit.gicv3_misr</code> contains the value of <code>ICH_MISR_EL2</code> at the time of the Realm exit.</p>

X _{DBGXB}	The Host could in principle infer the value of ICH_MISR_EL2 at the time of the Realm exit from the combination of <code>exit.gicv3_lrs[n]</code> and <code>exit.gicv3_hcr</code> . However, this would be cumbersome, error-prone, and diverge from the design of existing hypervisor software.
R _{QKZXD}	On REC exit, <code>exit.gicv3_lrs[n]</code> contains the value of ICH_LR<n>_EL2 at the time of the Realm exit, for all values of n supported by the PE.
R _{SNVZH}	On REC exit, the following fields in <code>exit.gicv3_hcr</code> contains the value of the corresponding field in ICH_HCR_EL2 at the time of the Realm exit: <ul style="list-style-type: none"> • EOICount • UIE • LRENPIE • NPIE • VGrp0EIE • VGrp0DIE • VGrp1EIE • VGrp1DIE • TDIR <p>All other fields contain zero.</p>
R _{FGQXT}	On REC exit, the values of the following registers may have changed: <ul style="list-style-type: none"> • ICH_AP0R<n>_EL2 • ICH_AP1R<n>_EL2 • ICH_LR<n>_EL2 • ICH_VMCR_EL2 • ICH_HCR_EL2
S _{QMJVJ}	It is the responsibility of the caller to save and restore GIC virtualization system control registers if their value needs to be preserved following execution of RMI_REC_ENTER.
X _{KDGHF}	On REC entry, the values of the GIC virtualization control system registers are overwritten. The Non-secure hypervisor runs at EL2 and therefore does not make direct use of the virtual GIC CPU interface for its own execution. This means that saving / restoring the caller's GIC virtualization control system registers would typically not be required and would add additional runtime overhead for each execution of RMI_REC_ENTER.
R _{VSBBS}	On REC exit, <code>ICH_HCR_EL2.En == '0'</code> .
X _{WLTX}	Disabling the virtual GIC CPU interface ensures that the caller does not receive unexpected GIC maintenance interrupts. A stronger constraint, for example stating that all GIC virtualization control system registers are zero on REC exit, was considered. However, this was rejected on the basis that it may preclude future optimisations, such as returning early from execution of RMI_REC_ENTER, without needing to first write zero to all GIC virtualization control system registers, if an interrupt is pending.

See also:

- [Arm Generic Interrupt Controller \(GIC\) Architecture Specification version 3 and version 4 \[6\]](#)
- [A4.2 REC entry](#)
- [A4.3 REC exit](#)
- [A10.4 Planes interrupts](#)
- [B4.3.30 RMI_REC_ENTER command](#)
- [B4.4.33 RmiRecEnter type](#)
- [B4.4.35 RmiRecExit type](#)
- [D1.6.1 Interrupt flow](#)

A6.2 Realm timers

This section describes the operation of architectural timers during Realm execution, including the following:

- The behavior of EL2 timers programmed by the Host
- The behavior of EL1 timers as perceived by the Realm
- The Realm timer state which is exposed to the Host on REC exit, in order to facilitate virtualization of timer interrupts

R _{LKNDV}	Architectural timers are available to a Realm and behave according to their architectural specification.
I _{VFYJV}	If the Host has programmed an EL1 timer to assert its output during Realm execution, that timer output is not guaranteed to assert.
R _{FKCHX}	If the Host has programmed an EL2 timer to assert its output during Realm execution, that timer output is guaranteed to assert.
R _{RJZRP}	Both the virtual and physical counter values are guaranteed to be monotonically increasing when read by a Realm, in accordance with the architectural counter behavior.
R _{JSMQP}	A read by a Realm of either the virtual or physical counter at the same place in the instruction flow would return the same value.
X _{YCDMW}	In order to ensure that the Realm has a consistent view of time, the virtual timer offset must be fixed for the lifetime of the Realm. The absolute value of the virtual timer offset is not important, so the value zero has been chosen for simplicity of both the specification and the implementation.
I _{FKMGZ}	The rule that virtual and physical counter values are identical may need to be amended if a future version of the specification supports migration and / or virtualization of time based on the virtual counter differing from the physical counter.
R _{SVCMR}	On a change in the output of an EL1 timer which requires a Realm-observable change to the state of virtual interrupts, a REC exit occurs.
R _{VWQDH}	On REC exit, Realm EL1 timer state is exposed via the RmiRecExit object: <ul style="list-style-type: none">• <code>exit.cntv_ctl</code> contains the value of <code>CNTV_CTL_EL0</code> at the time of the Realm exit.• <code>exit.cntv_cval</code> contains the value of <code>CNTV_CVAL_EL0</code> at the time of the Realm exit, expressed as if the virtual counter offset was zero.• <code>exit.cntp_ctl</code> contains the value of <code>CNTP_CTL_EL0</code> at the time of the Realm exit.• <code>exit.cntp_cval</code> contains the value of <code>CNTP_CVAL_EL0</code> at the time of the Realm exit, expressed as if the physical counter offset was zero.
S _{PYWWE}	The Host should check the Realm EL1 timer state on every return from <code>RMI_REC_ENTER</code> and update virtual interrupt state accordingly. This is true regardless of the value of <code>exit.exit_reason</code> : even if the return occurred for a reason unrelated to timers (for example, a REC exit due to Data Abort), the Realm EL1 timer state should be checked.
I _{VRWGS}	On REC entry, for both the EL1 Virtual Timer and the EL1 Physical Timer, if the EL1 timer asserts its output in the state described in the REC exit structure from the previous REC exit then the RMM masks the hardware timer signal before returning to the Realm.

This masking is done to allow the Realm to make forward progress, which would otherwise be prevented by the hardware timer generating a physical interrupt that would cause a Realm exit.

During Realm execution, when the hardware timer signal is masked, the Realm may write to the timer registers, causing the hardware timer to become de-asserted and possibly asserted again. Such changes in the output of the EL1 timer are not required to result in a REC exit if the RMM can infer that the change should not result in a Realm-observable change to the state of virtual interrupts.

It is only when a change in the hardware timer output means that the corresponding virtual interrupt needs to be made pending or idle, that a REC exit must occur.

See also:

- [A4.3 REC exit](#)
- [A10.5 Planes timers](#)
- [B4.4.35 RmiRecExit type](#)
- [D1.6.2 Timer interrupt delivery flow](#)

DRAFT

Chapter A7

Realm measurement and attestation

This section describes how the initial state of a Realm is measured and can be attested.

A7.1 Realm measurements

This section describes how Realm measurement values are calculated.

D_{SJWWS}	A Realm measurement value is a rolling hash.
D_{YKDBY}	A <i>Realm Hash Algorithm</i> (RHA) is an algorithm which is used to extend a Realm measurement value.
I_{NRKWB}	The RHA used by a Realm is selected via the <code>hash_algo</code> attribute.

See also:

- [A2.1.3 Realm attributes](#)
- [A3.2 Realm hash algorithm](#)
- [A7.2.3.1.4 Realm Initial Measurement claim](#)
- [A7.2.3.1.5 Realm Extensible Measurements claim](#)

A7.1.1 Realm Initial Measurement

This section describes how the Realm Initial Measurement (RIM) is calculated.

I_{XKSBZ}	The initial RIM value for a Realm is calculated from a subset of the Realm parameters.
I_{NCNDK}	A RIM is extended by applying the RHA to the inputs of RMM operations which are executed during Realm construction.
I_{NQQTF}	The following operations cause a RIM to be extended: <ul style="list-style-type: none"> • Creation of a DATA Granule during Realm construction • Creation of a runnable REC • Changes to RIPAS of Protected IPA during Realm construction
R_{VMPZG}	On execution of an operation which requires extension of a RIM, the RMM first constructs a <i>measurement descriptor</i> structure. The measurement descriptor contents include the current RIM value. The new RIM value is computed by applying the RHA to the measurement descriptor.

$$desc = MeasurementDescriptor(M_{i-1}, \dots)$$

$$M_i = RHA(desc)$$

I_{FQHFC}	A RIM is immutable while the state of the Realm is <code>REALM_ACTIVE</code> . This implies that a RIM reflects the configuration and contents of the Realm at the moment when it transitioned from the <code>REALM_NEW</code> to the <code>REALM_ACTIVE</code> state.
I_{DQGPT}	A RIM depends upon the order of the RMM operations which are executed during Realm construction.
S_{VZNCW}	The order in which RMM operations are executed during Realm construction must be agreed between the Realm owner (or a delegate of the Realm owner which will receive and validate the RIM) and the Host which executes the RMM commands. This ensures that a correctly-constructed Realm will have the expected measurement.
I_{LTWBL}	The value of a RIM can be read using the <code>RSI_MEASUREMENT_READ</code> command.

See also:

- [B4.3.1.4 RMI_DATA_CREATE extension of RIM](#)
- [B4.3.25.4 RMI_REALM_CREATE initialization of RIM](#)
- [B4.3.28.4 RMI_REC_CREATE extension of RIM](#)
- [B4.3.41.4 RMI_RTT_INIT_RIPAS extension of RIM](#)
- [B5.3.8 RSI_MEASUREMENT_READ command](#)

A7.1.2 Realm Extensible Measurement

This section describes the behavior of a Realm Extensible Measurement (REM).

I_{QJDDWM}

A REM is extended using the RSI_MEASUREMENT_EXTEND command.

I_{CTMBT}

The value of a REM can be read using the RSI_MEASUREMENT_READ command.

I_{MDQRP}

The initial value of a REM is zero.

See also:

- [B5.3.7 RSI_MEASUREMENT_EXTEND command](#)
- [B5.3.8 RSI_MEASUREMENT_READ command](#)

DRAFT

A7.2 Realm attestation

This section describes the primitives which are used to support remote Realm attestation.

A7.2.1 Attestation token

D_{VRRLN} A CCA attestation token is a collection of claims about the state of a Realm and of the CCA platform on which the Realm is running.

I_{BXSBD} A CCA attestation token consists of two parts:

- Realm token
 - Contains attributes of the Realm, including:
 - Realm Initial Measurement
 - Realm Extensible Measurements
- CCA platform token
 - Contains attributes of the CCA platform on which the Realm is running, including:
 - CCA platform identity
 - CCA platform lifecycle state
 - CCA platform software component measurements

I_{JKJCQ} The size of a CCA attestation token may be greater than 4KB.

See also:

- [A7.1.1 Realm Initial Measurement](#)
- [A7.1.2 Realm Extensible Measurement](#)

A7.2.2 Attestation token generation

I_{KMRH} The process for a Realm to obtain an attestation token is:

- Call `RSI_ATTESTATION_TOKEN_INIT` once
- Call `RSI_ATTESTATION_TOKEN_CONTINUE` in a loop, until the result is not `RSI_INCOMPLETE`

Each call to `RSI_ATTESTATION_TOKEN_CONTINUE` retrieves up to one Granule of the attestation token.

S_{XMLMF}

The following pseudocode illustrates the process of a Realm obtaining an attestation token.

```

int get_attestation_token(...)
{
    int ret;
    uint64_t size, max_size;
    uint64_t buf, granule;

    ret = RSI_ATTESTATION_TOKEN_INIT(challenge, &max_size);
    if (ret) {
        return ret;
    }

    buf = alloc(max_size);
    granule = buf;

    do { // Retrieve one Granule of data per loop iteration
        uint64_t offset = 0;

        do { // Retrieve sub-Granule chunk of data per loop iteration
            size = GRANULE_SIZE - offset;
            ret = RSI_ATTESTATION_TOKEN_CONTINUE(granule, offset, size, &len);
            offset += len;
        } while (ret == RSI_INCOMPLETE && offset < GRANULE_SIZE);

        // "offset" bytes of data are now ready for consumption from "granule"

        if (ret == RSI_INCOMPLETE) {
            granule += GRANULE_SIZE;
        }
    } while ((ret == RSI_INCOMPLETE) && (granule < buf + max_size));

    return ret;
}

```

I_{ZWQCB}

Up to one attestation token generation operation may be ongoing on a REC.

I_{TMJVG}

On execution of RSI_ATTESTATION_TOKEN_INIT, if an attestation token generation operation is ongoing on the calling REC, it is terminated.

I_{WTKDD}

The challenge value provided to RSI_ATTESTATION_TOKEN_INIT is included in the generated attestation token. This allows the relying party to establish freshness of the attestation token.

If the size of the challenge provided by the relying party is less than 64 bytes, it should be zero-padded prior to calling RSI_ATTESTATION_TOKEN_INIT. Arm recommends that the challenge should contain at least 32 bytes of unique data.

I_{GKDJW}

Generation of an attestation token can be a long-running operation, during which interrupts may need to be handled.

I_{CXSJP}

If a physical interrupt becomes pending during execution of RSI_ATTESTATION_TOKEN_CONTINUE, a REC exit due to IRQ can occur.

On the next entry to the REC:

- If a virtual interrupt is pending on that REC, it is taken to the REC's exception handler
- RSI_ATTESTATION_TOKEN_CONTINUE returns RSI_INCOMPLETE
- The REC should call RSI_ATTESTATION_TOKEN_CONTINUE again

See also:

- [A4.3.5 REC exit due to IRQ](#)
- [A6.1 Realm interrupts](#)

- [A7.2.3.1.1 Realm challenge claim](#)
- [B5.3.1 RSI_ATTESTATION_TOKEN_CONTINUE command](#)
- [B5.3.2 RSI_ATTESTATION_TOKEN_INIT command](#)
- [D1.7.1 Attestation token generation flow](#)
- [D1.7.2 Handling interrupts during attestation token generation flow](#)

A7.2.3 Attestation token format

<code>I_TFHGX</code>	The CCA attestation token is a profiled IETF Entity Attestation Token (EAT).
<code>I_LPTVH</code>	The CCA attestation token is a Concise Binary Object Representation (CBOR) map, in which the map values are the Realm token and the CCA platform token.
<code>I_YZPHG</code>	The Realm token contains structured data in CBOR, wrapped with a COSE_Sign1 envelope according to the CBOR Object Signing and Encryption (COSE) standard.
<code>I_MMQZG</code>	The Realm token is signed by the Realm Attestation Key (RAK).
<code>I_WBGNP</code>	The CCA platform token contains structured data in CBOR, wrapped with a COSE_Sign1 envelope according to the COSE standard.
<code>I_CGYKX</code>	The CCA platform token is signed by the Initial Attestation Key (IAK).
<code>I_CCGQH</code>	The CCA platform token contains a hash of RAK_pub. This establishes a cryptographic binding between the Realm token and the CCA platform token.
<code>I_PTKYD</code>	<p>The CCA attestation token is defined as follows:</p> <hr/> <pre> cca-token = #6.399(cca-token-collection) ; CMW Collection ; (draft-ietf-rats-msg-wrap) cca-platform-token = bstr .cbor COSE_Sign1_Tagged cca-realm-delegated-token = bstr .cbor COSE_Sign1_Tagged cca-token-collection = { 44234 => cca-platform-token ; 44234 = 0xACCA 44241 => cca-realm-delegated-token } ; EAT standard definitions COSE_Sign1_Tagged = #6.18(COSE_Sign1) ; Deliberately shortcut these definitions until EAT is finalised and able to ; pull in the full set of definitions COSE_Sign1 = "COSE-Sign1 placeholder" </pre> <hr/>
<code>I_HZNNH</code>	The composition of the CCA attestation token is summarised in the following figure.




Figure A7.1: Attestation token format

See also:

- [Arm CCA Security model](#) [4]
- [Concise Binary Object Representation \(CBOR\)](#) [7]
- [CBOR Object Signing and Encryption \(COSE\)](#) [8]
- [Entity Attestation Token \(EAT\)](#) [9]
- [A7.2.3.1 Realm claims](#)
- [A7.2.3.2 CCA platform claims](#)

A7.2.3.1 Realm claims

This section defines the format of the Realm token claim map. The format is described using a combination of Concise Data Definition Language (CDDL) and text description.

I_{HKBHC}

The Realm token claim map is defined as follows:

```
cca-realm-claims = (cca-realm-claim-map)
```

```
cca-realm-claim-map = {  
    cca-realm-challenge  
    ? cca-realm-profile  
    cca-realm-personalization-value  
    cca-realm-initial-measurement  
    cca-realm-extensible-measurements  
    cca-realm-hash-algo-id  
    cca-realm-public-key  
    cca-realm-public-key-hash-algo-id  
    cca-realm-mec-policy  
}
```

See also:

- [Concise Data Definition Language \(CDDL\) \[10\]](#)
- [A7.2.3.1.1 Realm challenge claim](#)
- [A7.2.3.1.2 Realm profile claim](#)
- [A7.2.3.1.3 Realm Personalization Value claim](#)
- [A7.2.3.1.4 Realm Initial Measurement claim](#)
- [A7.2.3.1.5 Realm Extensible Measurements claim](#)
- [A7.2.3.1.6 Realm hash algorithm ID claim](#)
- [A7.2.3.1.7 Realm MEC policy claim](#)
- [A7.2.3.1.8 Realm public key claim](#)
- [A7.2.3.1.9 Realm public key hash algorithm identifier claim](#)
- [A7.2.3.1.10 Collated CDDL for Realm claims](#)
- [A7.2.3.1.11 Example Realm claims](#)

A7.2.3.1.1 Realm challenge claim

I_{TFWXQ}

The Realm challenge claim is used to carry the challenge provided by the caller to demonstrate freshness of the generated token.

I_{RVLZK}

The Realm challenge claim is identified using the EAT_{nonce} label (10).

I_{MNVNP}

The length of the Realm challenge is 64 bytes.

I_{PXMXF}

The Realm challenge claim must be present in a Realm token.

I_{BXGFN}

The format of the Realm challenge claim is defined as follows:

```
cca-realm-challenge-label = 10  
cca-realm-challenge-type = bytes .size 64  
  
cca-realm-challenge = (  
    cca-realm-challenge-label => cca-realm-challenge-type  
)
```

See also:

- [A7.2.2 Attestation token generation](#)
- [B5.3.2 RSI_ATTESTATION_TOKEN_INIT command](#)

A7.2.3.1.2 Realm profile claim

I _{CVNNV}	The Realm profile claim identifies the EAT profile to which the Realm token conforms.
I _{SMSCF}	The Realm profile claim is identified using the EAT _{profile} label (265).
I _{XSSJY}	The Realm profile claim is optional in a CCA Realm token.
I _{GQTJT}	If the Realm profile is not included in a CCA Realm token then the profile value used in the CCA Platform token should refer to a profile that describes both Platform and Realm claims.
I _{SWDJM}	The format of the Realm profile claim is defined as follows:

```
cca-realm-profile-label = 265 ; EAT profile

cca-realm-profile-type = "tag:arm.com,2024:realm#1.1.0"

cca-realm-profile = (
    cca-realm-profile-label => cca-realm-profile-type
)
```

A7.2.3.1.3 Realm Personalization Value claim

I _{SCNXB}	The Realm Personalization Value claim contains the RPV which was provided at Realm creation.
I _{BKZPD}	The Realm Personalization Value claim must be present in a Realm token.
I _{QKNDV}	The format of the Realm Personalization Value claim is defined as follows:

```
cca-realm-personalization-value-label = 44235
cca-realm-personalization-value-type = bytes .size 64

cca-realm-personalization-value = (
    cca-realm-personalization-value-label => cca-realm-personalization-value-type
)
```

See also:

- [A2.1.3 Realm attributes](#)

A7.2.3.1.4 Realm Initial Measurement claim

I _{BXKGD}	The Realm Initial Measurement claim contains the values of the Realm Initial Measurement.
I _{FZQSM}	The Realm Initial Measurement claim must be present in a Realm token.
I _{GGTNH}	The format of the Realm Initial Measurement claim is defined as follows:

```
cca-realm-measurement-type = bytes .size 32 / bytes .size 48 / bytes .size 64
cca-realm-initial-measurement-label = 44238

cca-realm-initial-measurement = (
    cca-realm-initial-measurement-label => cca-realm-measurement-type
)
```

See also:

- [A7.1 Realm measurements](#)
- [A7.2.3.1.5 Realm Extensible Measurements claim](#)

A7.2.3.1.5 Realm Extensible Measurements claim

I _{KFNMV}	The Realm Extensible Measurements claim contains the values of the Realm Extensible Measurements.
I _{DSNFB}	The Realm Extensible Measurements claim must be present in a Realm token.

I _{ZKVMN}	<p>The format of the Realm measurements claim is defined as follows:</p> <hr/> <pre>cca-realm-measurement-type = bytes .size 32 / bytes .size 48 / bytes .size 64 cca-realm-extensible-measurements-label = 44239</pre> <pre>cca-realm-extensible-measurements = (cca-realm-extensible-measurements-label => [4*4 cca-realm-measurement-type])</pre> <hr/> <p>See also:</p> <ul style="list-style-type: none"> • A7.1 Realm measurements • A7.2.3.1.4 Realm Initial Measurement claim <p>A7.2.3.1.6 Realm hash algorithm ID claim</p>
I _{DGCGG}	<p>The Realm hash algorithm ID claim identifies the algorithm used to calculate all hash values which are present in the Realm token.</p>
I _{PVLCJ}	<p>Arm recommends that the value of the Realm hash algorithm ID claim is an IANA Hash Function name IANA Named Information Hash Algorithm Registry [11].</p>
I _{WKVCQ}	<p>The Realm hash algorithm ID claim must be present in a Realm token.</p>
I _{PWPLJ}	<p>The format of the Realm hash algorithm ID claim is defined as follows:</p> <hr/> <pre>cca-realm-hash-algo-id-label = 44236</pre> <pre>cca-realm-hash-algo-id = (cca-realm-hash-algo-id-label => text)</pre> <hr/> <p>A7.2.3.1.7 Realm MEC policy claim</p>
I ₀₀₅₃	<p>The Realm MEC policy identifies the MEC policy of the Realm.</p>
I ₀₀₅₄	<p>The Realm MEC policy claim must be present in a Realm token.</p>
R ₀₀₅₅	<p>On a platform which does not implement FEAT_MEC, the value of the Realm MEC policy claim is “shared”.</p>
I ₀₀₅₆	<p>The format of the Realm MEC policy claim is defined as follows:</p> <hr/> <pre>cca-realm-mec-policy-label = 44241</pre> <pre>cca-realm-mec-policy = (cca-realm-mec-policy-label => "shared" / "private")</pre> <hr/> <p>See also:</p> <ul style="list-style-type: none"> • Chapter A11 Realm memory encryption <p>A7.2.3.1.8 Realm public key claim</p>
I _{ZCFMQ}	<p>The Realm public key claim identifies the key which is used to sign the Realm token.</p>
I _{WBNHC}	<p>The value of the Realm public key claim is a CBOR bstr of a COSE_Key structure. The parameters used for the COSE_Key are profile-specific.</p>
I _{LSNPQ}	<p>The Realm public key claim must be present in a Realm token.</p>
I _{NNNDS}	<p>The format of the Realm public key claim is defined as follows:</p> <hr/> <pre>cca-realm-public-key-label = 44237</pre> <pre>cca-realm-public-key-type = bstr .cbor COSE_Key</pre> <hr/>

```

cca-realm-public-key = (
    cca-realm-public-key-label => cca-realm-public-key-type
)
COSE_Key-label = int / tstr

COSE_Key-values = any

; See RFC8152 for full definition of COSE_Key
COSE_Key = {
    1 => tstr / int,          ; kty
    ? 2 => bstr,              ; kid
    ? 3 => tstr / int,        ; alg
    ? 4 => [+ (tstr / int) ], ; key_ops
    ? 5 => bstr,              ; Base IV
    * COSE_Key-label => COSE_Key-values
}

```

See also:

- [SEC 1: Elliptic Curve Cryptography, version 2.0 \[12\]](#)
- [A7.2.3.1.9 Realm public key hash algorithm identifier claim](#)
- [A7.2.3.2.2 CCA platform challenge claim](#)

A7.2.3.1.9 Realm public key hash algorithm identifier claim

I_{WWSLP}

The Realm public key hash algorithm identifier claim identifies the algorithm used to calculate H(RAK_{pub}).

I_{TNRBN}

The Realm public key hash algorithm identifier claim must be present in a Realm token.

I_{NNPVX}

The format of the Realm public key hash algorithm identifier claim is defined as follows:

```

cca-realm-public-key-hash-algo-id-label = 44240

cca-realm-public-key-hash-algo-id = (
    cca-realm-public-key-hash-algo-id-label => text
)

```

See also:

- [SEC 1: Elliptic Curve Cryptography, version 2.0 \[12\]](#)
- [A7.2.3.1.8 Realm public key claim](#)
- [A7.2.3.2.2 CCA platform challenge claim](#)

A7.2.3.1.10 Collated CDDL for Realm claims

D_{DCYXZ}

The format of the Realm token claim map is defined as follows:

```
cca-realm-claims = (cca-realm-claim-map)

cca-realm-claim-map = {
    cca-realm-challenge
    ? cca-realm-profile
    cca-realm-personalization-value
    cca-realm-initial-measurement
    cca-realm-extensible-measurements
    cca-realm-hash-algo-id
    cca-realm-public-key
    cca-realm-public-key-hash-algo-id
    cca-realm-mec-policy
}
cca-realm-challenge-label = 10
cca-realm-challenge-type = bytes .size 64

cca-realm-challenge = (
    cca-realm-challenge-label => cca-realm-challenge-type
)
cca-realm-profile-label = 265 ; EAT profile

cca-realm-profile-type = "tag:arm.com,2024:realm#1.1.0"

cca-realm-profile = (
    cca-realm-profile-label => cca-realm-profile-type
)
cca-realm-personalization-value-label = 44235
cca-realm-personalization-value-type = bytes .size 64

cca-realm-personalization-value = (
    cca-realm-personalization-value-label => cca-realm-personalization-value-type
)
cca-realm-measurement-type = bytes .size 32 / bytes .size 48 / bytes .size 64
cca-realm-initial-measurement-label = 44238

cca-realm-initial-measurement = (
    cca-realm-initial-measurement-label => cca-realm-measurement-type
)
cca-realm-extensible-measurements-label = 44239

cca-realm-extensible-measurements = (
    cca-realm-extensible-measurements-label => [ 4*4 cca-realm-measurement-type ]
)
cca-realm-hash-algo-id-label = 44236

cca-realm-hash-algo-id = (
    cca-realm-hash-algo-id-label => text
)
cca-realm-public-key-label = 44237

cca-realm-public-key-type = bstr .cbor COSE_Key

cca-realm-public-key = (
    cca-realm-public-key-label => cca-realm-public-key-type
)
COSE_Key-label = int / tstr
```



```
COSE_Key-values = any

; See RFC8152 for full definition of COSE_Key
COSE_Key = {
    1 => tstr / int,          ; kty
    ? 2 => bstr,              ; kid
    ? 3 => tstr / int,        ; alg
    ? 4 => [+ (tstr / int) ], ; key_ops
    ? 5 => bstr,              ; Base IV
    * COSE_Key-label => COSE_Key-values
}
cca-realm-public-key-hash-algo-id-label = 44240

cca-realm-public-key-hash-algo-id = (
    cca-realm-public-key-hash-algo-id-label => text
)
cca-realm-mec-policy-label = 44241

cca-realm-mec-policy = (
    cca-realm-mec-policy-label => "shared" / "private"
)
```

DRAFT

A7.2.3.1.11 Example Realm claims

 $\mathbb{I}_{\text{CPTFR}}$

An example Realm claim map is shown below in COSE-DIAG format:

[illegible]

A7.2.3.2 CCA platform claims

This section defines the format of the CCA platform token claim map. The format is described using a combination of Concise Data Definition Language (CDDL) and text description.

I_{FJKFY}

The CCA platform token claim map is defined as follows:

```
cca-platform-claims = (cca-platform-claim-map)

cca-platform-claim-map = {
    cca-platform-profile
    cca-platform-challenge
    cca-platform-implementation-id
    cca-platform-instance-id
    cca-platform-config
    cca-platform-lifecycle
    cca-platform-sw-components
    ? cca-platform-verification-service
    cca-platform-hash-algo-id
}
```

See also:

- [Concise Data Definition Language \(CDDL\) \[10\]](#)
- [A7.2.3.2.1 CCA platform profile claim](#)
- [A7.2.3.2.2 CCA platform challenge claim](#)
- [A7.2.3.2.3 CCA platform Implementation ID claim](#)
- [A7.2.3.2.4 CCA platform Instance ID claim](#)
- [A7.2.3.2.5 CCA platform config claim](#)
- [A7.2.3.2.6 CCA platform lifecycle claim](#)
- [A7.2.3.2.7 CCA platform software components claim](#)
- [A7.2.3.2.8 CCA platform verification service claim](#)
- [A7.2.3.2.9 CCA platform hash algorithm ID claim](#)
- [A7.2.3.3.1 Collated CDDL for CCA platform claims](#)
- [A7.2.3.3.2 Example CCA platform claims](#)

A7.2.3.2.1 CCA platform profile claim

I_{FQYTP}

The CCA platform profile claim identifies the EAT profile to which the CCA platform token conforms. Note that because the platform token is expected to be issued when bound to a Realm token, the profile document should also include the relevant Realm profile or a reference to that profile.

I_{XMVFR}

The CCA platform profile claim is identified using the EAT `profile` label (265).

I_{GKKNR}

The CCA platform profile claim must be present in a CCA platform token.

I_{MHRTD}

The format of the CCA platform profile claim is defined as follows:

```
cca-platform-profile-label = 265 ; EAT profile

cca-platform-profile-type = "tag:arm.com,2024:cca_platform#1.1.0"

cca-platform-profile = (
    cca-platform-profile-label => cca-platform-profile-type
)
```

A7.2.3.2.2 CCA platform challenge claim

I_{TKTWZ}

The CCA platform challenge claim contains a hash of the public key used to sign the Realm token.

I_{CLJJK}

The CCA platform challenge claim is identified using the EAT `nonce` label (10).

I_{XHLVJ}

The length of the CCA platform challenge is either 32, 48 or 64 bytes.

I_{GVHNX} The CCA platform challenge claim must be present in a CCA platform token.

I_{LRWHR} The format of the CCA platform challenge claim is defined as follows:

```
cca-hash-type = bytes .size 32 / bytes .size 48 / bytes .size 64
cca-platform-challenge-label = 10
```

```
cca-platform-challenge = (
    cca-platform-challenge-label => cca-hash-type
)
```

See also:

- [A7.2.3.1.8 Realm public key claim](#)

A7.2.3.2.3 CCA platform Implementation ID claim

I_{SMWND} The CCA platform Implementation ID claim uniquely identifies the implementation of the CCA platform.

I_{NDVFB} The value of the CCA platform Implementation ID claim can be used by a verification service to locate the details of the CCA platform implementation from an endorser or manufacturer. Such details are used by a verification service to determine the security properties or certification status of the CCA platform implementation.

I_{RXPVW} The semantics of the CCA platform Implementation ID value are defined by the manufacturer or a particular certification scheme. For example, the ID could take the form of a product serial number, database ID, or other appropriate identifier.

I_{SRPZY} The CCA platform Implementation ID claim does not identify a particular instance of the CCA implementation.

I_{NTCFY} The CCA platform Implementation ID claim must be present in a CCA platform token.

I_{DHYDG} The format of the CCA platform Implementation ID claim is defined as follows:

```
cca-platform-implementation-id-label = 2396 ; PSA implementation ID
cca-platform-implementation-id-type = bytes .size 32

cca-platform-implementation-id = (
    cca-platform-implementation-id-label => cca-platform-implementation-id-type
)
```

See also:

- [Arm CCA Security model \[4\]](#)
- [A7.2.3.2.4 CCA platform Instance ID claim](#)

A7.2.3.2.4 CCA platform Instance ID claim

I_{ZYZB} The CCA platform Instance ID claim represents the unique identifier of the Initial Attestation Key (IAK) for the CCA platform.

I_{XVLLN} The CCA platform Instance ID claim is identified using the EAT ueid label (256).

R_{HVTNC} The first byte of the CCA platform Instance ID value must be 0x01.

I_{ZNGDF} The CCA platform Instance ID claim must be present in a CCA platform token.

I_{VPKJN} The format of the CCA platform Instance ID claim is defined as follows:

```
cca-platform-instance-id-label = 256 ; EAT ueid

; TODO: require that the first byte of cca-platform-instance-id-type is 0x01
; EAT UEIDs need to be 7 - 33 bytes
cca-platform-instance-id-type = bytes .size 33

cca-platform-instance-id = (
    cca-platform-instance-id-label => cca-platform-instance-id-type
)
```

See also:

- [Arm CCA Security model \[4\]](#)
- [A7.2.3.2.3 CCA platform Implementation ID claim](#)

A7.2.3.2.5 CCA platform config claim

<code>I_WVQJT</code>	The CCA platform config claim describes the set of chosen implementation options of the CCA platform. As an example, these may include a description of the level of physical memory protection which is provided.
<code>U_GPXWX</code>	The CCA platform config claim is expected to contain the System Properties field which is present in the Root Non-volatile Storage (RNVS) public parameters.
<code>I_MJHQQ</code>	The CCA platform config claim must be present in a CCA platform token.

```
cca-platform-config-label = 2401 ; PSA platform range
                                ; TBD: add to IANA registration
cca-platform-config-type = bytes

cca-platform-config = (
    cca-platform-config-label => cca-platform-config-type
)
```

See also:

- [RME system architecture spec \[13\]](#)

A7.2.3.2.6 CCA platform lifecycle claim

<code>I_SYKFY</code>	The CCA platform lifecycle claim identifies the lifecycle state of the CCA platform.
<code>R_NBFVV</code>	The value of the CCA platform lifecycle claim is an integer which is divided as follows: <ul style="list-style-type: none"> • value[15:8]: CCA platform lifecycle state • value[7:0]: IMPLEMENTATION DEFINED
<code>I_WFZHV</code>	The CCA platform lifecycle claim must be present in a CCA platform token.
<code>I_QFYLF</code>	A non debugged CCA platform will be in psa-lifecycle-secured state. Realm Management Security Domain debug is always recoverable, and would therefore be represented by psa-lifecycle-non-psa-rot-debug state. Root world debug is recoverable on a HES system and would be represented by psa-lifecycle-recoverable-psa-rot state. On a non-HES system Root world debug is usually non-recoverable, and would be represented by psa-lifecycle-lifecycle-decommissioned state.
<code>I_HMZLL</code>	The format of the CCA platform lifecycle claim is defined as follows:

```
cca-platform-lifecycle-label = 2395 ; PSA lifecycle

cca-platform-lifecycle-unknown-type = 0x0000..0x00ff
cca-platform-lifecycle-assembly-and-test-type = 0x1000..0x10ff
cca-platform-lifecycle-cca-platform-rot-provisioning-type = 0x2000..0x20ff
cca-platform-lifecycle-secured-type = 0x3000..0x30ff
cca-platform-lifecycle-non-cca-platform-rot-debug-type = 0x4000..0x40ff
cca-platform-lifecycle-recoverable-cca-platform-rot-debug-type = 0x5000..0x50ff
cca-platform-lifecycle-decommissioned-type = 0x6000..0x60ff

cca-platform-lifecycle-type =
    cca-platform-lifecycle-unknown-type /
    cca-platform-lifecycle-assembly-and-test-type /
    cca-platform-lifecycle-cca-platform-rot-provisioning-type /
    cca-platform-lifecycle-secured-type /
    cca-platform-lifecycle-non-cca-platform-rot-debug-type /
    cca-platform-lifecycle-recoverable-cca-platform-rot-debug-type /
```

```
cca-platform-lifecycle-decommissioned-type

cca-platform-lifecycle = (
    cca-platform-lifecycle-label => cca-platform-lifecycle-type
)
```

See also:

- [Arm CCA Security model \[4\]](#)

A7.2.3.2.7 CCA platform software components claim

I_{PJCSC} The CCA platform software components claim is a list of software components which can affect the behavior of the CCA platform. It is expected that an implementation will describe the expected software component values within the profile.

U₀₀₅₇ In some implementations, a software component may consist of a configuration data item.

I_{TJTXG} The CCA platform software components claim must be present in a CCA platform token.

I_{DPSKT} The format of the CCA platform software components claim is defined as follows:

```
cca-platform-sw-components-label = 2399 ; PSA software components

cca-platform-sw-component = {
    ? 1 => text,                ; component type
    2 => cca-hash-type,         ; measurement value
    ? 4 => text,                ; version
    5 => cca-hash-type,         ; signer id
    ? 6 => text,                ; hash algorithm identifier
    ? 7 => bool,                ; live firmware activation supported
    ? 8 => [ + cca-hash-type ], ; list of countersigner ids
}

cca-platform-sw-components = (
    cca-platform-sw-components-label => [ + cca-platform-sw-component ]
)
```

CCA platform software component type

I_{PDNCF} The CCA platform software component type is a string which represents the role of the software component.

I_{TPSYF} The CCA platform software component type is intended for use as a hint to help the relying party understand how to evaluate the CCA platform software component measurement value.

R_{RSNBH} The CCA platform software component type is optional in a CCA platform token.

U₀₀₅₈ If the CCA platform supports Live Firmware Activation, one entry in the platform software component table is reserved to act as a measurement register for the Firmware Activity Log. This entry is identified by having a software component type of “FAL”.

See also:

- [A3.12 Live Firmware Activation](#)

CCA platform software component measurement value

I_{RWDKD} The CCA platform software component measurement value represents a hash of the state of the software component in memory at the time it was initialized.

R_{TVXRZ} The CCA platform software component measurement value must be a hash of 256 bits or stronger.

R_{LGBCM} The CCA platform software component measurement value must be present in a CCA platform token.

CCA platform software component version

I_{JVJFW} The CCA platform software component version is a text string whose meaning is defined by the software component vendor.

R_{CZRXB} The CCA platform software component version is optional in a CCA platform token.

CCA platform software component signer ID

I_{DCDMR} The CCA platform software component signer ID is the hash of a signing authority public key for the software component. It can be used by a verifier to ensure that the software component was signed by an expected trusted source.

R_{PXRMC} The CCA platform software component signer ID value must be a hash of 256 bits or stronger.

R_{XPHQC} The CCA platform software signer ID must be present in a CCA platform token.

CCA platform software component hash algorithm ID

I_{TQWZX} The CCA platform software component hash algorithm ID identifies the way in which the hash algorithm used to measure the CCA platform software component.

I_{HBBHG} Arm recommends that the value of the CCA platform software component hash algorithm ID is an IANA Hash Function name [IANA Named Information Hash Algorithm Registry](#) [11].

I_{NJYCM} Arm recommends that the hash algorithm used to measure the CCA platform software component is one of the algorithms listed in the [Arm CCA Security model](#) [4].

I_{HPHCD} The CCA platform software component hash algorithm ID is optional in a CCA platform token.

CCA platform software component Live Firmware Activation support

I₀₀₅₉ The CCA platform software component Live Firmware Activation support attribute declares whether an individual component is subject to Live Firmware Activation. If the attribute is False, the component will not be updated before the next CCA platform reset.

I₀₀₆₀ The CCA platform software component Live Firmware activation support attribute is optional in a CCA platform token.

See also:

- [A3.12 Live Firmware Activation](#)

CCA platform software component countersigner ID list

I₀₀₆₁ The CCA platform software component countersigner ID list contains hashes of public keys which identify signing authorities that provides additional trustworthiness information for the software component. These signatures are provided in addition to the primary signature, which is identified by the CCA platform software component signer ID.

U₀₀₆₂ Example use cases for CCA platform software component countersignatures include:

- An indication of approval for the component, provided by the owner of the CCA platform
- An indication of approval for the component, provided by a third party auditor

U₀₀₆₃ The order of multiple entries within the countersigner ID list may imply a hierarchy. The existence and meaning of any such hierarchy is IMPLEMENTATION DEFINED.

I₀₀₆₄ The CCA platform software component countersigner ID list is optional in a CCA platform token.

A7.2.3.2.8 CCA platform verification service claim

I_{NSTDP} The CCA platform verification service claim is a hint which can be used by a relying party to locate a verifier for the token.

I_{RZJSQ} The value of the CCA platform verification service claim is a text string which can be used to locate the service or a URL specifying the address of the service.

I_{MFYCX} The CCA platform verification service claim may be ignored by a relying party in favor of other information.

I_{MRSXY} The CCA platform verification service claim is optional in a CCA platform token.

I_{WRJSX} The format of the CCA platform verification service claim is defined as follows:

```
cca-platform-verification-service-label = 2400 ; PSA verification service
cca-platform-verification-service-type = text

cca-platform-verification-service = (
    cca-platform-verification-service-label =>
    cca-platform-verification-service-type
)
```

A7.2.3.2.9 CCA platform hash algorithm ID claim

I_{VDZMF} The CCA platform hash algorithm ID claim identifies the default algorithm used to calculate measurements in the CCA platform token.

I_{XHJFX} The default hash algorithm may be overridden for an individual software component, by the CCA platform software component hash algorithm ID claim.

I_{YRPYY} Arm recommends that the value of the CCA platform hash algorithm ID claim is an IANA Hash Function name [IANA Named Information Hash Algorithm Registry \[11\]](#).

I_{TQSTK} The CCA platform hash algorithm ID claim must be present in a CCA platform token.

I_{RKZJT} The format of the CCA platform hash algorithm ID claim is defined as follows:

```
cca-platform-hash-algo-id-label = 2402 ; PSA platform range
                                   ; TBD: add to IANA registration

cca-platform-hash-algo-id = (
    cca-platform-hash-algo-id-label => text
)
```

A7.2.3.3 Attestation token format compatibility

I₀₀₆₅ The table below summarises the support for each claim, depending on the version of the profile of the attestation token. The profile version is reported in the cca-platform-profile claim as <http://arm.com/CCA-SSD/VERSION>.

Claim	Version	
	1.0.0	1.1.0
cca-platform-profile	Required	Required
cca-platform-challenge	Required	Required
cca-platform-implementation-id	Required	Required
cca-platform-instance-id	Required	Required
cca-platform-config	Required	Required
cca-platform-lifecycle	Required	Required

Claim	Version	
	1.0.0	1.1.0
cca-platform-sw-components	Required: <ul style="list-style-type: none"> • Measurement value • Signer ID Optional: <ul style="list-style-type: none"> • Type • Version • Hash algorithm identifier 	Required <ul style="list-style-type: none"> • Measurement value • Signer ID Optional: <ul style="list-style-type: none"> • Type • Version • Hash algorithm identifier • LFA supported • Countersigner IDs
cca-platform-verification-service	Optional	Optional
cca-platform-hash-algo-id	Required	Required
cca-realm-challenge	Required	Required
cca-realm-personalization-value	Required	Required
cca-realm-initial-measurement	Required	Required
cca-realm-extensible-measurements	Required	Required
cca-realm-hash-algo-id	Required	Required
cca-realm-mec-policy	Not supported	Required
cca-realm-public-key	Required	Required
cca-realm-public-key-hash-algo-id	Required	Required

A7.2.3.3.1 Collated CDDL for CCA platform claims

D_{DVMJZ}

The format of the CCA platform token claim map is defined as follows:

```
cca-platform-claims = (cca-platform-claim-map)

cca-platform-claim-map = {
    cca-platform-profile
    cca-platform-challenge
    cca-platform-implementation-id
    cca-platform-instance-id
    cca-platform-config
    cca-platform-lifecycle
    cca-platform-sw-components
    ? cca-platform-verification-service
    cca-platform-hash-algo-id
}
cca-platform-profile-label = 265 ; EAT profile

cca-platform-profile-type = "tag:arm.com,2024:cca_platform#1.1.0"

cca-platform-profile = (
    cca-platform-profile-label => cca-platform-profile-type
)
cca-hash-type = bytes .size 32 / bytes .size 48 / bytes .size 64
cca-platform-challenge-label = 10

cca-platform-challenge = (
    cca-platform-challenge-label => cca-hash-type
)
cca-platform-implementation-id-label = 2396 ; PSA implementation ID
cca-platform-implementation-id-type = bytes .size 32

cca-platform-implementation-id = (
    cca-platform-implementation-id-label => cca-platform-implementation-id-type
)
cca-platform-instance-id-label = 256 ; EAT uuid

; TODO: require that the first byte of cca-platform-instance-id-type is 0x01
; EAT UEIDs need to be 7 - 33 bytes
cca-platform-instance-id-type = bytes .size 33

cca-platform-instance-id = (
    cca-platform-instance-id-label => cca-platform-instance-id-type
)
cca-platform-config-label = 2401 ; PSA platform range
                                ; TBD: add to IANA registration
cca-platform-config-type = bytes

cca-platform-config = (
    cca-platform-config-label => cca-platform-config-type
)
cca-platform-lifecycle-label = 2395 ; PSA lifecycle

cca-platform-lifecycle-unknown-type = 0x0000..0x00ff
cca-platform-lifecycle-assembly-and-test-type = 0x1000..0x10ff
cca-platform-lifecycle-cca-platform-rot-provisioning-type = 0x2000..0x20ff
cca-platform-lifecycle-secured-type = 0x3000..0x30ff
cca-platform-lifecycle-non-cca-platform-rot-debug-type = 0x4000..0x40ff
cca-platform-lifecycle-recoverable-cca-platform-rot-debug-type = 0x5000..0x50ff
cca-platform-lifecycle-decommissioned-type = 0x6000..0x60ff
```

```
cca-platform-lifecycle-type =  
    cca-platform-lifecycle-unknown-type /  
    cca-platform-lifecycle-assembly-and-test-type /  
    cca-platform-lifecycle-cca-platform-rot-provisioning-type /  
    cca-platform-lifecycle-secured-type /  
    cca-platform-lifecycle-non-cca-platform-rot-debug-type /  
    cca-platform-lifecycle-recoverable-cca-platform-rot-debug-type /  
    cca-platform-lifecycle-decommissioned-type  
  
cca-platform-lifecycle = (  
    cca-platform-lifecycle-label => cca-platform-lifecycle-type  
)  
cca-platform-sw-components-label = 2399 ; PSA software components  
  
cca-platform-sw-component = {  
    ? 1 => text,                ; component type  
    ? 2 => cca-hash-type,        ; measurement value  
    ? 4 => text,                ; version  
    ? 5 => cca-hash-type,        ; signer id  
    ? 6 => text,                ; hash algorithm identifier  
    ? 7 => bool,                ; live firmware activation supported  
    ? 8 => [ + cca-hash-type ], ; list of countersigner ids  
}  
  
cca-platform-sw-components = (  
    cca-platform-sw-components-label => [ + cca-platform-sw-component ]  
)  
cca-platform-verification-service-label = 2400 ; PSA verification service  
cca-platform-verification-service-type = text  
  
cca-platform-verification-service = (  
    cca-platform-verification-service-label =>  
    cca-platform-verification-service-type  
)  
cca-platform-hash-algo-id-label = 2402 ; PSA platform range  
                                     ; TBD: add to IANA registration  
  
cca-platform-hash-algo-id = (  
    cca-platform-hash-algo-id-label => text  
)
```

A7.2.3.3.2 Example CCA platform claims

I_{TVHKL}

An example CCA platform claim map is shown below in COSE-DIAG format:

```
/ CCA platform claim map /
{
  / cca-platform-profile /
  265: "tag:arm.com,2024:cca_platform#1.1.0",

  / cca-platform-challenge /
  10: h'AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
      AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA',

  / cca-platform-implementation-id /
  2396: h'AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA',

  / cca-platform-instance-id /
  256: h'010BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
      BB',

  / cca-platform-config /
  2401: h'CFCFCFCF',

  / cca-platform-lifecycle /
  2395: 12288,

  / cca-platform-sw-components /
  2399: [
    {
      / measurement value /
      2: h'AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
          AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA',

      / signer id /
      5: h'BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
          BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB',

      / version /
      4: "1.0.0",

      / hash algorithm identifier /
      6: "sha-256"
    },
    {
      / measurement value /
      2: h'CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
          CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC',

      / signer id /
      5: h'DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
          DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD',

      / version /
      4: "1.0.0",

      / hash algorithm identifier /
      6: "sha-256"
    }
  ],

  / cca-platform-verification-service /
```

```
2400: "https://cca_verifier.org",  
  
/ cca-platform-hash-algo-id /  
2402: "sha-256"  
}
```

DRAFT

Chapter A8

Realm debug and performance monitoring

This section describes the debug and performance monitoring features which are available to a Realm.

A8.1 Realm PMU

This section describes the programming model for usage of PMU by a Realm.

R_{DNNQO}

On REC entry, Realm PMU state is restored from the REC object.

R_{LHRYJ}

On REC exit, all Realm PMU state is saved to the REC object.

R_{WXTZF}

On REC exit, `exit.pmu_ovf_status` indicates the status of the PMU overflow at the time of the Realm exit.

See also:

- [A3.6 Realm support for Performance Monitors Extension](#)
- [A4.3 REC exit](#)
- [B4.4.35 RmiRecExit type](#)

DRAFT

Chapter A9

Realm device assignment

This section describes how devices are assigned to Realms, attested and granted permission to access Realm-owned memory.

A9.1 Realm device assignment overview

- I₀₀₆₆ The RMM allows a device to be assigned to a Realm in a trustworthy manner, allowing the Realm to attest the identity and configuration of the device before it is permitted to access the Realm's memory.
- I₀₀₆₇ From the Host point of view, devices are managed using two types of RMM object:
- *Physical Device* (PDEV)
Represents a communication channel between the RMM and a physical device, for example a PCIe device.
 - *Virtual Device* (VDEV)
Represents the binding between a device function and a Realm. For example, a VDEV can represent a physical function of a PCIe device or a virtual function of a multi-function PCIe device. Every VDEV is associated with one PDEV.
- Interaction between the Host and a PDEV object or VDEV object is performed via RMI commands.
- I₀₀₆₈ From the Realm point of view, an assigned device function is represented by a *Realm Device* (RDEV).
Interaction between the Realm and an RDEV is performed via RSI commands.
- U₀₀₆₉ Arm expects that a VDEV and an RDEV will both refer to the same data structure inside the RMM.
- D₀₀₇₀ A *platform-attested device* is a device whose identity and firmware are attested as part of the CCA platform.
- I₀₀₇₁ A device which is physically integrated into the platform and which does not implement its own Device Security Manager (DSM) is a platform-attested device.
- D₀₀₇₂ An *independently-attested device* is a device whose identity and firmware are attested separately from the CCA platform.
- I₀₀₇₃ An independently-attested device implements its own DSM.
See also:
- [A3.10 Support for Realm device assignment](#)
 - [A9.3 Physical device object](#)
 - [A9.4 Virtual device object](#)

A9.1.1 Assignment of an independently-attested device

- I₀₀₇₄ The communication channel between the RMM and an independently-attested device is managed by the Host.
- I₀₀₇₅ Communication between the RMM and an independently-attested device is protected using the Security Protocol and Data Model (SPDM) protocol.
- I₀₀₇₆ Identify of an independently-attested device is described by a device certificate.
- I₀₀₇₇ Firmware measurements for an independently-attested device are reported in an SPDM measurement block.
- I₀₀₇₈ Assignment of an independently-attested device to a Realm involves the following steps:
1. The Host creates a PDEV object, associated with the target physical device.
 2. The Host initializes the PDEV object, causing the following to happen:
 - A secure channel is established between the RMM and the device.
 - The physical link between the device and memory is secured. For example, for an off-chip PCIe device, this is achieved using the Integrity and Data Encryption (IDE) standard. See [PCI Express 6.0 specification](#) [14].
 - The device certificate is provided to the Host. The Host is expected to store this information, and to later present it to the Realm.

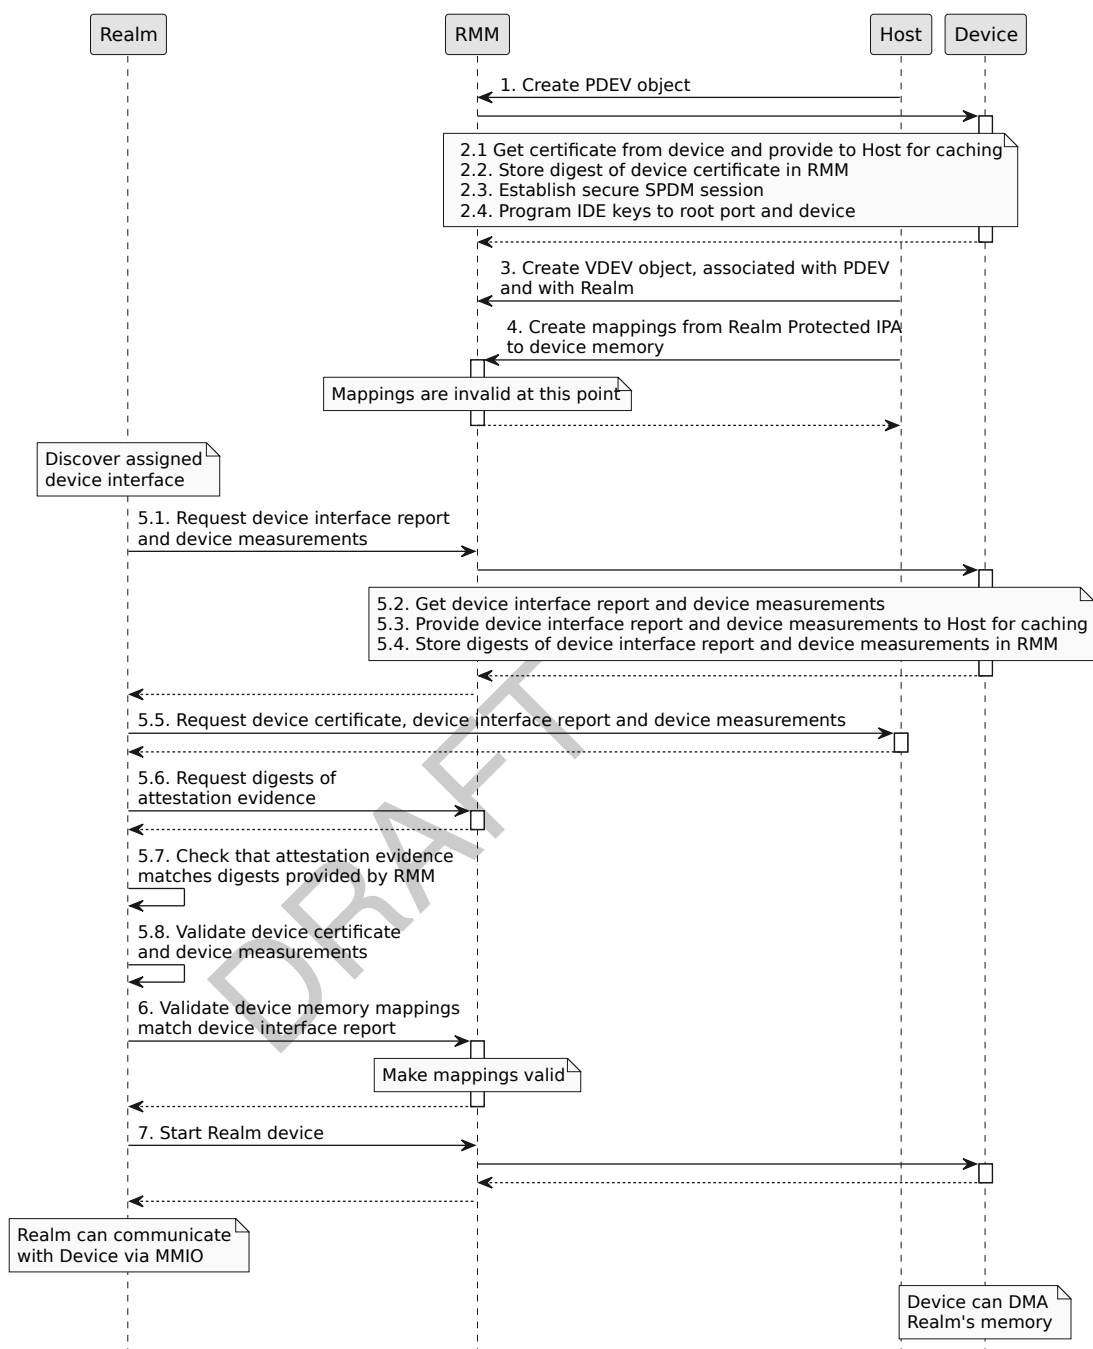
- A digest of the device certificate is stored by the RMM. This is used later to check integrity of the attestation evidence provided by the Host to the Realm.
3. The Host creates a VDEV object, which represents a binding between a function of the target device, and a Realm. At this stage, the target device is not granted access to the Realm-owned memory.
 4. The Host maps memory regions of the target device function into the Protected IPA space of the Realm. At this stage, the mappings are invalid, so the Realm cannot yet access the device's memory regions.
 5. Device information provided by the RMM tells the Realm that the device is independently-attested. The Realm requests the RMM to retrieve device attestation evidence (device interface report and device measurements.) As with the device certificate, the evidence is provided to the Host for caching, and the RMM stores digests of the evidence. The Realm later requests device attestation evidence from the Host, and verifies that this matches the corresponding digests stored by the RMM.
 6. The Realm verifies that the device identity (represented by the device certificate) and device measurements are acceptable.
 7. The Realm verifies that device memory mappings created by the Host match those described in the device interface report. Once each mapping has been verified, it is made valid by the RMM, so the Realm can access the device's memory regions.
 8. The Realm instructs the RMM to grant the device access to Realm-owned memory.

X₀₀₇₉

Requiring the Host to store device attestation evidence means that storage for this information, whose size may not be known ahead of time, does not need to be allocated in RMM memory. The RMM therefore only needs to store a digest of the Host-cached data, which can be used by the Realm to check integrity of data retrieved from the Host.

I₀₀₈₀

Assignment of an independently-attested device to a Realm is illustrated in the following sequence diagram.



Note that “secure SPDM” means that the requester (the RMM) has verified that the responder holds the private key which was used to sign the device certificate. The identity and trustworthiness of the responder are not evaluated until the Realm receives attestation evidence for the device.

See also:

- [Secured Messages using SPDM Specification version 1.1.0 \[15\]](#)

A9.1.2 Assignment of a platform-attested device

I₀₀₈₁

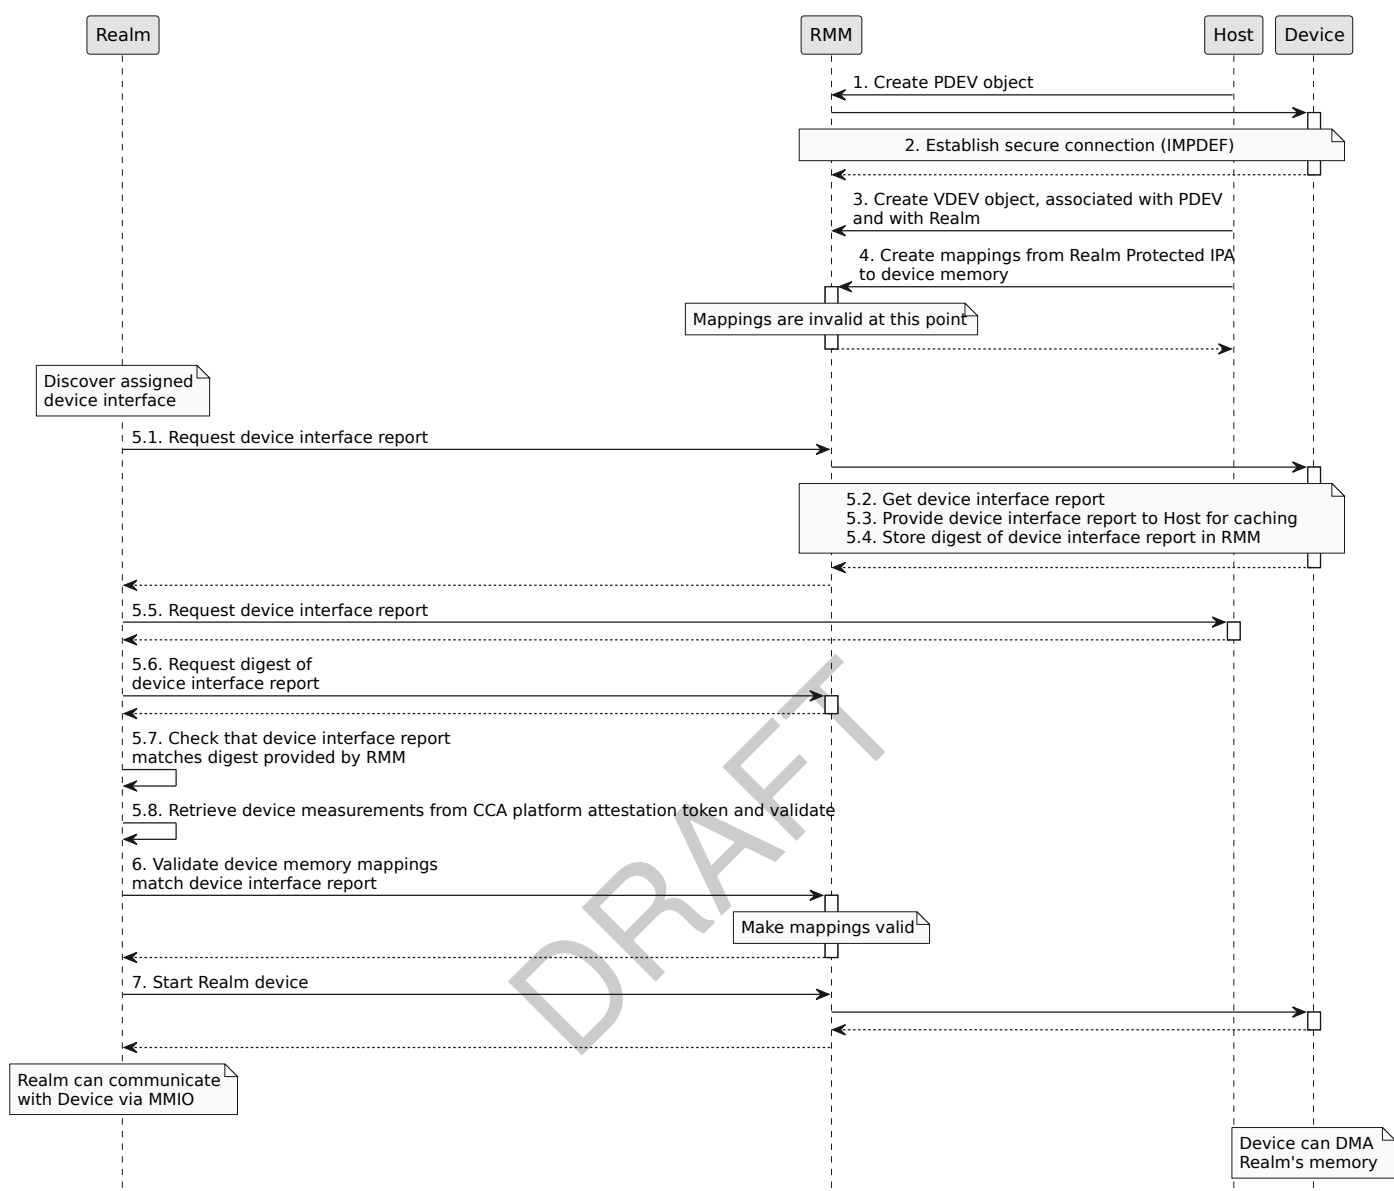
The communication channel between the RMM and a platform-attested device is IMPLEMENTATION DEFINED.

I₀₀₈₂ Firmware measurements for a platform-attested device are reported in the CCA platform software components claim.

I₀₀₈₃ Assignment of a platform-attested device to a Realm involves the following steps:

1. The Host creates a PDEV object, associated with the target physical device.
2. The Host initializes the PDEV object, causing the following to happen:
 - A secure channel is established between the RMM and the device.
3. The Host creates a VDEV object, which represents a binding between a function of the target device, and a Realm. At this stage, the target device is not granted access to the Realm-owned memory.
4. The Host maps memory regions of the target device function into the Protected IPA space of the Realm. At this stage, the mappings are invalid, so the Realm cannot yet access the device's memory regions.
5. Device information provided by the RMM tells the Realm that the device is platform-attested. The Realm requests the RMM to retrieve a device interface report. The device interface report is provided to the Host for caching, and the RMM stores digests of the device interface report. The Realm later requests the device interface report from the Host, and verifies that this matches the corresponding digest stored by the RMM.
6. The Realm retrieves device measurements from the CCA platform attestation token, and verifies that the measurements are acceptable.
7. The Realm verifies that device memory mappings created by the Host match those described in the device interface report. Once each mapping has been verified, it is made valid by the RMM, so the Realm can access the device's memory regions.
8. The Realm instructs the RMM to grant the device access to Realm-owned memory.

I₀₀₈₄ Assignment of a platform-attested device to a Realm is illustrated in the following sequence diagram.



See also:

- [A7.2.3.2.7 CCA platform software components claim](#)

A9.2 Communication between RMM and a device

A9.2.1 Device requests and responses

D₀₀₈₅ Communication between the RMM and a device consists of a series of *device requests* sent from the RMM to the device and *device responses* returned by the device to the RMM.

D₀₀₈₆ A *device transaction* is a series of one or more (device request, device response) tuples.

I₀₀₈₇ At the requester side (that is, at the RMM), a device transaction is associated with either a PDEV or a VDEV, depending on the event which triggered the device transaction:

- If the device transaction was triggered by one of the following RMI commands then the device transaction is associated with the PDEV.
 - RMI_PDEV_CREATE
 - RMI_PDEV_SET_PUBKEY
 - RMI_PDEV_NOTIFY
 - RMI_PDEV_STOP
- If the device transaction was triggered by one of the following RMI commands then the device transaction is associated with the VDEV.
 - RMI_VDEV_STOP
- If the device transaction was triggered by one of the following RSI commands then the device transaction is associated with the VDEV.
 - RSI_RDEV_GET_INTERFACE_REPORT
 - RSI_RDEV_GET_MEASUREMENTS
 - RSI_RDEV_LOCK
 - RSI_RDEV_START
 - RSI_RDEV_STOP

I₀₀₈₈ A PDEV is associated with at most one device transaction at a time.

I₀₀₈₉ A VDEV is associated with at most one device transaction at a time.

I₀₀₉₀ Communication between the RMM and an off-chip device consists of SPDm messages, transported via Non-secure memory.

I₀₀₉₁ Communication between the RMM and an on-chip device can occur via one of two paths:

- SPDm messages, transported via Non-secure memory.
- Platform communication, via an IMPLEMENTATION DEFINED channel.

I₀₀₉₂ When the RMM requires the Host to send a device request, the “send” flag is set in the DevCommExitFlags fieldset.

I₀₀₉₃ When the RMM sends a device request via an IMPLEMENTATION DEFINED channel, the “send” flag is clear and the “wait” flag is set in the DevCommExitFlags fieldset. This informs the Host that it should either:

- Register for an IMPLEMENTATION DEFINED notification that the request has been completed, and call RMI_PDEV_COMMUNICATE or RMI_VDEV_COMMUNICATE when this notification is received, or
- Poll for request completion by periodically calling RMI_PDEV_COMMUNICATE or RMI_VDEV_COMMUNICATE.

D₀₀₉₄ The states of a Device communication are listed below.

State	Description
DEV_COMM_IDLE	The RMM is not communicating with the device.
DEV_COMM_PENDING	The RMM has a device request which is ready to be sent to the device.

State	Description
DEV_COMM_ACTIVE	The RMM has initiated a device transaction. One or more device requests associated with this device transaction have been sent from the RMM to the device. The RMM has not received all the expected device responses associated with this device transaction.
DEV_COMM_ERROR	The RMM encountered an error during communication with the device.

I₀₀₉₅

Device communication state transitions are shown in the following figure. Each arc is labeled with the events which can cause the corresponding state transition.

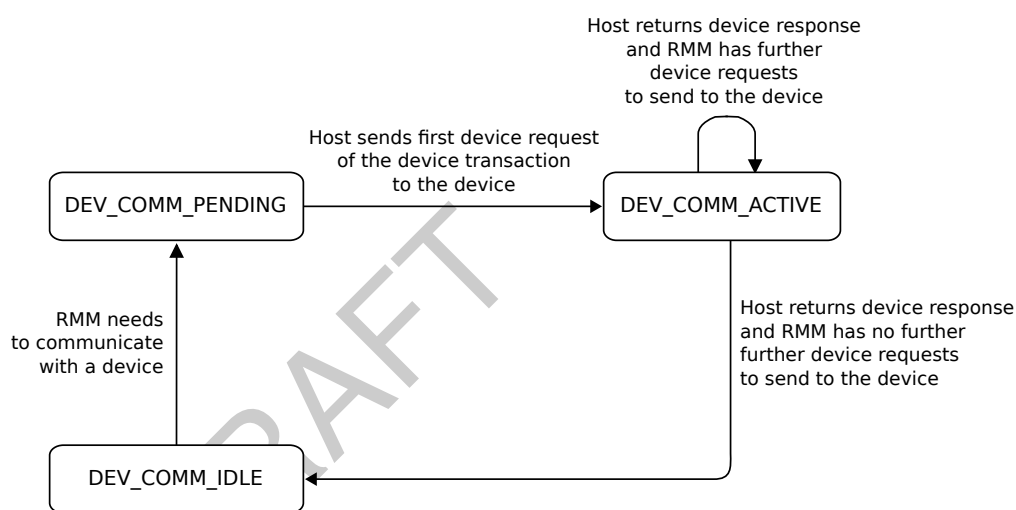


Figure A9.1: Device communication state transitions

See also:

- [Secured Messages using SPDm Specification version 1.1.0 \[15\]](#)
- [A9.5.1 Interruptible Realm device operations](#)
- [B4.4.9 RmiDevCommExitFlags type](#)

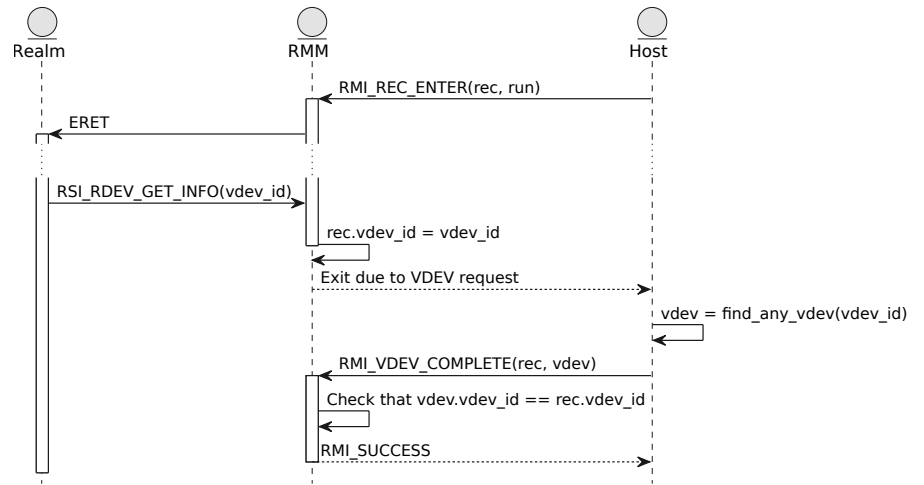
A9.2.2 Mapping from virtual device ID to VDEV object

I₀₀₉₆

The RMM does not maintain a mapping from virtual device ID to VDEV object. Consequently, when a Realm executes the RSI_RDEV_GET_INFO command, passing a virtual device ID, the RMM must request the Host to provide a corresponding VDEV object.

I₀₀₉₇

The following sequence diagram shows how the virtual device ID passed by a Realm is mapped to a corresponding VDEV object.



See also:

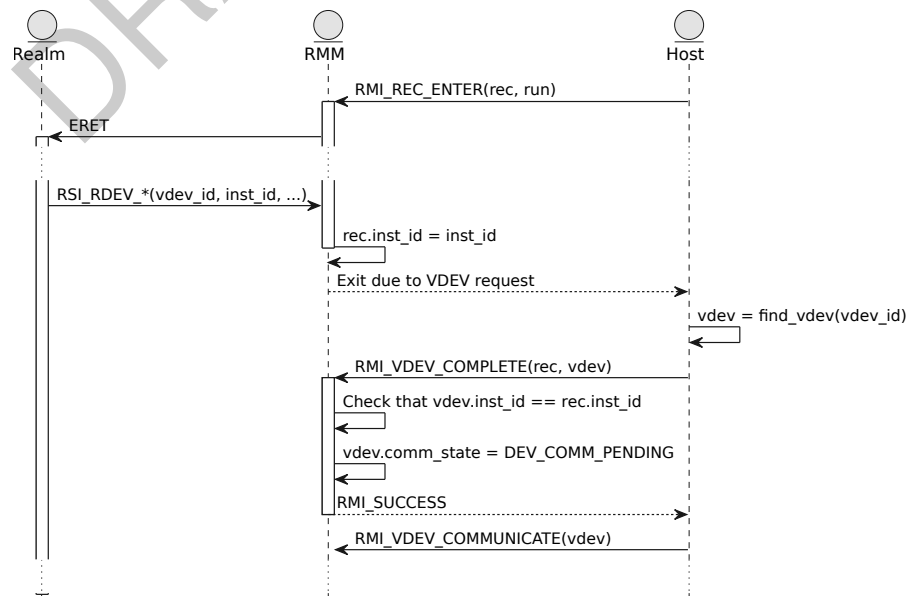
- [A4.3.14 REC exit due to VDEV request](#)
- [B4.3.50 RMI_VDEV_COMPLETE command](#)
- [B5.3.16 RSI_RDEV_GET_INFO command](#)

I₀₀₉₈

The data structure returned from RSI_RDEV_GET_INFO includes a “device instance ID” value. The device instance ID is guaranteed to be unique among VDEVs owned by the calling Realm, unlike the virtual device ID. When the Realm calls any RSI_RDEV command other than RSI_RDEV_GET_INFO, it passes the (virtual device ID, instance ID) tuple.

I₀₀₉₉

The following sequence diagram shows how the (virtual device ID, instance ID) tuple passed by a Realm is mapped to the corresponding VDEV object.



See also:

- [A4.3.14 REC exit due to VDEV request](#)
- [B4.3.50 RMI_VDEV_COMPLETE command](#)
- [B5.3.15 RSI_RDEV_CONTINUE command](#)
- [B5.3.17 RSI_RDEV_GET_INTERFACE_REPORT command](#)
- [B5.3.18 RSI_RDEV_GET_MEASUREMENTS command](#)

- [B5.3.19 RSI_RDEV_GET_STATE command](#)
- [B5.3.20 RSI_RDEV_LOCK command](#)
- [B5.3.21 RSI_RDEV_START command](#)
- [B5.3.22 RSI_RDEV_STOP command](#)
- [B5.3.23 RSI_RDEV_VALIDATE_MAPPING command](#)

A9.2.3 Host-side device communication flow

- I₀₁₀₀ The RMI_PDEV_COMMUNICATE command is used to send a PDEV-associated device request from the RMM to a device, and / or to return a device response from the device to the RMM.
- I₀₁₀₁ The RMI_VDEV_COMMUNICATE command is used to send a VDEV-associated device request from the RMM to a device, and / or to return a device response from the device to the RMM.
- I₀₁₀₂ The RMI_PDEV_COMMUNICATE and RMI_VDEV_COMMUNICATE commands have identical programming models. Hereafter, they are referred to collectively as “device communication commands”.
- R₀₁₀₃ For a given physical device, at most one device transaction can be active.
- I₀₁₀₄ The RMI_PDEV_ABORT command or RMI_VDEV_ABORT command is used to abort an DEV_COMM_ACTIVE device transaction.
- I₀₁₀₅ At the responder side (that is, at the device), device transactions associated with a PDEV and device transactions associated with its child VDEVs all terminate at the same physical device.

A9.2.3.1 Communication flow for devices which use SPD

- I₀₁₀₆ The overall flow for communication between the RMM and a device which uses SPD is as follows:
1. The RMM indicates to the Host that a device transaction, associated with a specified PDEV or VDEV, is DEV_COMM_PENDING.
 - If the device transaction was triggered by execution of an RMI command, this indication is provided via the command’s output values.
 - If the device transaction was triggered by execution of an RSI command, this indication is provided via a REC exit due to device communication.

The RMM also indicates to the Host whether the pending transaction will contain more than one (device request, device response) tuple.
 2. The Host calls the appropriate device communication command.

Input values to the command include a *device request buffer*, which is a pointer to a Granule of NS memory.
 3. The RMM writes a device request to the device request buffer and returns to the Host, indicating that data is ready to be sent to the device. The device request is guaranteed to be no larger than 4KB.
 4. The Host sends the device request to the device.

Details of how this is performed are out of scope of this specification. For an off-chip PCIe device, this could be done by copying from the device request buffer to a Data Object Exchange (DOE) mailbox.

The device communication state becomes DEV_COMM_ACTIVE.
 5. The device indicates to the Host that it has responded to the request.

The Host copies the device response to a *device response buffer* in NS memory. As with sending the request, details of how this are out of scope of this specification.
 6. The Host calls the device communication command, providing a pointer to the device response buffer.
 7. The return value indicates that either:
 - The RMM has another device request to send within the same device transaction (in which case the flow returns to step 2), or

- The device transaction is complete.

8. If the device transaction is incomplete, the Host checks the device state to determine whether an error has occurred.

I₀₁₀₇ When multiple device transactions destined for the same physical device are DEV_COMM_PENDING, the Host is free to choose which of them to transition to DEV_COMM_ACTIVE.

I₀₁₀₈ When a device transaction is DEV_COMM_ACTIVE, the Host must not send to that physical device a device request associated with any other device transaction.

U₀₁₀₉ The SPDM session which is established between the RMM and a device has the following characteristics:

- SPDM measurements use DMTF format
- SPDM heartbeat is not supported
- SPDM key update is not supported

U₀₁₁₀ In order for its functions to be assignable to Realms, a device must provide the following functionality:

- SPDM version required by PCIe TDISP and IDE_KM specifications.
- Identity and authentication including key exchange.

See also:

- [PCI Express 6.0 specification \[14\]](#)
- [A4.3.11 REC exit due to device communication](#)

A9.2.3.2 Communication flow for devices which do not use SPDM

I₀₁₁₁ The overall flow for communication between the RMM and a device which does not use SPDM is as follows:

1. The RMM indicates to the Host that a device transaction, associated with a specified PDEV or VDEV, is DEV_COMM_PENDING.
 - If the device transaction was triggered by execution of an RMI command, this indication is provided via the command's output values.
 - If the device transaction was triggered by execution of an RSI command, this indication is provided via a REC exit due to device communication.
2. The Host calls the appropriate device communication command repeatedly, until either:
 - The return value indicates that the device transaction is complete.
 - The device enters an error state.

See also:

- [A4.3.11 REC exit due to device communication](#)

A9.2.4 Host caching of device attestation evidence

I₀₁₁₂ On execution of a device communication command, the RMM can indicate to the Host that the Host should cache data from the response buffer, for later retrieval by the Realm.

The identity of the data which the Host is requested to cache is implied by the command which triggered the device transaction, as follows:

- If the device transaction was triggered while the PDEV state was PDEV_NEW then the cached data is a device certificate.
- If the device transaction was triggered by RSI_RDEV_GET_INTERFACE_REPORT then the cached data is a device interface report.
- If the device transaction was triggered by RSI_RDEV_GET_MEASUREMENTS then the cached data is a device measurement block.

See also:

- [A4.3.11 REC exit due to device communication](#)
- [A9.5.2 Realm retrieval of device attestation evidence](#)
- [B4.3.16 RMI_PDEV_CREATE command](#)
- [B5.3.17 RSI_RDEV_GET_INTERFACE_REPORT command](#)
- [B5.3.18 RSI_RDEV_GET_MEASUREMENTS command](#)

A9.2.5 Device communication data structures

A9.2.5.1 Device communication exit data structure

D₀₁₁₃ An *RmiDevCommExit* object is a data structure which is passed from the RMM to the Host during a device transaction.

I₀₁₁₄ The attributes of an *RmiDevCommExit* object tell the Host the following:

- Whether there is data in the device response buffer which the Host is requested to cache. This is indicated by the *RmiDevCommExitFlags::cache* flag.
- Whether the device request buffer contains a device request which the Host is requested to send to the device. This is indicated by the *RmiDevCommExitFlags::send* flag.
- Whether the RMM is waiting for a response from the device. This is indicated by the *RmiDevCommExitFlags::wait* flag.
- Whether the device transaction contains more than one (device request, device response) tuple. This is indicated by the *RmiDevCommExitFlags::multi* flag.

I₀₁₁₅ *RmiDevCommExitFlags::send* and *RmiDevCommExitFlags::wait* are never set together.

I₀₁₁₆ *RmiDevCommExitFlags::multi* is only set when *RmiDevCommExitFlags::send* is set.

I₀₁₁₇ During communication between the RMM and a device which uses SPDH:

- *RmiDevCommExitFlags::send* indicates that the Host is requested to send a device request to the device.
- *RmiDevCommExitFlags::wait* indicates that the Host is requested to return a device response to the RMM.

I₀₁₁₈ During communication between the RMM and a device which does not use SPDH:

- *RmiDevCommExitFlags::wait* indicates that the Host is requested to notify the RMM when a device response is available.

D₀₁₁₉ The attributes of an *RmiDevCommExit* object are summarized in the following table.

Name	Byte offset	Type	Description
flags	0x0	RmiDevCommExitFlags	Flags indicating action(s) which the Host is requested to perform
cache_offset	0x8	UInt64	If flags.cache is true, offset in the device response buffer to the start of data to be cached, in bytes
cache_len	0x10	UInt64	If flags.cache is true, amount of data to be cached, in bytes
protocol	0x18	RmiDevCommProtocol	If flags.send is true, protocol to use
req_len	0x20	UInt64	If flags.send is true, amount of valid data in request buffer in bytes
timeout	0x28	UInt64	If flags.wait is true, amount of time to wait for device response in milliseconds

See also:

- [B4.4.8 RmiDevCommExit type](#)
- [B4.4.9 RmiDevCommExitFlags type](#)

A9.2.5.2 Device communication enter data structure

D₀₁₂₀ An *RmiDevCommEnter* object is a data structure which is passed from the Host to the RMM during a device transaction.

I₀₁₂₁ The attributes of an *RmiDevCommEnter* object tell the RMM the following:

- Whether the device reported an error.
- Whether the device response buffer contains a device response.
- The location of the device request buffer, into which the RMM can write a device request.

D₀₁₂₂ The attributes of an *RmiDevCommEnter* object are summarized in the following table.

Name	Byte offset	Type	Description
status	0x0	RmiDevCommStatus	Status of device transaction
req_addr	0x8	Address	Address of request buffer
resp_addr	0x10	Address	Address of response buffer
resp_len	0x18	UInt64	Amount of valid data in response buffer in bytes

See also:

- [B4.4.7 RmiDevCommEnter type](#)

A9.3 Physical device object

D₀₁₂₃ A *Physical Device* (PDEV) represents a communication channel between the RMM and a physical device, for example a PCIe device.

A9.3.1 Physical device attributes

D₀₁₂₄ The attributes of a PDEV are summarized in the following table.

Name	Type	Description
pdev_id	Bits64	Device identifier
prot_config	RmmPdevProtConfig	Configuration of protection between system and device
segment_id	Bits16	Segment identifier PCIe Segment identifier of the Root Port and endpoint.
root_id	Bits16	Root Port identifier Physical PCIe routing identifier of the Root Port to which the endpoint is connected.
cert_id	UInt64	Certificate identifier
rid_base	UInt64	Base of requester ID range (inclusive)
rid_top	UInt64	Top of requester ID range (exclusive)
hash_algo	RmmHashAlgorithm	Algorithm used to generate device digests
ide_sid	UInt64	IDE stream ID
iocoh_num_addr_range	UInt64	Number of IO-coherent address ranges
iocoh_addr_range	RmmAddressRange[16]	IO-coherent address range
fcoh_num_addr_range	UInt64	Number of fully-coherent address ranges
fcoh_addr_range	RmmAddressRange[4]	Fully-coherent address range
aux	Address[32]	Addresses of auxiliary Granules
num_aux	UInt64	Number of auxiliary Granules
state	RmmPdevState	Lifecycle state
comm_state	RmmDevCommState	Device communication state
num_vdevs	UInt64	Number of VDEVs associated with this PDEV whose state is not VDEV_STOPPED

I₀₁₂₅ pdev.root_id and pdev.cert_id attributes are ignored for platform-attested devices.

R₀₁₂₆ If pdev.prot_config is PDEV_IOCOH_E2E_IDE or PDEV_FCOH_E2E_IDE then all of the following are true:

- The RMM checks that pdev.ide_sid is not used by any other PDEV which is connected to the same Root Port.
- The RMM configures an IDE selective stream with IDE selective stream ID pdev.ide_sid and IDE selective stream address ranges pdev.iocoh_addr_range.

R₀₁₂₇ The RMM checks that all entries in (pdev.rid_base, pdev.rid_top] are not used by any other PDEV within the same PCIe segment.

R₀₁₂₈ The RMM checks that all entries in pdev.iocoh_addr_range have the following properties:

- Fall within non-coherent IO ranges of the system address map.
- Not used by any other PDEV.

R₀₁₂₉

If pdev.prot_config is PDEV_FCOH_E2E_IDE or PDEV_FCOH_E2E_SYS then the RMM checks that all entries in pdev.fcoh_addr_range have the following properties:

- Fall within coherent IO ranges of the system address map.
- Not used by any other PDEV.

I₀₁₃₀

The following table summarises which PDEV attributes are valid for each value of pdev.prot_config.

prot_config	Device type	ide_sid	(rid_base, rid_top]	iocoh_addr_range	fcoh_addr_range
PDEV_IOCOH_E2E_IDE	PCIe off-chip	Y	Y	Y	
PDEV_IOCOH_E2E_SYS	PCIe on-chip		Y	Y	
PDEV_FCOH_E2E_IDE	CHI off-chip	Y	Y	Y	Y
PDEV_FCOH_E2E_SYS	CHI on-chip		Y	Y	Y

A9.3.2 Physical device lifecycle

A9.3.2.1 States

D₀₁₃₁

The states of a PDEV are listed below.

State	Description
PDEV_NEW	Initial state of the device.
PDEV_NEEDS_KEY	RMM needs device public key.
PDEV_HAS_KEY	RMM has device public key.
PDEV_READY	Secure connection between the RMM and the device has been established. Physical link between the device and memory is secured. Ready for creation of VDEV instances.
PDEV_IDE_RESETTING	The PDEV's IDE link is being reset.
PDEV_COMMUNICATING	The RMM is communicating with the device.
PDEV_STOPPING	The RMM is communicating with the device to terminate the secure connection between the RMM and the device.
PDEV_STOPPED	Secure connection between the RMM and the device has been terminated.
PDEV_ERROR	Device has reported a fatal error.

A9.3.2.2 State transitions

I₀₁₃₂

Permitted PDEV Realm state transitions are shown in the following figure. Each arc is labeled with the events which can cause the corresponding state transition.

A transition from the pseudo-state *NULL* represents creation of a PDEV. A transition to the pseudo-state *NULL* represents destruction of a PDEV.

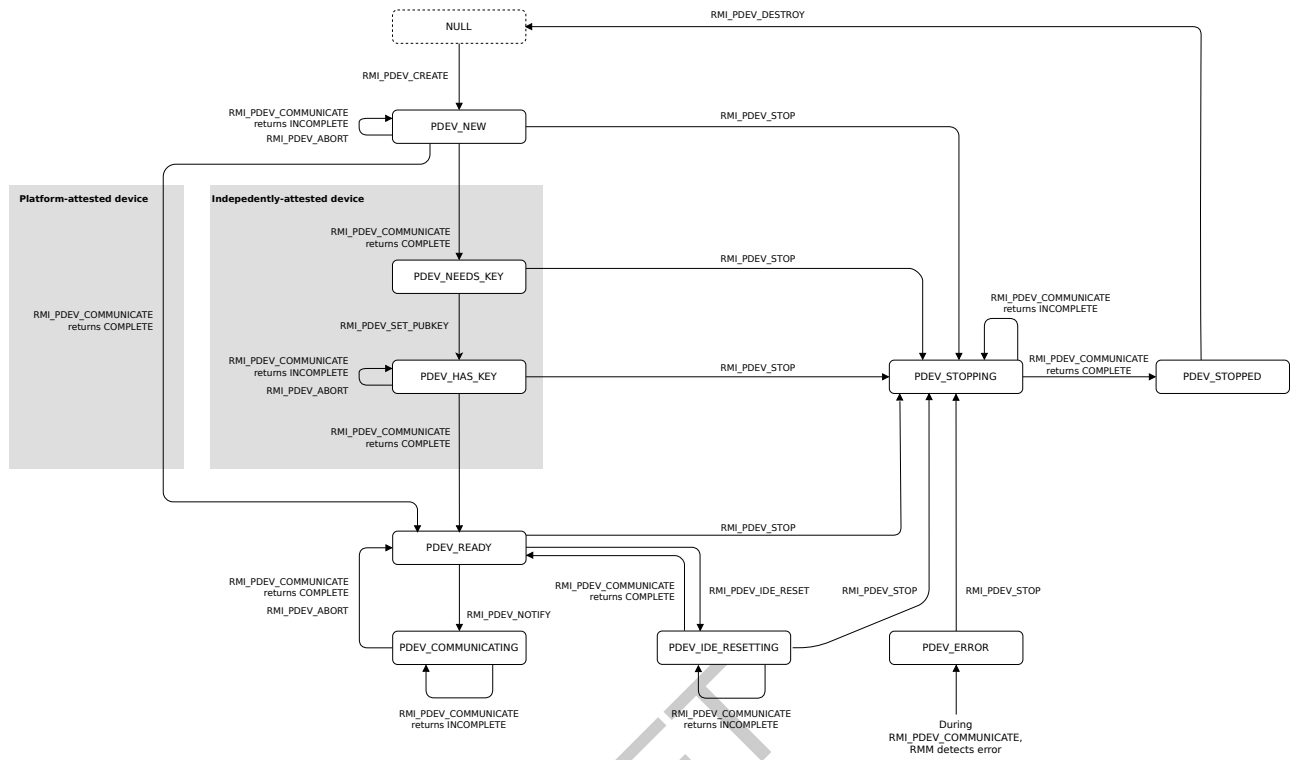


Figure A9.2: PDEV state transitions

A9.3.2.2.1 State transitions for an independently-attested device

I₀₁₃₃ While a PDEV associated with an independently-attested device is in PDEV_NEW state, execution of RMI_PDEV_COMMUNICATE causes the RMM to fetch the device certificate. The RMM stores a digest of the device certificate for later retrieval by the Realm. The Host is expected to cache the device certificate for later retrieval by the Realm.

Once device certificate retrieval is complete, the PDEV moves to PDEV_NEEDS_KEY state.

I₀₁₃₄ While a PDEV associated with an independently-attested device is in PDEV_HAS_KEY state, execution of RMI_PDEV_COMMUNICATE causes the RMM to perform secure SPDM session establishment and IDE key programming.

Once secure SPDM session establishment and IDE key programming is complete, the PDEV moves to PDEV_READY state.

A9.3.2.2.2 State transitions for a platform-attested device

I₀₁₃₅ While a PDEV associated with a platform-attested device is in PDEV_NEW state, execution of RMI_PDEV_COMMUNICATE causes the RMM to establish a communication channel with the device.

Once device certificate retrieval is complete, the PDEV moves to PDEV_READY state.

A9.3.2.2.3 State transitions for devices which use IDE

I₀₁₃₆ While a PDEV is in PDEV_READY state, the Host can notify the RMM of an event relevant to the device by executing RMI_PDEV_NOTIFY.

If the return value is RMI_DEV_COMM_PENDING then the PDEV moves to PDEV_COMMUNICATING state.

I₀₁₃₇ While a PDEV is in PDEV_READY state, if the device class is PCIe then the Host can request the device's IDE link to be reset by executing RMI_PDEV_IDE_RESET.

The PDEV moves to PDEV_IDE_RESETTING state.

I₀₁₃₈ While a PDEV is in PDEV_IDE_RESETTING state, the Host can enact the “IDE reset” device transaction by executing RMI_PDEV_COMMUNICATE.

If the return value indicates that the device transaction is complete then the PDEV moves to PDEV_READY state.

A9.3.2.2.4 State transitions for all devices

I₀₁₃₉ While a PDEV is in PDEV_COMMUNICATING state, the Host can either:

- Transfer device requests and device responses by executing RMI_PDEV_COMMUNICATE. If the return value indicates that the device transaction is complete then the PDEV moves to PDEV_READY state.
- Abort the device transaction by executing RMI_PDEV_ABORT. On successful execution of this command, the PDEV moves to PDEV_READY state.

I₀₁₄₀ On execution of RMI_PDEV_COMMUNICATE, if the RMM detects a fatal error (such as an unexpected response or a protocol error) then the PDEV moves to PDEV_ERROR state.

I₀₁₄₁ While a PDEV is in any of the following state, the Host can request the RMM to stop the device by executing RMI_PDEV_STOP:

- PDEV_NEW
- PDEV_NEEDS_KEY
- PDEV_HAS_KEY
- PDEV_READY
- PDEV_IDE_RESETTING
- PDEV_ERROR

The PDEV moves to PDEV_STOPPING state.

I₀₁₄₂ While a PDEV is in PDEV_STOPPING state, the Host can enact the “stop” device transaction by executing RMI_PDEV_COMMUNICATE.

If the return value indicates that the device transaction is complete then the PDEV moves to PDEV_STOPPED state.

I₀₁₄₃ While a PDEV is in PDEV_STOPPING state, if the device fails to respond or reports an error then the Host can call RMI_PDEV_COMMUNICATE, passing RMI_DEV_COMM_ERROR.

On successful execution of this command, the PDEV moves to PDEV_STOPPED state.

I₀₁₄₄ While a PDEV is in PDEV_STOPPED state, the Host can reclaim resources by executing RMI_PDEV_DESTROY. This command will fail if the PDEV is associated with any VDEVs.

I₀₁₄₅ Permitted PDEV state transitions are shown in the following table. The rightmost column lists the events which can cause the corresponding state transition.

A transition from the pseudo-state *NULL* represents creation of a PDEV object. A transition to the pseudo-state *NULL* represents destruction of a PDEV object.

From state	To state	Events
<i>NULL</i>	PDEV_NEW	RMI_PDEV_CREATE
PDEV_NEW	<i>NULL</i>	RMI_PDEV_DESTROY
PDEV_NEW	PDEV_NEW	RMI_PDEV_ABORT RMI_PDEV_COMMUNICATE
PDEV_NEW	PDEV_NEEDS_KEY	RMI_PDEV_COMMUNICATE
PDEV_NEEDS_KEY	PDEV_HAS_KEY	RMI_PDEV_SET_PUBKEY
PDEV_HAS_KEY	PDEV_HAS_KEY	RMI_PDEV_ABORT RMI_PDEV_COMMUNICATE

From state	To state	Events
PDEV_HAS_KEY	PDEV_READY	RMI_PDEV_COMMUNICATE
PDEV_READY	PDEV_COMMUNICATING	RMI_PDEV_NOTIFY
PDEV_COMMUNICATING	PDEV_COMMUNICATING	RMI_PDEV_COMMUNICATE
PDEV_COMMUNICATING	PDEV_READY	RMI_PDEV_ABORT RMI_PDEV_COMMUNICATE
PDEV_READY	PDEV_IDE_RESETTING	RMI_PDEV_IDE_RESET
PDEV_IDE_RESETTING	PDEV_READY	RMI_PDEV_COMMUNICATE
PDEV_NEW	PDEV_STOPPING	RMI_PDEV_STOP
PDEV_NEEDS_KEY	PDEV_STOPPING	RMI_PDEV_STOP
PDEV_HAS_KEY	PDEV_STOPPING	RMI_PDEV_STOP
PDEV_READY	PDEV_STOPPING	RMI_PDEV_STOP
PDEV_IDE_RESETTING	PDEV_STOPPING	RMI_PDEV_STOP
PDEV_ERROR	PDEV_STOPPING	RMI_PDEV_STOP
PDEV_STOPPING	PDEV_STOPPED	RMI_PDEV_COMMUNICATE
PDEV_STOPPED	NULL	RMI_PDEV_DESTROY

See also:

- [B4.3.13 RMI_PDEV_ABORT command](#)
- [B4.3.15 RMI_PDEV_COMMUNICATE command](#)
- [B4.3.16 RMI_PDEV_CREATE command](#)
- [B4.3.17 RMI_PDEV_DESTROY command](#)
- [B4.3.19 RMI_PDEV_IDE_RESET command](#)
- [B4.3.20 RMI_PDEV_NOTIFY command](#)
- [B4.3.21 RMI_PDEV_SET_PUBKEY command](#)
- [B4.3.22 RMI_PDEV_STOP command](#)

See also:

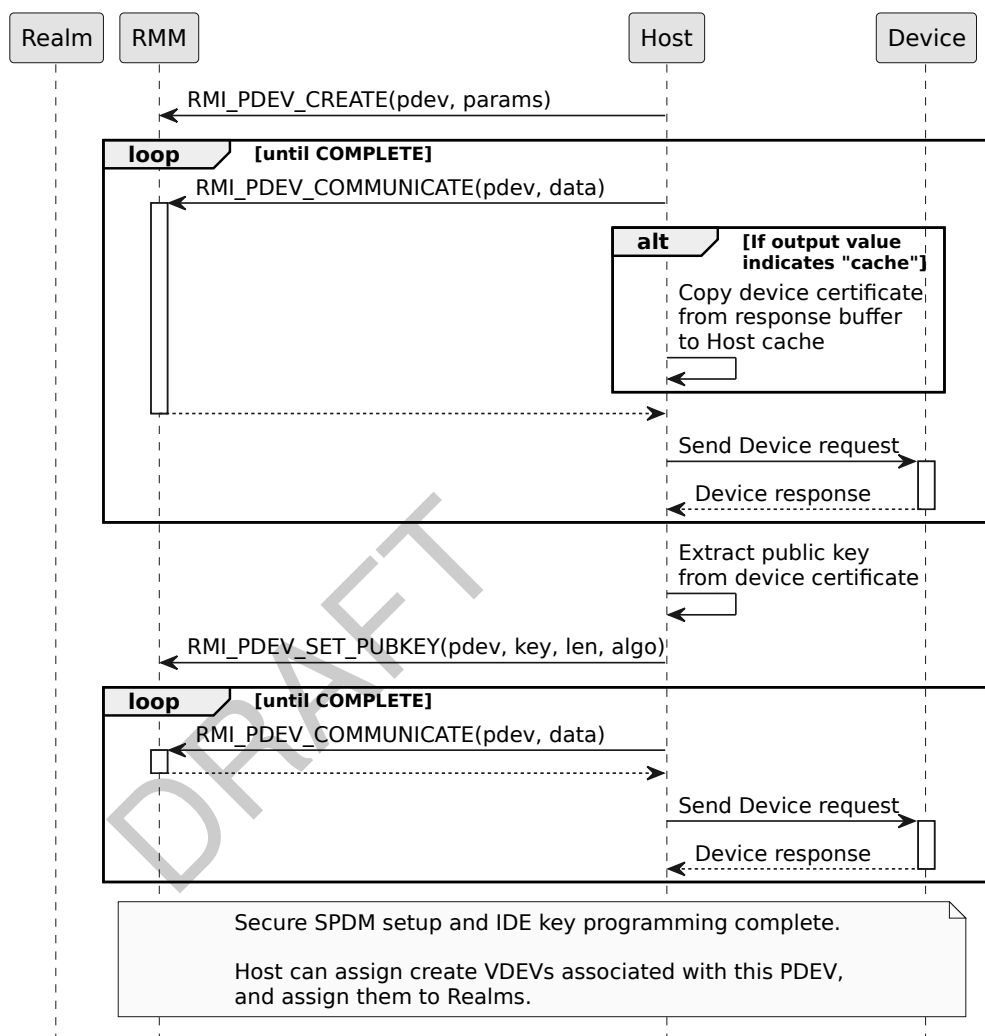
- [PCI Express 6.0 specification \[14\]](#)
- [A9.2 Communication between RMM and a device](#)

A9.3.3 Physical device flows

A9.3.3.1 Physical device setup flow

I0146

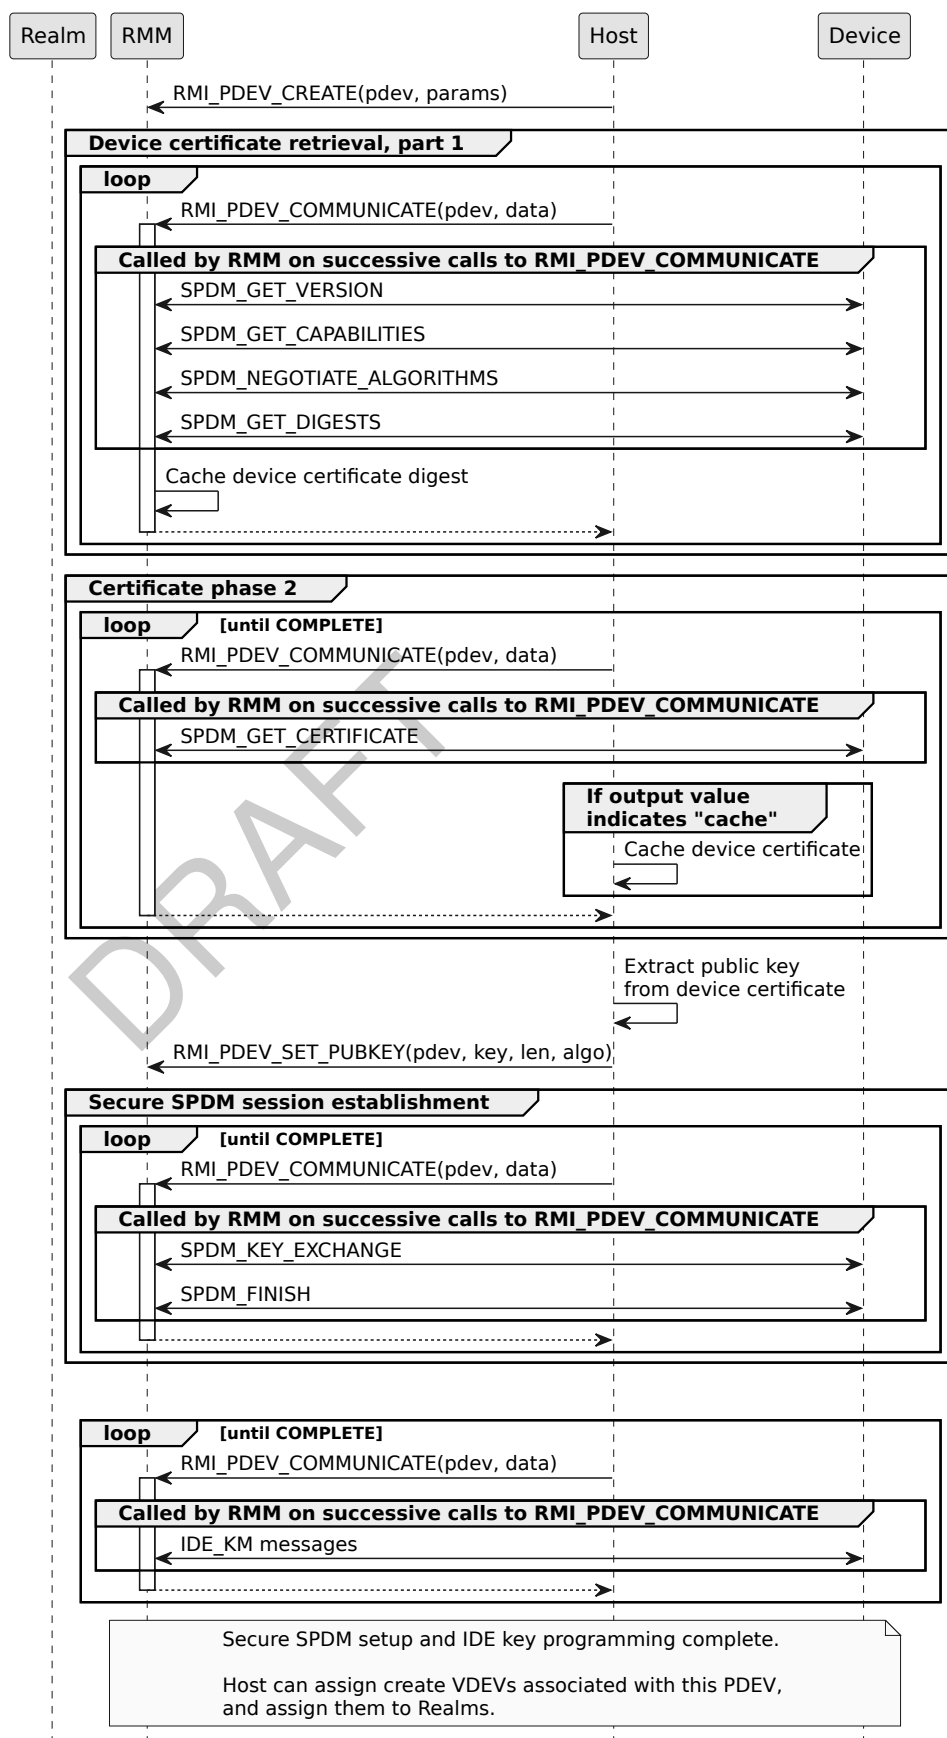
Setup of a PDEV is illustrated in the following sequence diagram.



I₀₁₄₇

Mapping of the PDEV setup flow onto SPDm and IDE communication with a TDISP PCIe device is illustrated in the following sequence diagram.

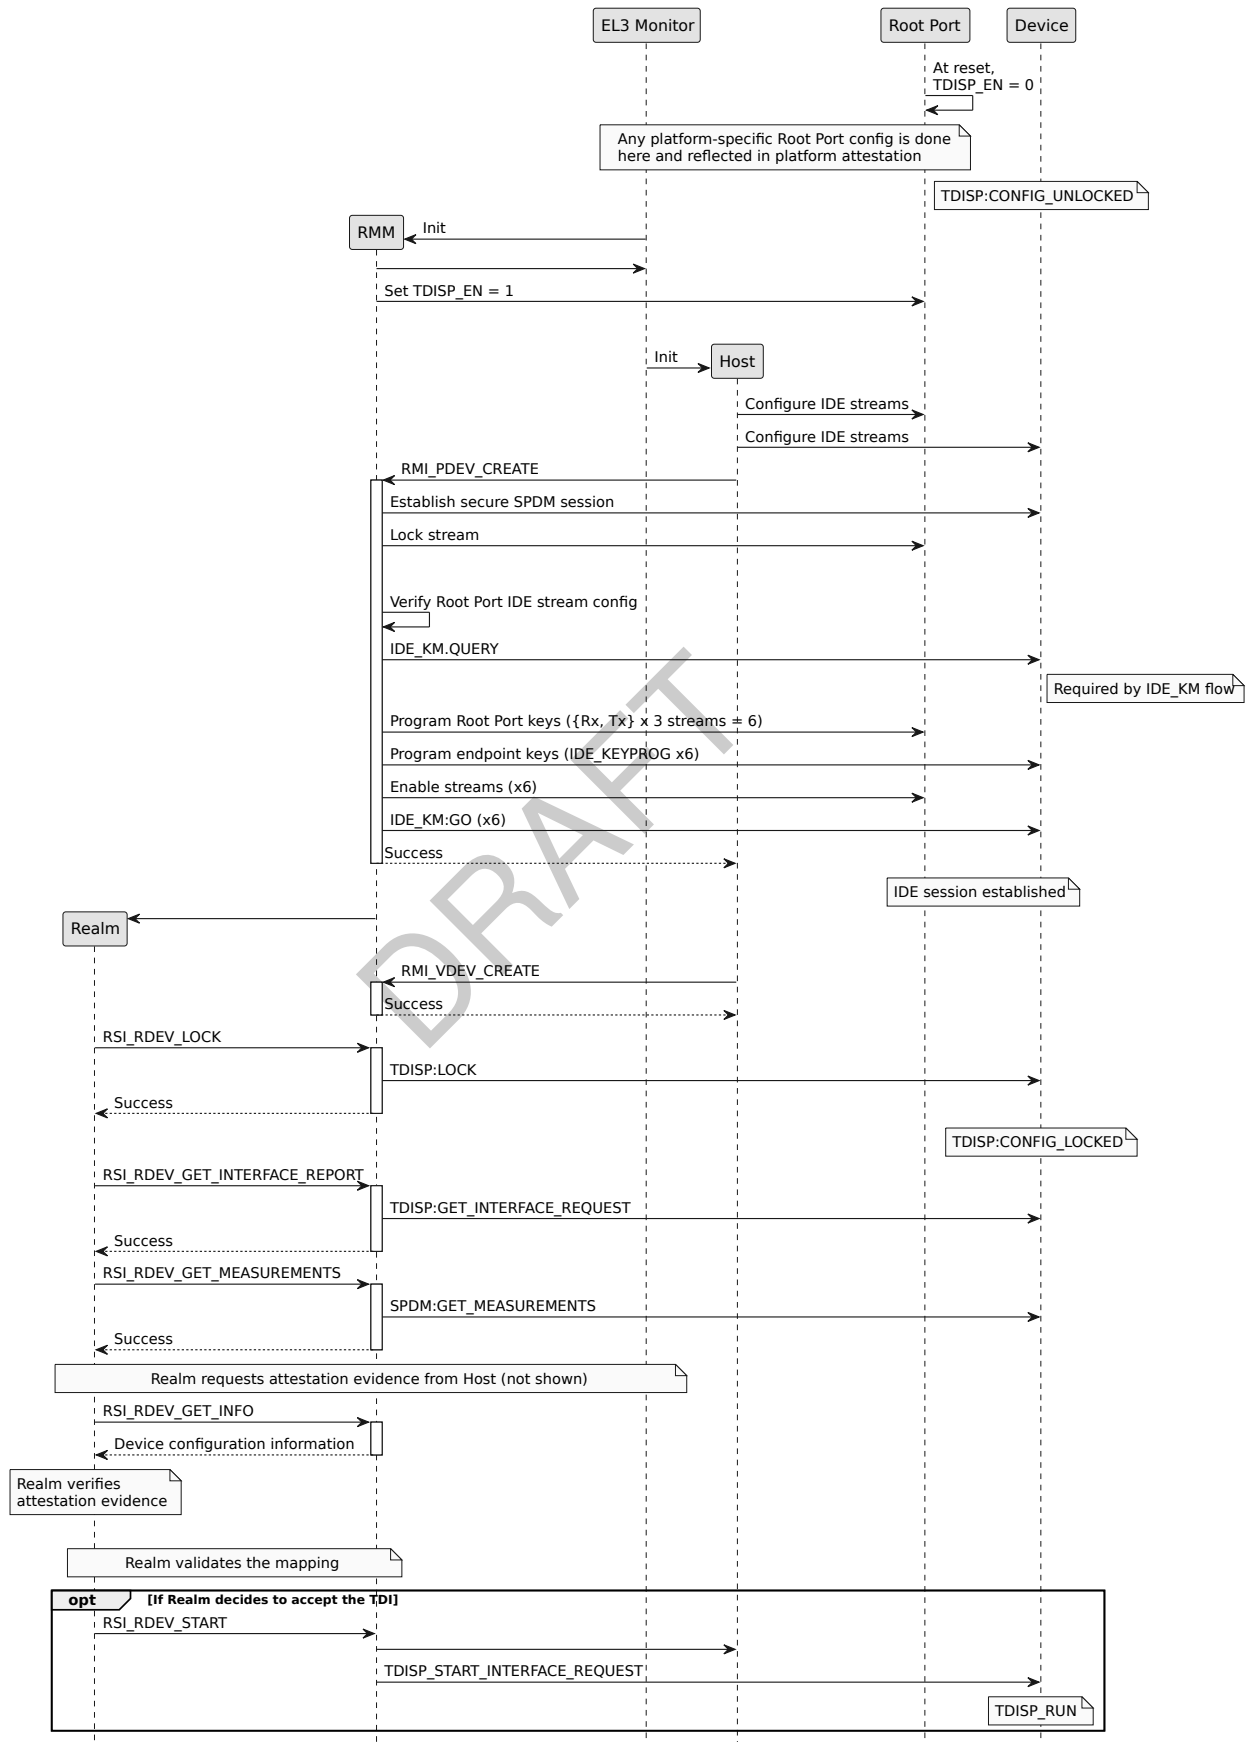
DRAFT



I₀₁₄₈

IDE setup for a TDISP PCIe device is illustrated in the following sequence diagram.

DRAFT



See also:

- [PCI Express 6.0 specification](#) [14]

DRAFT

A9.4 Virtual device object

D₀₁₄₉ A *virtual device* (VDEV) represents the binding between a device function and a Realm.

A9.4.1 Virtual device attributes

D₀₁₅₀ The attributes of a VDEV are summarized in the following table.

Name	Type	Description
vdev_id	Bits64	Virtual device identifier
tdi_id	Bits64	TDI identifier
inst_id	UInt64	Instance identifier
pdev	Address	PA of parent PDEV
realm	Address	PA of RD of Realm which owns this REC
state	RmmVdevState	Lifecycle state
comm_state	RmmDevCommState	Device communication state
aux	Address[32]	Addresses of auxiliary Granules
num_aux	UInt64	Number of auxiliary Granules

D₀₁₅₁ A *Device identifier* is a value which uniquely identifies a VDEV on a system, and is agreed between the Host and the Realm to which the VDEV is assigned.

I₀₁₅₂ The choice of device identifier depends upon the interface provided by the Host to the Realm. For example, if the Host provides a PCIe interface, then identifier is a PCIe (bus, device, function) tuple.

I₀₁₅₃ The Host provides the device identifier when executing RMI_VDEV_CREATE.

I₀₁₅₄ The Realm provides the device identifier when executing any RSI_RDEV command.

See also:

- [B4.3.51 RMI_VDEV_CREATE command](#)

A9.4.2 Virtual device invariants

R₀₁₅₅ Providing a tdi_id which is already used by another VDEV within the same segment causes RMI_VDEV_CREATE to fail.

U₀₁₅₆ The RMM can track usage of tdi_id values within a segment by using SW_RESERVED bits in the SMMU Stream Table Entry.

R₀₁₅₇ Providing a tdi_id which is outside the RID range of the parent PDEV causes RMI_VDEV_CREATE to fail.

A9.4.3 Virtual device lifecycle

A9.4.3.1 States

D₀₁₅₈ The states of a VDEV are listed below.

State	Description
VDEV_READY	No device transaction is associated with the VDEV.
VDEV_COMMUNICATING	The RMM is communicating with the VDEV.
VDEV_STOPPING	The RMM is communicating with the VDEV to stop the device interface.
VDEV_STOPPED	Device interface is stopped.
VDEV_ERROR	Device interface has reported a fatal error.

A9.4.3.2 State transitions

I₀₁₅₉

Permitted VDEV state transitions are shown in the following figure. Each arc is labeled with the events which can cause the corresponding state transition.

A transition from the pseudo-state *NULL* represents creation of a VDEV. A transition to the pseudo-state *NULL* represents destruction of a VDEV.

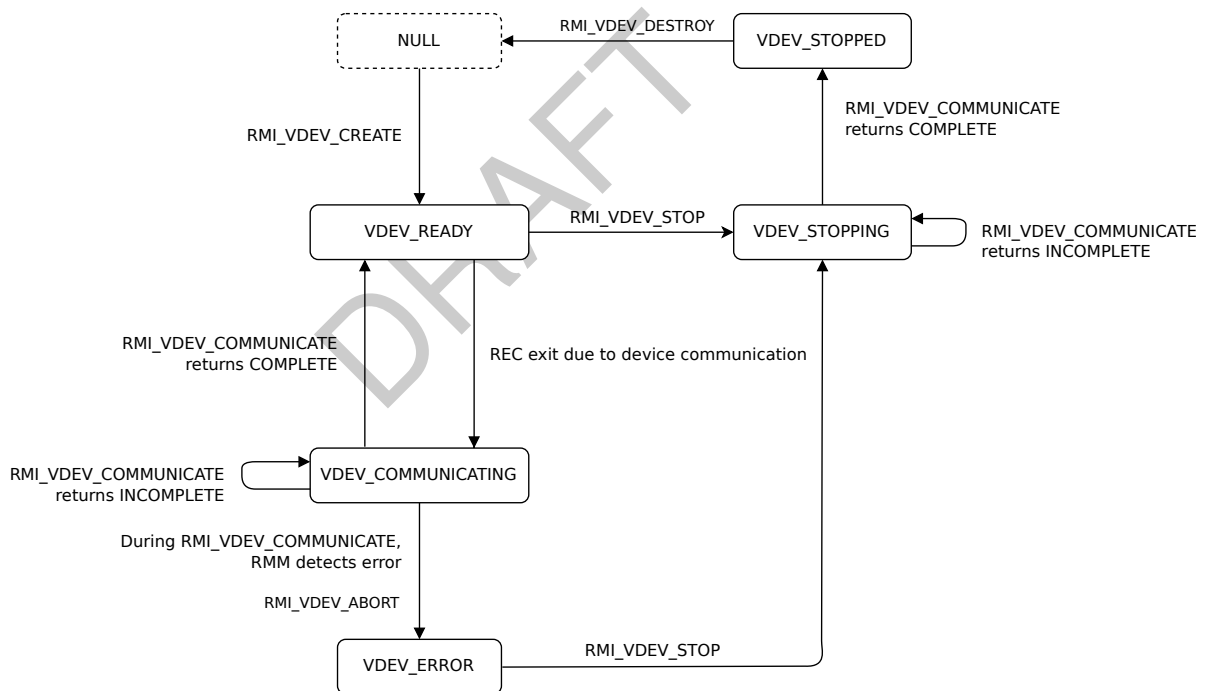


Figure A9.3: VDEV state transitions

I₀₁₆₀

While a VDEV is in *VDEV_READY* state, on a REC exit due to device communication, the VDEV moves to *VDEV_COMMUNICATING* state.

I₀₁₆₁

While a VDEV is in *VDEV_COMMUNICATING* state, the Host can either:

- Transfer device requests and device responses by executing *RMI_VDEV_COMMUNICATE*. If the return value indicates that the device transaction is complete then the VDEV moves to *VDEV_READY* state.
- Abort the device transaction by executing *RMI_VDEV_ABORT*. On successful execution of this command, the VDEV moves to *VDEV_ERROR* state.

- I₀₁₆₂ On execution of `RMI_VDEV_COMMUNICATE`, if the RMM detects a fatal error (such as an unexpected response or a protocol error) then the VDEV moves to `VDEV_ERROR` state.
- I₀₁₆₃ While a VDEV is in `VDEV_READY` state or `VDEV_ERROR` state, the Host can request the RMM to stop the device by executing `RMI_VDEV_STOP`.
If the return value is `RMI_DEV_COMM_PENDING` then the VDEV moves to `VDEV_STOPPING` state.
- I₀₁₆₄ While a VDEV is in `VDEV_STOPPING` state, the Host can enact the “stop” device transaction by executing `RMI_VDEV_COMMUNICATE`.
If the return value indicates that the device transaction is complete then the VDEV moves to `VDEV_STOPPED` state.
- I₀₁₆₅ While a VDEV is in `VDEV_STOPPING` state, if the device fails to respond or reports an error then the Host can call `RMI_VDEV_COMMUNICATE`, passing `RMI_DEV_COMM_ERROR`.
On successful execution of this command, the VDEV moves to `VDEV_STOPPED` state.
- I₀₁₆₆ While a VDEV is in `VDEV_STOPPED` state, the Host can reclaim resources by executing `RMI_VDEV_DESTROY`.
- I₀₁₆₇ Permitted VDEV state transitions are shown in the following table. The rightmost column lists the events which can cause the corresponding state transition.

A transition from the pseudo-state `NULL` represents creation of a VDEV object. A transition to the pseudo-state `NULL` represents destruction of a VDEV object.

From state	To state	Events
<code>NULL</code>	<code>VDEV_READY</code>	RMI_VDEV_CREATE
<code>VDEV_READY</code>	<code>VDEV_COMMUNICATING</code>	REC exit due to device communication
<code>VDEV_READY</code>	<code>VDEV_STOPPING</code>	RMI_VDEV_STOP
<code>VDEV_COMMUNICATING</code>	<code>VDEV_COMMUNICATING</code>	RMI_VDEV_COMMUNICATE
<code>VDEV_COMMUNICATING</code>	<code>VDEV_ERROR</code>	RMI_VDEV_ABORT
<code>VDEV_ERROR</code>	<code>VDEV_STOPPING</code>	RMI_VDEV_STOP
<code>VDEV_STOPPING</code>	<code>VDEV_STOPPING</code>	RMI_VDEV_COMMUNICATE
<code>VDEV_STOPPING</code>	<code>VDEV_STOPPED</code>	RMI_VDEV_COMMUNICATE
<code>VDEV_STOPPED</code>	<code>NULL</code>	RMI_VDEV_DESTROY

See also:

- [B4.3.47 RMI_VDEV_ABORT command](#)
- [B4.3.49 RMI_VDEV_COMMUNICATE command](#)
- [B4.3.51 RMI_VDEV_CREATE command](#)
- [B4.3.52 RMI_VDEV_DESTROY command](#)
- [B4.3.54 RMI_VDEV_STOP command](#)

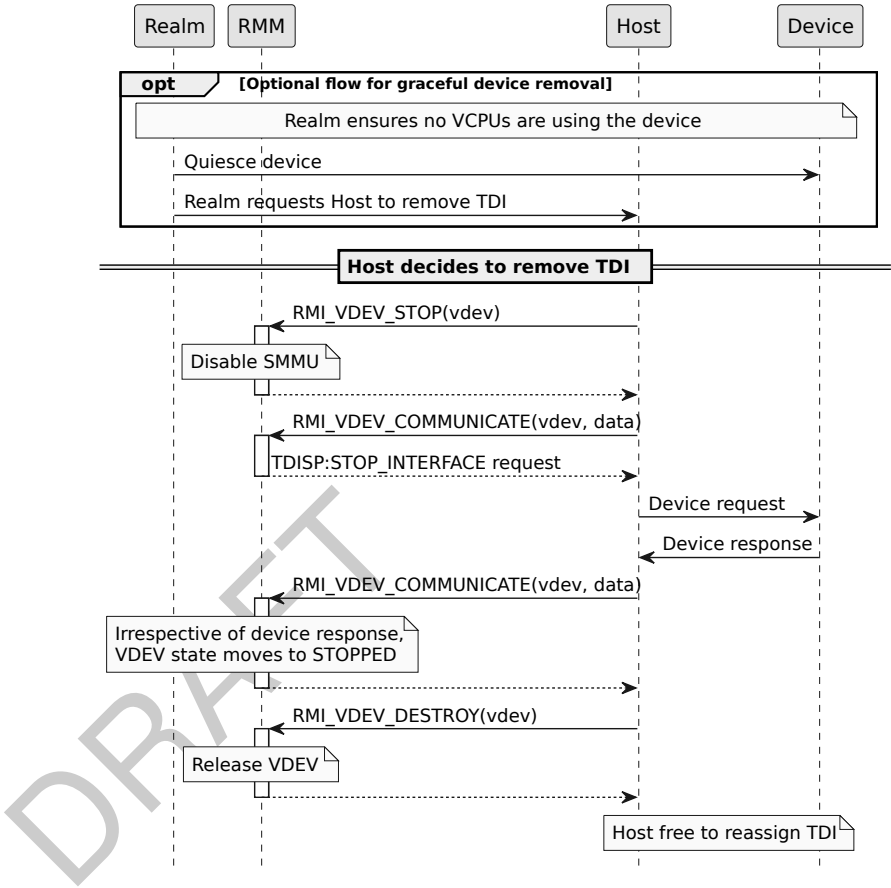
See also:

- [A4.3.11 REC exit due to device communication](#)

A9.4.4 Virtual device flows

A9.4.4.1 Virtual device teardown flow

I0168 Teardown of a VDEV is illustrated in the following sequence diagram.



A9.5 Realm management of an assigned device interface

This section describes interaction between a Realm and the RMM to manage an RDEV.

A9.5.1 Interruptible Realm device operations

D₀₁₆₉ The following are *interruptible Realm device operations*:

- Requesting digests of device attestation evidence
- Requesting a device interface report
- Requesting device measurements
- Locking a device interface
- Starting a device interface
- Stopping a device interface

I₀₁₇₀ Interruptible Realm device operation are:

- Initiated by the Realm executing RSI commands
- Enacted via a REC exit due to device communication, and by subsequent Host calls to RMI_VDEV_COMMUNICATE.

I₀₁₇₁ The Realm programming model for interruptible Realm device operation consists of two phases:

1. The Realm initiates the operation by executing an RSI command. Successful execution of this command causes the RDEV to transition to a *busy* state.
2. The Realm executes RSI_RDEV_CONTINUE in a loop, until the result is not RSI_INCOMPLETE.

The following pseudocode illustrates this programming model, using RSI_RDEV_GET_INTERFACE_REPORT as an example.

```
int get_interface_report(...)
{
    int ret;

    ret = RSI_RDEV_GET_INTERFACE_REPORT(dev_id, ...);
    if (ret) {
        return ret;
    }

    do {
        ret = RSI_RDEV_CONTINUE(dev_id);
    } while (ret == RSI_INCOMPLETE);

    return ret;
}
```

I₀₁₇₂ Once an interruptible Realm device operation has been initiated by the Realm, it must either be:

- Completed by the Realm, via execution of RSI_RDEV_CONTINUE and subsequent calls by the Host to RMI_VDEV_COMMUNICATE, or
- Aborted by the Host, via execution of RMI_VDEV_ABORT.

If the Realm wishes to abort an interruptible Realm device operation, it must request the Host to do so.

I₀₁₇₃ The set of commands which can initiate an interruptible Realm device operation are as follows:

- RSI_RDEV_GET_INTERFACE_REPORT
- RSI_RDEV_GET_MEASUREMENTS
- RSI_RDEV_LOCK
- RSI_RDEV_START
- RSI_RDEV_STOP

See also:

- [A4.3.11 REC exit due to device communication](#)
- [A9.2 Communication between RMM and a device](#)
- [B4.3.49 RMI_VDEV_COMMUNICATE command](#)
- [B5.3.17 RSI_RDEV_GET_INTERFACE_REPORT command](#)
- [B5.3.18 RSI_RDEV_GET_MEASUREMENTS command](#)
- [B5.3.20 RSI_RDEV_LOCK command](#)
- [B5.3.21 RSI_RDEV_START command](#)
- [B5.3.22 RSI_RDEV_STOP command](#)

A9.5.2 Realm retrieval of device attestation evidence

I₀₁₇₄ A Realm is expected to retrieve cached device attestation evidence from the Host via an RSI_HOST_CALL interface. Details of this interface are out of scope of this specification.

I₀₁₇₅ A Realm can retrieve digests of cached device attestation evidence from the RMM by executing the RSI_RDEV_GET_INFO command.

See also:

- [A9.2.4 Host caching of device attestation evidence](#)
- [B5.3.4 RSI_HOST_CALL command](#)
- [B5.3.16 RSI_RDEV_GET_INFO command](#)

A9.5.3 Realm validation of device memory mappings

I₀₁₇₆ A Realm device interface report describes the memory regions of the device. Each device memory region has the following attributes in the report:

- PA base address of the region.
This is obfuscated by addition of a random offset value to the system physical address. The offset value is known to the RMM.
- Size of the region.
- Whether the mapping is private or shared.
- Whether the output address of the mapping is within the system coherent memory space.

I₀₁₇₇ A Realm can validate by execution of RSI_RDEV_VALIDATE_MAPPING that each MMIO region in the Realm device interface report has been correctly mapped into the Realm's Protected IPA space.

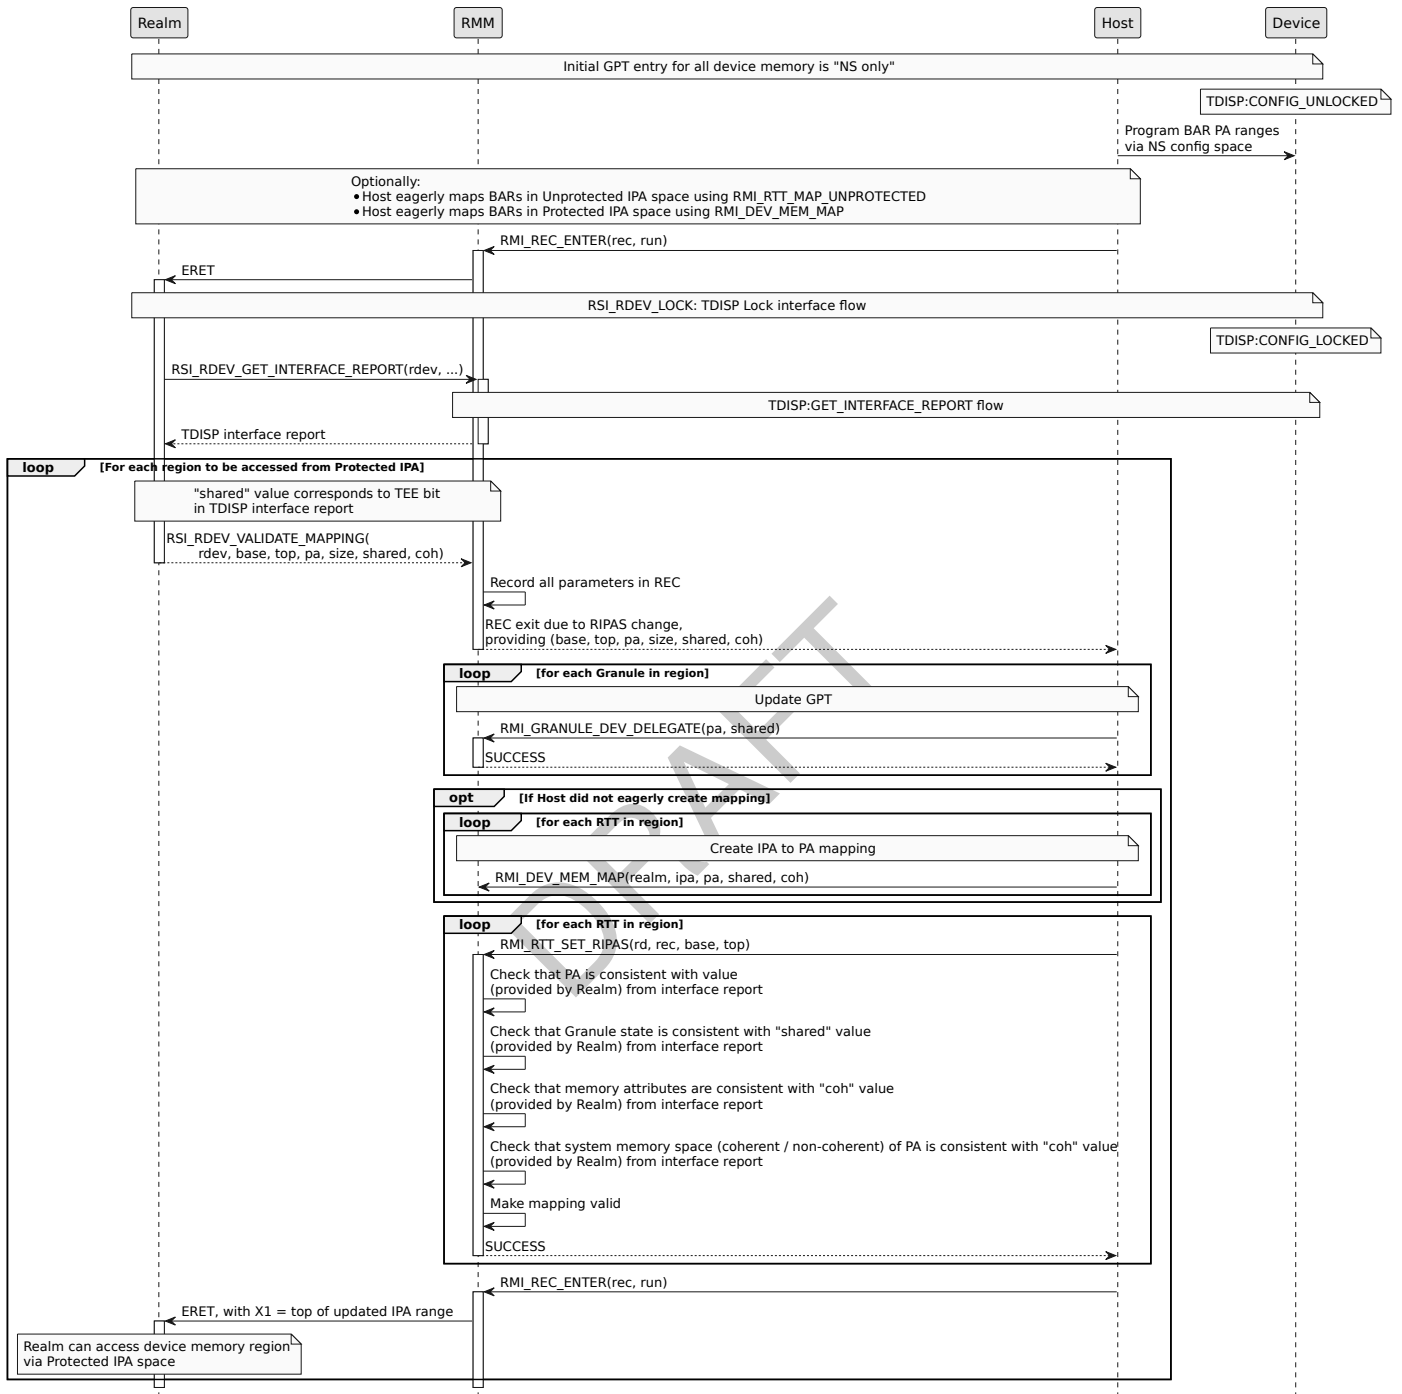
I₀₁₇₈ Execution of RSI_RDEV_VALIDATE_MAPPING initiates a RIPAS change request. On execution of RSI_RDEV_VALIDATE_MAPPING, the RMM records the PA base address of the region and the flags provided by the Realm in the REC. On subsequent execution of RMI_RTT_SET_RIPAS, the RMM validates that the contents of the target RTTE are compatible with the PA base address of the region and the flags provided by the Realm. If this validation passes then the RMM sets the RIPAS of the RTTE to DEV.

R₀₁₇₉ The RTTE attributes checked by RMI_RTT_SET_RIPAS, when the target RIPAS is RIPAS_DEV, are as follows:

- The IPA to PA mapping is consistent with the PA base address of the region provided by the Realm.
- The value of the GPT entry for the PA is consistent with the “shared” flag provided by the Realm.
- The memory attributes are consistent with the value of the “coh” flag provided by the Realm.
- The system memory space (coherent or non-coherent) of the PA is consistent with the “coh” flag provided by the Realm.

I_0180

Creation and validation of device memory mappings is illustrated in the following sequence diagram.



See also:

- [A5.4 RIPAS change](#)
- [A5.5.11 RTT entry attributes](#)
- [B5.3.23 RSI_RDEV_VALIDATE_MAPPING command](#)

A9.5.4 Realm device attributes

D_0181

The attributes of an RDEV are summarized in the following table.

Name	Type	Description
state	RmmRdevState	Lifecycle state
operation	RmmRdevOperation	Operation being performed by the RDEV
vdev_ptr	Address	PA of VDEV associated with this RDEV

A9.5.5 Realm device lifecycle

A9.5.5.1 States

D₀₁₈₂

The states of a RDEV are listed below.

State	Description
RDEV_NEW	Device interface is unlocked.
RDEV_NEW_BUSY	Device interface is unlocked and is handling an interruptible Realm device operation.
RDEV_LOCKED	Device interface is locked.
RDEV_LOCKED_BUSY	Device interface is locked and is handling an interruptible Realm device operation.
RDEV_STARTED	Device interface is started.
RDEV_STARTED_BUSY	Device interface is started and is handling an interruptible Realm device operation.
RDEV_STOPPING	Device interface is stopping.
RDEV_STOPPED	Device interface is stopped.
RDEV_ERROR	Device interface has reported a fatal error.

A9.5.5.2 State transitions

I₀₁₈₃

Permitted RDEV Realm state transitions are shown in the following figure. Each arc is labeled with the events which can cause the corresponding state transition.

A transition from the pseudo-state *NULL* represents creation of an RDEV.

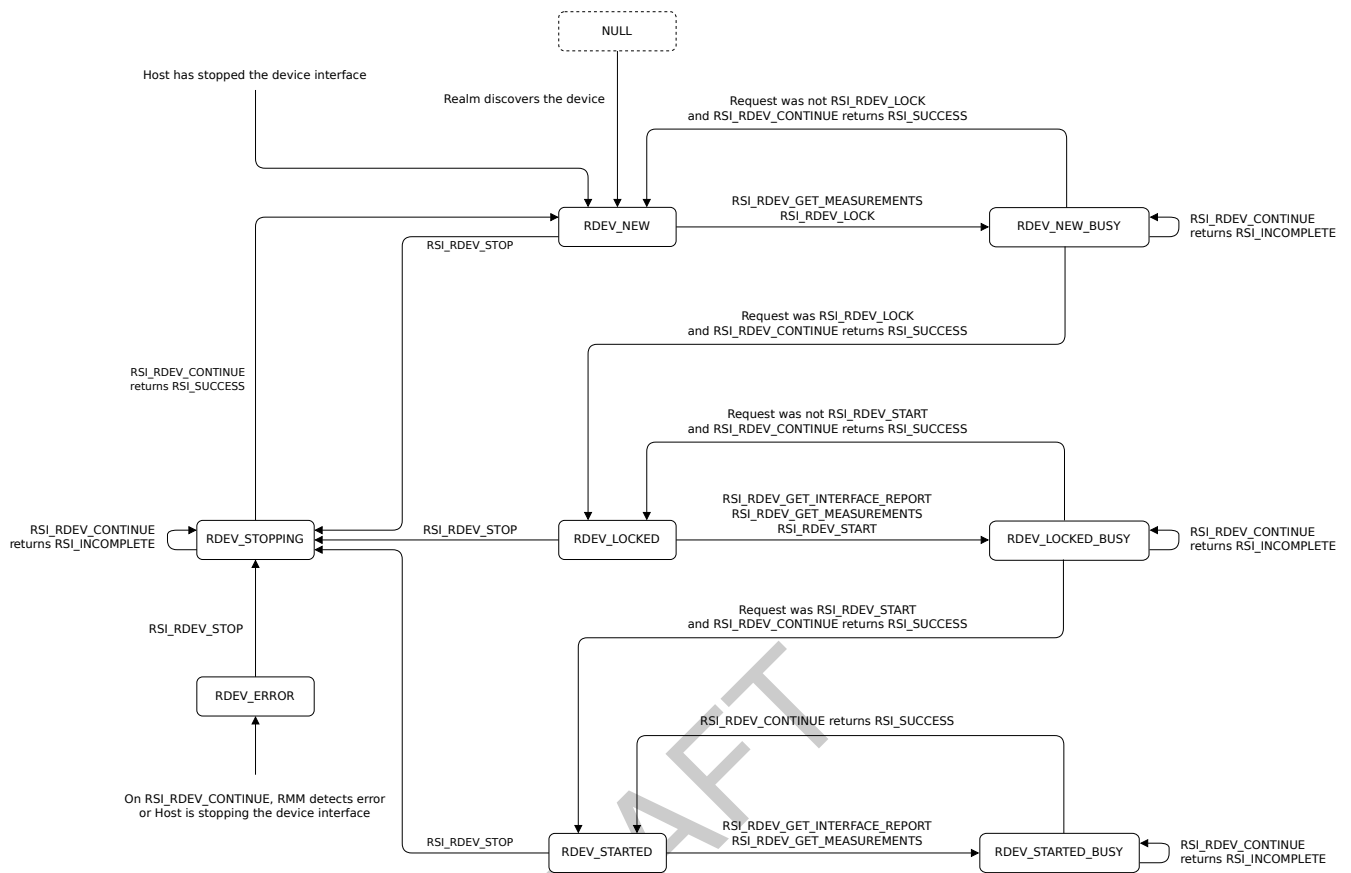


Figure A9.4: RDEV state transitions

I0184 The Realm discovers an assigned PCIe TDI by probing PCIe config space. Arm expects that this will be emulated by the Host, via Unprotected IPA space.

I0185 While an RDEV is in RDEV_NEW state, the Realm can execute any of the following commands:

- RSI_RDEV_GET_MEASUREMENTS
- RSI_RDEV_LOCK

On successful execution of this command, the RDEV moves to RDEV_NEW_BUSY state and records the requested operation.

X0186 Permitting `RSI_RDEV_GET_MEASUREMENTS` while the RDEV is in `RDEV_NEW` state allows the Realm to perform the following sequence:

1. Request a measurement of the initial firmware present in a device.
2. Transition the device to RDEV_LOCKED.
3. Upload new or additional firmware to the device.
4. Verify that the firmware has been successfully uploaded, by requesting a new measurement from the device.

While an RDEV is in RDEV_NEW_BUSY state, the Realm can execute RSI_RDEV_CONTINUE.

Once the requested operation is complete:

- If the requested operation was `RSI_RDEV_LOCK`, the RDEV moves to `RDEV_LOCKED` state.
- Otherwise the RDEV moves to `RDEV_NEW` state.

While an RDEV is in RDEV_LOCKED state, the Realm can execute any of the following commands:

- RSI_RDEV_GET_INTERFACE_REPORT
- RSI_RDEV_GET_MEASUREMENTS
- RSI_RDEV_START

On successful execution of this command, the RDEV moves to RDEV_LOCKED_BUSY state and records the requested operation.

I₀₁₈₉

While an RDEV is in RDEV_LOCKED_BUSY state, the Realm can execute RSI_RDEV_CONTINUE.

Once the requested operation is complete:

- If the requested operation was RSI_RDEV_START, the RDEV moves to RDEV_STARTED state.
- Otherwise the RDEV moves to RDEV_LOCKED state.

I₀₁₉₀

While an RDEV is in RDEV_STARTED state, the Realm can execute any of the following commands:

- RSI_RDEV_GET_INTERFACE_REPORT
- RSI_RDEV_GET_MEASUREMENTS

On successful execution of this command, the RDEV moves to RDEV_STARTED_BUSY state and records the requested operation.

I₀₁₉₁

While an RDEV is in RDEV_STARTED_BUSY state, the Realm can execute RSI_RDEV_CONTINUE.

Once the requested operation is complete, the RDEV moves to RDEV_STARTED state.

I₀₁₉₂

On execution of RSI_RDEV_CONTINUE, if the RMM detects a fatal error (such as an unexpected response or a protocol error) then the RDEV moves to RDEV_ERROR state.

I₀₁₉₃

While an RDEV is in any state, the Realm can execute RSI_RDEV_STOP.

On successful execution of this command the RDEV moves to RDEV_STOPPING state.

Following successful execution of this command, accesses from the device to Realm memory are blocked.

I₀₁₉₄

While an RDEV is in RDEV_STOPPING state, the Realm can execute RSI_RDEV_CONTINUE.

Once the requested operation is complete, the RDEV moves to RDEV_STOPPED state.

I₀₁₉₅

Permitted RDEV state transitions are shown in the following table. The rightmost column lists the events which can cause the corresponding state transition.

From state	To state	Events
RDEV_NEW	RDEV_NEW_BUSY	RSI_RDEV_GET_MEASUREMENTS RSI_RDEV_LOCK
RDEV_NEW_BUSY	RDEV_NEW_BUSY	RSI_RDEV_CONTINUE
RDEV_NEW_BUSY	RDEV_LOCKED	RSI_RDEV_CONTINUE
RDEV_LOCKED	RDEV_LOCKED_BUSY	RSI_RDEV_GET_INTERFACE_REPORT RSI_RDEV_GET_MEASUREMENTS RSI_RDEV_START
RDEV_LOCKED_BUSY	RDEV_LOCKED_BUSY	RSI_RDEV_CONTINUE
RDEV_LOCKED_BUSY	RDEV_STARTED	RSI_RDEV_CONTINUE
RDEV_STARTED	RDEV_STARTED_BUSY	RSI_RDEV_GET_INTERFACE_REPORT RSI_RDEV_GET_MEASUREMENTS
RDEV_STARTED_BUSY	RDEV_STARTED_BUSY	RSI_RDEV_CONTINUE
RDEV_NEW	RDEV_STOPPING	RSI_RDEV_STOP
RDEV_LOCKED	RDEV_STOPPING	RSI_RDEV_STOP

From state	To state	Events
RDEV_STARTED	RDEV_STOPPING	RSI_RDEV_STOP
RDEV_NEW	RDEV_ERROR	RSI_RDEV_CONTINUE Host is stopping device interface
RDEV_NEW_BUSY	RDEV_ERROR	RSI_RDEV_CONTINUE Host is stopping device interface
RDEV_NEW_BUSY	RDEV_NEW	RSI_RDEV_CONTINUE Host has stopped device interface
RDEV_LOCKED	RDEV_ERROR	RSI_RDEV_CONTINUE Host is stopping device interface
RDEV_LOCKED_BUSY	RDEV_ERROR	RSI_RDEV_CONTINUE Host is stopping device interface
RDEV_LOCKED_BUSY	RDEV_NEW	RSI_RDEV_CONTINUE Host has stopped device interface
RDEV_STARTED	RDEV_ERROR	RSI_RDEV_CONTINUE Host is stopping device interface
RDEV_STARTED_BUSY	RDEV_ERROR	RSI_RDEV_CONTINUE Host is stopping device interface
RDEV_STARTED_BUSY	RDEV_NEW	RSI_RDEV_CONTINUE Host has stopped device interface
RDEV_ERROR	RDEV_STOPPING	RSI_RDEV_STOP
RDEV_STOPPING	RDEV_STOPPING	RSI_RDEV_CONTINUE
RDEV_STOPPING	RDEV_NEW	RSI_RDEV_CONTINUE

See also:

- [B5.3.15 RSI_RDEV_CONTINUE command](#)
- [B5.3.17 RSI_RDEV_GET_INTERFACE_REPORT command](#)
- [B5.3.18 RSI_RDEV_GET_MEASUREMENTS command](#)
- [B5.3.20 RSI_RDEV_LOCK command](#)
- [B5.3.21 RSI_RDEV_START command](#)
- [B5.3.22 RSI_RDEV_STOP command](#)

A9.5.5.3 Relationship between RDEV state and TDISP TDI state

I₀₁₉₆ The lifecycle of an RDEV closely resembles the lifecycle of a TDISP TDI. There cannot exist a direct mapping between the two, because changes in TDI state (for example due to Host action) may not be immediately observed by either the RMM or the Realm. However, the two states are sufficiently closely coupled to provide the Realm with important guarantees regarding the TDI state.

R₀₁₉₇ On transition of an RDEV to RDEV_LOCKED state, the corresponding TDI transitions to CONFIG_LOCKED state.

R₀₁₉₈ On transition of an RDEV to RDEV_STARTED state, the corresponding TDI transitions to RUN state.

R₀₁₉₉ On detection by the RMM that a TDI is in ERROR state, the corresponding RDEV transitions to RDEV_ERROR state.

See also:

- [PCI Express 6.0 specification \[14\]](#)

A9.5.5.4 Relationship between RDEV state and SMMU enablement

R0200

When the state of an RDEV is RDEV_STARTED, the SMMU permits traffic from the device to the Realm's Protected IPA space. When the state of an RDEV is not RDEV_STARTED, the SMMU blocks traffic from the device to the Realm's Protected IPA space.

DRAFT

A9.5.6 Realm device flows

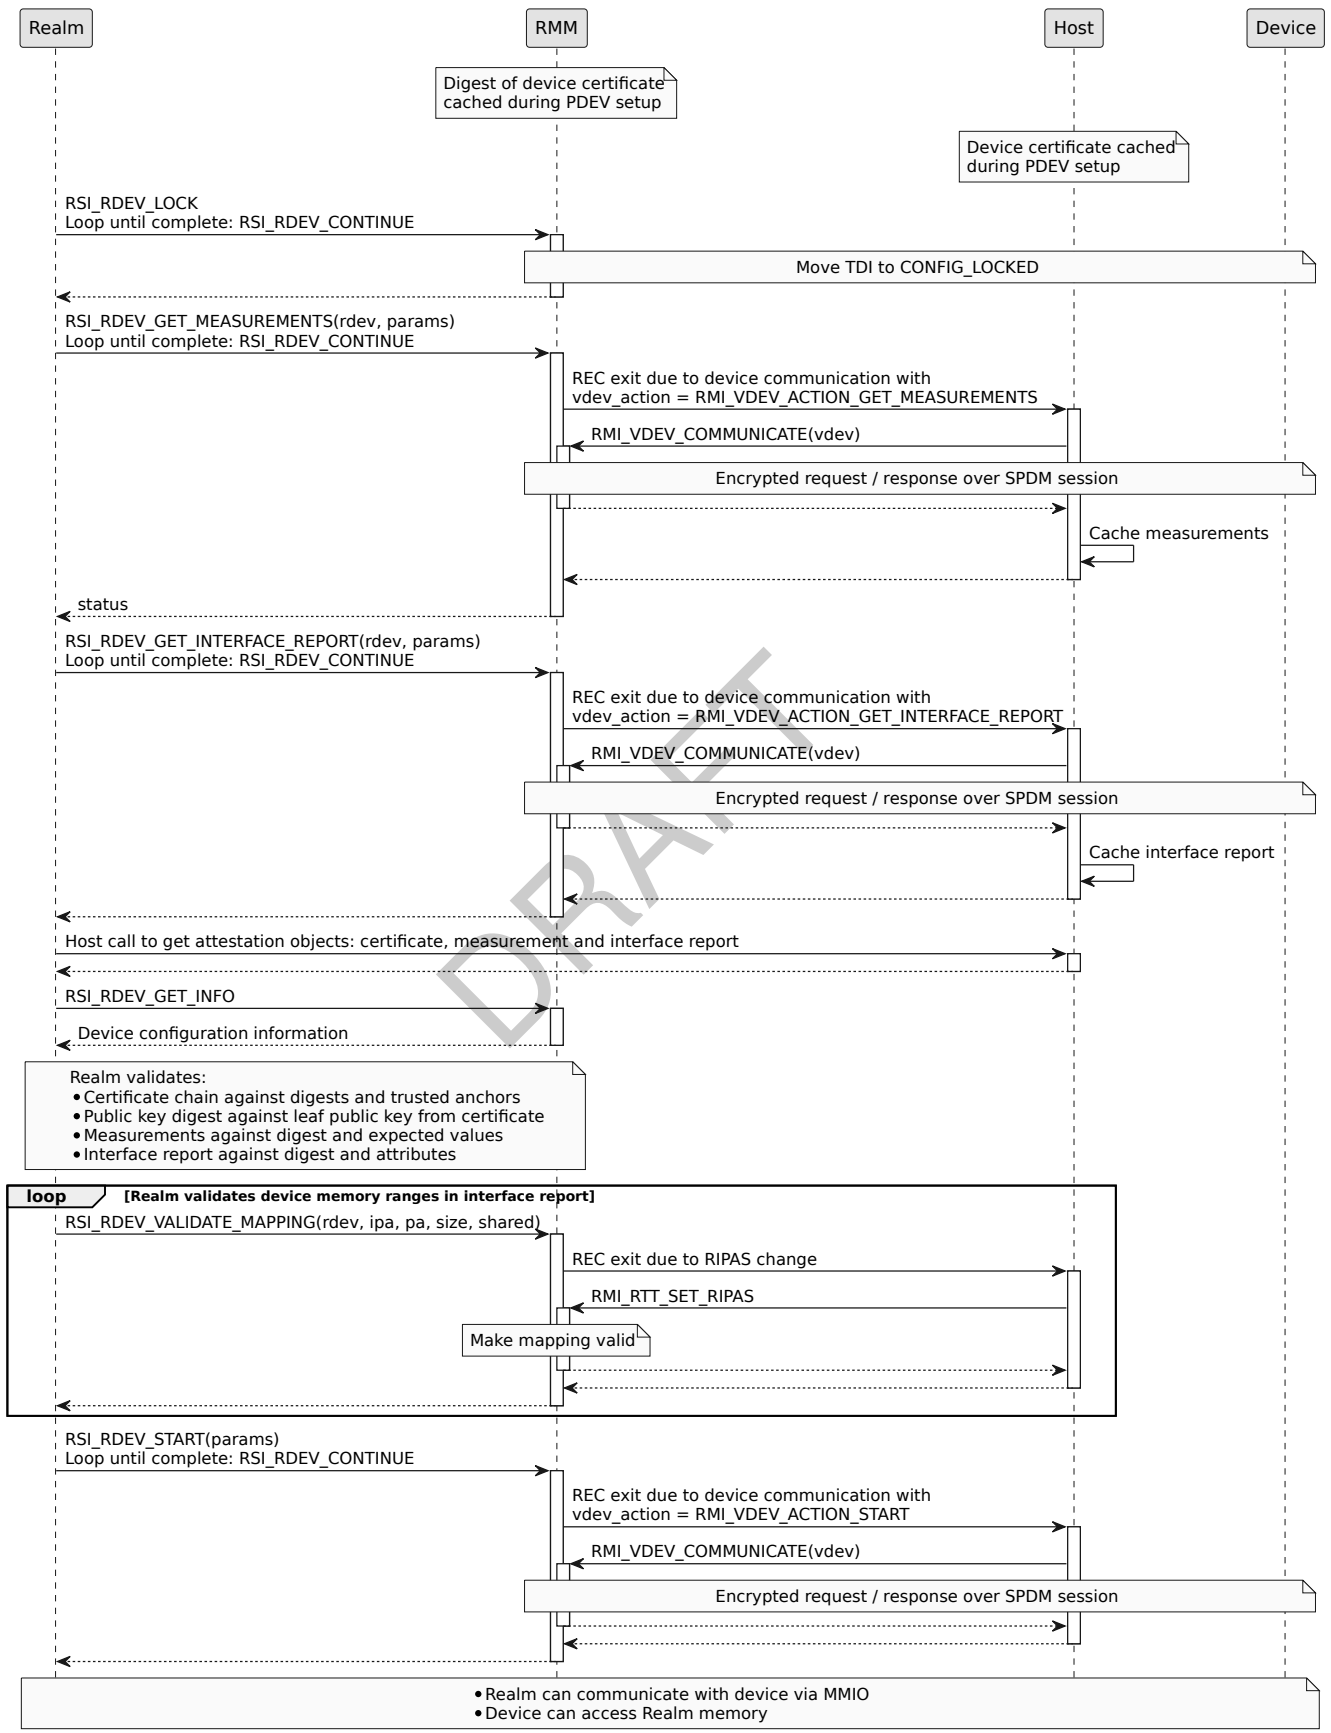
A9.5.6.1 Realm device setup flow

I_0201

Setup of an RDEV is illustrated in the following sequence diagram.

DRAFT

Chapter A9. Realm device assignment
A9.5. Realm management of an assigned device interface



A9.6 Device access to a Protected IPA

I₀₂₀₂

Device access to a Protected IPA follows the same rules as Realm data access to a Protected IPA.

See also:

- [A5.2.3 Realm access to a Protected IPA](#)

DRAFT

Chapter A10

Planes

This section describes how a Realm can be divided into multiple mutually isolated execution environments, called Planes.

Provisional

Decide whether this section should be:

- Moved to the [Concepts](#) chapter, alongside the [Realm](#) section.
- Left in this position, on the grounds that, like [Realm device assignment](#), Planes is an optional feature.

A10.1 Planes overview

Provisional

The following aspects and implications of Planes will be described in a future release of this specification:

- Access from P_n to a device assigned to the Realm

D₀₂₀₃

A Realm contains:

- a single *primary Plane*
- zero or more *auxiliary Planes*.

A Realm with a non-zero number of auxiliary Planes is said to contain *multiple Planes*.

I₀₂₀₄

The number of auxiliary Planes is specified by the Host at Realm creation.

D₀₂₀₅

Planes within a Realm are identified using a zero-based *Plane index*.

The Plane index of the primary Plane is zero.

D₀₂₀₆

When referring to the primary Plane of a Realm, this specification uses the term *Plane 0*, or *P0*.

When referring to any auxiliary Plane, this specification uses the term *P_n*.

I₀₂₀₇

All Planes within a Realm share a single IPA space.

Stage 2 memory access permissions for a given IPA can differ between Planes.

I₀₂₀₈

Each Plane has a VMID which is unique both within the owning Realm and among all Realms.

I₀₂₀₉

A Realm with multiple Planes may either have:

- An RTT tree per Plane, or
- A single RTT tree, with per-Plane access permissions being managed indirectly.

I₀₂₁₀

On REC exit due to Data Abort, Instruction Abort or RTT request, the index of the RTT tree used by the exited Plane is provided to the Host.

This allows the Host to know which RTT tree must be modified in order to service a fault.

D₀₂₁₁

A given VPE executes in one Plane at a time.

This is referred to as the *active Plane* of the VPE.

I₀₂₁₂

The following capabilities are available only to P0:

- Change the active Plane of the current VPE
- Read and write register state of other Planes in the current VPE
- Configure and take traps from other Planes in the current VPE
- Control delivery of virtual interrupts to other Planes in the current VPE

See also:

- [A3.11 Support for auxiliary Planes](#)
- [A10.2 Planes exception model](#)
- [A10.3 Planes memory management](#)
- [A10.4 Planes interrupts](#)

A10.2 Planes exception model

Provisional

Decide whether this section should be integrated into the main [Realm exception model](#) chapter.

A10.2.1 Plane exception model overview

- D₀₂₁₃ A *Plane entry* is a transition from P0 to Pn, due to execution of RSI_PLANE_ENTER.
- I₀₂₁₄ P0 provides the index of the target Plane (Pn) as an input to the RSI_PLANE_ENTER command.
- D₀₂₁₅ A *Plane exit* is return to P0 from an execution of RSI_PLANE_ENTER which caused a Plane entry.
- D₀₂₁₆ A *PlaneRun* object is a data structure used to pass values between the RMM and P0 on Plane entry and on Plane exit.
- I₀₂₁₇ A PlaneRun object is stored in Realm memory.
- I₀₂₁₈ Between a Plane entry and a Plane exit, a REC exit and REC entry may occur.
- As an example:
1. Running in REC A, P0 executes RSI_PLANE_ENTER, passing target Plane index 1.
This causes a Realm exit to the RMM, followed by a Realm entry to P1, within REC A.
 2. Running in REC A, P1 accesses an IPA which is not mapped (HIPAS = UNASSIGNED, RIPAS = RAM).
This causes a REC exit to the Host.
 3. The Host executes RMI_REC_ENTER, passing the address of REC A.
This causes the RMM to return to P1 within REC A.
 4. Running in REC A, P1 accesses an IPA whose RIPAS is EMPTY.
This causes a Plane exit to P0.
- I₀₂₁₉ Following a REC exit from P0, on the next entry to the same REC, control returns to P0.
- I₀₂₂₀ Following a REC exit from Pn, on the next entry to the same REC, the Host determines whether:
- Control returns to Pn. This is the default behaviour.
 - A Plane exit occurs, returning control to P0. This can be triggered for example due to the Host injecting a virtual interrupt into the REC.
- See also:
- [A4.1 Realm exception model overview](#)
 - [B5.3.12 RSI_PLANE_ENTER command](#)
 - [B5.4.18 RsiPlaneRun type](#)

A10.2.2 Plane entry

- D₀₂₂₁ An *RsiPlaneEnter object* is a data structure used to pass values from P0 to the RMM on Plane entry.
- I₀₂₂₂ An RsiPlaneEnter object is stored in the RsiPlaneRun object which is passed by P0 as an input to the RSI_PLANE_ENTER command.
- I₀₂₂₃ In this chapter, both `plane_enter` and “the RsiPlaneEnter object” refer to the RsiPlaneEnter object which is provided to the RSI_PLANE_ENTER command.
- I₀₂₂₄ On Plane entry, execution state is restored from the RsiPlaneEnter object to the PE.

I₀₂₂₅ An RsiPlaneEnter object contains attributes which are used to manage Pn virtual interrupts.

D₀₂₂₆ The attributes of an RsiPlaneEnter object are summarized in the following table.

Name	Byte offset	Type	Description
flags	0x0	RsiPlaneEnterFlags	Flags
pc	0x8	Bits64	Program counter
gprs[31]	0x100	Bits64	Registers
gicv3_hcr	0x200	Bits64	GICv3 Hypervisor Control Register value
gicv3_lrs[16]	0x208	Bits64	GICv3 List Register values

A10.2.3 Plane exit

D₀₂₂₇ An RsiPlaneExit object is a data structure used to pass values from the RMM to P0 on Plane exit.

I₀₂₂₈ An RsiPlaneExit object is stored in the RsiPlaneRun object which is passed by P0 as an input to the RSI_PLANE_ENTER command.

I₀₂₂₉ In this chapter, both plane_exit and “the RsiPlaneExit object” refer to the RsiPlaneExit object which is provided to the RSI_PLANE_ENTER command.

I₀₂₃₀ On Plane exit, execution state is saved from the PE to the RsiPlaneExit object.

D₀₂₃₁ The attributes of an RsiPlaneExit object are summarized in the following table.

Name	Byte offset	Type	Description
reason	0x0	RsiPlaneExitReason	Exit reason
elr_el2	0x100	Bits64	Exception Link Register
esr_el2	0x108	Bits64	Exception Syndrome Register
far_el2	0x110	Bits64	Fault Address Register
hpfar_el2	0x118	Bits64	Hypervisor IPA Fault Address register
gprs[31]	0x200	Bits64	Registers
gicv3_hcr	0x300	Bits64	GICv3 Hypervisor Control Register value
gicv3_lrs[16]	0x308	Bits64	GICv3 List Register values
gicv3_misr	0x388	Bits64	GICv3 Maintenance Interrupt State Register value
gicv3_vmcr	0x390	Bits64	GICv3 Virtual Machine Control Register value
cntp_ctl	0x400	Bits64	Counter-timer Physical Timer Control Register value
cntp_cval	0x408	Bits64	Counter-timer Physical Timer CompareValue Register value
cntv_ctl	0x410	Bits64	Counter-timer Virtual Timer Control Register value

Name	Byte offset	Type	Description
cntv_cval	0x418	Bits64	Counter-timer Virtual Timer CompareValue Register value

I₀₂₃₂ RsiPlaneExit uses architectural encodings (of ESR, FAR, HPFAR) which are normally observable only to EL2; however, the exit is taken to P0 at EL1. This is justified on the grounds that P0 exists essentially in order to allow part of the job of the hypervisor to be performed inside the Realm.

I₀₂₃₃ On Plane exit, all RsiPlaneExit fields are zero unless specified otherwise.

A10.2.3.1 Plane exit due to synchronous exception

R₀₂₃₄ An exception due to any of the following in Pn causes a Plane exit due to Synchronous Exception:

- Trapped WF* instruction execution
- Data Abort at a Protected IPA
 - Permission fault
 - Access to an IPA whose RIPAS is EMPTY
- Instruction Abort at a Protected IPA
 - Permission fault
 - Access to an IPA whose RIPAS is EMPTY
- HVC instruction execution
- RSI_HOST_CALL execution, if `plane_enter.flags.trap_hc == RSI_TRAP`
- Any other SMC instruction execution

R₀₂₃₅ On Plane exit due to Synchronous Exception, all of the following are true:

- `plane_exit.exit_reason` is `RSI_EXIT_SYNC`.
- `plane_exit.esr_el2` contains the value of `ESR_EL2` at the time of the Realm exit.
- If `plane_exit.esr_el2.EC` indicates Data Abort from a lower Exception level and `plane_exit.esr_el2.ISV == 1` then `plane_exit.far_el2` contains the value of `FAR_EL2` at the time of the Realm exit.
- If `plane_exit.esr_el2.EC` indicates Data Abort from a lower Exception level or Instruction Abort from a lower Exception level, `plane_exit.hpfar_el2` contains the value of `HPFAR_EL2` at the time of the Realm exit.

See also:

- [A4.3.9 REC exit due to Host call](#)

A10.2.4 REC exit from Pn

R₀₂₃₆ An exception due to any of the following in Pn cause a REC exit to the Host:

- The following Synchronous Exceptions:
 - Access to an IPA whose RIPAS is DESTROYED
 - Access to an IPA whose HIPAS is UNASSIGNED and whose RIPAS is not EMPTY
 - Synchronous External Abort
- The following Asynchronous Exceptions:
 - IRQ
 - FIQ
 - SError

- RSI_HOST_CALL execution, if `plane_enter.flags.trap_hc == RSI_NO_TRAP`. In this case, the result is a REC exit due to Host call.

I₀₂₃₇ Any other exception during execution of Pn causes a Plane exit to P0.

A10.2.5 Pn execution of HVC and SMC

I₀₂₃₈ On Plane exit due to execution by Pn of an HVC instruction, possible actions taken by P0 include the following:

- Emulate the instruction
- Forward the request to the Host using RSI_HOST_CALL
- Return SMCCC_NOT_SUPPORTED to Pn.

I₀₂₃₉ On Plane exit due to execution by Pn of an RSI command, possible actions taken by P0 include the following:

- Emulate the RSI command
- Return SMCCC_NOT_SUPPORTED to Pn.

See also:

- [Chapter B5 Realm Services Interface](#)
- [B5.3.4 RSI_HOST_CALL command](#)

A10.2.6 Pn system registers

R₀₂₄₀ On Realm creation, all Pn EL0 and EL1 system register values take architecturally-defined reset values.

I₀₂₄₁ P0 can access Pn EL0 and EL1 system register values using the RSI_PLANE_REG_READ and RSI_PLANE_REG_WRITE commands.

See also:

- [B5.3.13 RSI_PLANE_REG_READ command](#)
- [B5.3.14 RSI_PLANE_REG_WRITE command](#)

A10.3 Planes memory management

Provisional

Decide whether this section should be integrated into the main [Realm memory management](#) chapter.

- I₀₂₄₂ All Planes within a Realm have the same IPA size.
- I₀₂₄₃ If a given IPA x is mapped to a given PA y in one Plane, x is not mapped to a different PA z in any other Plane within the Realm.

A10.3.1 Auxiliary RTT

- I₀₂₄₄ A Realm which is configured to have an RTT tree per Plane has one *primary RTT tree* and (number of auxiliary Planes) *auxiliary RTT trees*.
- D₀₂₄₅ For a Realm which is configured to have an RTT tree per Plane, RTT trees are identified using a zero-based *RTT tree index*.
- The RTT tree index of the primary RTT tree is zero.
- D₀₂₄₆ For a Realm which is configured to have an RTT tree per Plane, the mapping from Plane index to RTT tree index is as follows:

Plane index	RTT tree index
0	1
1	2
...	...
n-2	n-1
n-1	0

For a Realm with no auxiliary Planes, $n == 1$ and therefore the index of the single Plane (0) maps to the index of the single RTT tree (0).

- I₀₂₄₇ Creation, destruction and folding of primary RTTs is independent of creation, destruction and folding of auxiliary RTTs for the same IPA range.
- I₀₂₄₈ If a primary RTT entry is live and its state is not TABLE then the Host can create a mapping to the same output address in an auxiliary RTT by executing RMI_RTT_AUX_MAP_PROTECTED or RMI_RTT_AUX_MAP_UNPROTECTED.
- D₀₂₄₉ An IPA is *auxiliary-live* if any of the entries identified by that IPA in auxiliary RTTs are live.
- I₀₂₅₀ In a Realm which is configured to have an RTT tree per Plane, a given Unprotected IPA may be mapped to different output addresses in different Planes.
- U₀₂₅₁ The absence of a rule which states that a given Unprotected IPA must map to the same output address in different Planes avoids the need for the RMM to manage reference counts for NS Granules.
- I₀₂₅₂ If a Protected IPA is auxiliary-live then the corresponding entry in the primary RTT is live.
- This invariant is preserved by blocking any actions which would make the primary RTT entry non-live, including the following:
- RMI_DATA_DESTROY
 - RMI_DEV_MEM_UNMAP

- `RMI_RTT_SET_RIPAS`

I₀₂₅₃ If an IPA is auxiliary-live then its RIPAS cannot be changed.

X₀₂₅₄ Specifying that RIPAS is invariant while an IPA is auxiliary-live avoids an implementation of `RMI_RTT_SET_RIPAS` having to walk multiple auxiliary RTT trees.

See also:

- [A5.5.8 RTTE liveness and RTT liveness](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.31 RMI_RTT_AUX_CREATE command](#)
- [B4.3.32 RMI_RTT_AUX_DESTROY command](#)
- [B4.3.33 RMI_RTT_AUX_FOLD command](#)
- [B4.3.34 RMI_RTT_AUX_MAP_PROTECTED command](#)
- [B4.3.35 RMI_RTT_AUX_MAP_UNPROTECTED command](#)
- [B4.3.36 RMI_RTT_AUX_UNMAP_PROTECTED command](#)
- [B4.3.37 RMI_RTT_AUX_UNMAP_UNPROTECTED command](#)
- [B4.3.44 RMI_RTT_SET_RIPAS command](#)

A10.3.2 Stage 2 access permissions

This section describes how Stage 2 access permissions (S2AP) are managed for Realm IPA space.

A10.3.2.1 Stage 2 access permissions overview

I₀₂₅₅ The programming model for control of S2AP is based on indirection, as follows:

- Each IPA has an *S2AP base index*.
 - For a Protected IPA, the S2AP base index is fixed.
 - For an Unprotected IPA, the S2AP base index is controlled by the Host via RMI commands.
- The mapping from S2AP base index to S2AP base value is defined by this specification.
- Each IPA has an *S2AP overlay index*.
 - For a Protected IPA, the S2AP overlay index is controlled by P0, via RSI commands.
 - For an Unprotected IPA, the S2AP overlay index is fixed.
- Each { Plane, S2AP overlay index } tuple maps to an *S2AP overlay value*.
 - For Pn, this mapping is controlled by P0 via RSI commands, subject to certain constraints imposed by the RMM.
 - For P0, this mapping is architecturally fixed.
- The S2AP which applies to an access from a given Plane to a given IPA are defined by the combination of the S2AP base value and the S2AP overlay value, following the rules for Stage 2 permission indirection in the [Arm Architecture Reference Manual for A-Profile architecture](#) [3].

I₀₂₅₆ The S2AP which applies to an access from P0 to a Protected IPA is architecturally fixed to *RW+puX*.

I₀₂₅₇ The S2AP which applies to an access from Pn to a Protected IPA is controlled by P0.

I₀₂₅₈ The S2AP which applies to an access from Pn to a Protected IPA can be different for each Pn within the Realm.

I₀₂₅₉ An access by Pn to a Protected IPA which violates S2AP causes a Plane exit taken to P0.

I₀₂₆₀ The S2AP which applies to a Realm access to an Unprotected IPA is controlled by the Host.

I₀₂₆₁ The S2AP which applies to a Realm access to an Unprotected IPA is the same for all Planes within the Realm.

I₀₂₆₂ A data access by Pn to an Unprotected IPA which violates S2AP causes a REC exit taken to the Host.

I₀₂₆₃ An instruction fetch by Pn to an Unprotected IPA causes a Plane exit taken to P0.

U₀₂₆₄ On a platform which implements FEAT_S2PIE and FEAT_S2POE, the RMM can utilise these architecture features to store per-Plane S2AP in a single RTT tree.

On a platform which does not implement these architecture features, a separate RTT tree is required for each Plane.

A10.3.2.2 Stage 2 access permissions for a Protected IPA

R₀₂₆₅ The S2AP overlay index of a Protected IPA is between 0 and 14.

I₀₂₆₆ S2AP overlay index 15 is RESERVED.

R₀₂₆₇ At Realm activation, S2AP overlay indices 0 to 14 map to S2AP overlay value *RW+puX* for P0.

R₀₂₆₈ At Realm activation, S2AP overlay indices 0 to 14 map to S2AP overlay value *NoAccess* for all Planes other than P0.

D₀₂₆₉ For each S2AP overlay index there is an associated lock bit which applies to S2AP overlay values for Planes other than P0.

- If the lock bit is LOCKED then the S2AP overlay values for all Planes are immutable.
- If the lock bit is UNLOCKED then the S2AP overlay values for Planes other than P0 can be changed by P0.

R₀₂₇₀ At Realm activation, S2AP overlay index 0 is LOCKED.

R₀₂₇₁ At Realm activation, S2AP overlay indices 1 to 14 are UNLOCKED.

R₀₂₇₂ At Realm activation, the S2AP overlay index of all Protected IPAs is 0.

I₀₂₇₃ The following table summarises the attributes of all S2AP overlay indices for Protected IPA space, at Realm activation.

S2AP overlay index	P0 S2AP overlay value	Pn S2AP overlay values	Lock status for Pn S2AP overlay values
0	<i>RW+puX</i>	<i>NoAccess</i>	LOCKED
1 to 14	<i>RW+puX</i>	<i>NoAccess</i>	UNLOCKED

I₀₂₇₄ The RSI_MEM_SET_PERM_INDEX command can be used by P0 to change the S2AP overlay index for a Protected IPA.

I₀₂₇₅ The RSI_MEM_SET_PERM_VALUE command can be used by P0 to change the mapping from a { Plane, S2AP overlay index } tuple to an S2AP overlay value.

See also:

- [Chapter A5 Realm memory management](#)
- [A10.3.1 Auxiliary RTT](#)
- [B5.3.10 RSI_MEM_SET_PERM_INDEX command](#)
- [B5.3.11 RSI_MEM_SET_PERM_VALUE command](#)

A10.3.2.3 Stage 2 access permissions for an Unprotected IPA

I₀₂₇₆ The programming model for control by the Host of S2AP for Unprotected IPA space depends on the configuration of the Realm.

If the Realm is configured to use an RTT tree per Plane then S2AP are specified directly, as follows:

- The Host provides the RW field in the descriptor passed to RMI_RTT_MAP_UNPROTECTED to create the mapping in the primary RTT tree.
- Execution of RMI_RTT_AUX_MAP_UNPROTECTED causes the RW and XN fields to be copied from the primary RTT tree into an auxiliary RTT tree.

- The RMM enforces that the XN field of the descriptors in all RTT trees never grants execute permission.

If the Realm is configured to use a single RTT tree then then S2AP are specified indirectly as follows:

- The Host provides an *S2AP base index* in the descriptor passed to RMI_RTT_MAP_UNPROTECTED.
- The RMM configures the S2AP overlay index to provide a S2AP overlay value of RW.
- The observed S2AP is defined by the combination of the S2AP base value and the S2AP overlay value, following the rules for Stage 2 permission indirection in the [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#).

In this way, the RMM ensures that neither unprivileged execute permission (*uX*) nor privileged execute permission (*pX*) is observed on Realm access to an Unprotected IPA.

See also:

- [Chapter A5 Realm memory management](#)
- [A10.2 Planes exception model](#)
- [A10.3.1 Auxiliary RTT](#)
- [B4.3.35 RMI_RTT_AUX_MAP_UNPROTECTED command](#)
- [B4.3.42 RMI_RTT_MAP_UNPROTECTED command](#)
- [B4.4.48 RmiUnprotectedS2AP type](#)

A10.3.2.4 Stage 2 access permissions change

D₀₂₇₇ An *S2AP change* is a process via which the S2AP of a region of Protected IPA space is changed.

I₀₂₇₈ An S2AP change consists of actions taken both by P0 within the Realm and by the Host:

- P0 issues an *S2AP change request* by executing RSI_MEM_SET_PERM_INDEX.
 - The input values to this command include:
 - * The requested IPA range: [base, top)
 - * The requested S2AP overlay index
 - * A “cookie”, whose initial value is zero. Across successive calls to RSI_MEM_SET_PERM_INDEX, the cookie is used by the RMM to store an IMPLEMENTATION DEFINED value which tracks progress of the request.
 - The RMM records these values in the REC, and then performs a REC exit due to S2AP change pending.
- In response, the Host executes zero or more RMI_RTT_SET_S2AP commands.
- If the requested RIPAS value was not EMPTY then at the next RMI_REC_ENTER the Host can optionally indicate that it rejects the S2AP change request.

U₀₂₇₉ The purpose of the cookie is to record the progress of the S2AP change request across multiple RTT trees. For a Realm which is configured to use a shared RTT tree, the RMM should return a cookie value of zero.

X₀₂₈₀ Rejection by the Host of an S2AP change request is intended to be used if the target IPA range extends beyond the agreed DRAM range for the Realm. In this situation, accepting the request may impose unplanned resource costs on the Host, by requiring allocation of additional RTTs.

I₀₂₈₁ The S2AP change process ensures that a Realm can always reliably determine the maximum S2AP which can be observed at the next access to any Protected IPA.

I₀₂₈₂ An S2AP change is applied by one or more calls to the RMI_RTT_SET_S2AP command.

I₀₂₈₃ The order in which the S2AP of Planes and pages within the target IPA range are changed during execution of RSI_MEM_SET_PERM_INDEX is IMPLEMENTATION DEFINED.

I₀₂₈₄ If the input arguments of RSI_MEM_SET_PERM_INDEX are modified by the caller during the loop, it is IMPLEMENTATION DEFINED whether the S2AP of Planes and pages within the target IPA range are changed.

I₀₂₈₅ The P0 programming model for changing S2AP is to call RSI_MEM_SET_PERM_INDEX in a loop until progress reaches the top of the target IPA range, as shown in the following pseudocode:

```

int realm_set_s2ap(unsigned long base, unsigned long top,
                  unsigned int index)
{
    unsigned long new_base, response;
    unsigned long cookie = 0, new_cookie;
    int ret = RSI_SUCCESS;

    while (base != top) {
        ret = rsi_mem_set_perm_index(base, top, index, cookie,
                                    &new_base, &response,
                                    &new_cookie);

        if (ret != RSI_SUCCESS) {
            return ret;
        }

        if (response == RSI_REJECT) {
            return RSI_ERROR_INPUT;
        }

        base = new_base;
        cookie = new_cookie;
    }

    return RSI_SUCCESS;
}

```

I0286

The Host programming model for handling an S2AP change request is to call RMI_RTT_SET_S2AP in a loop until progress reaches the top of the target IPA range, as shown in the following pseudocode:

```

int host_set_s2ap(unsigned long base, unsigned long top)
{
    unsigned long out_top, index, rtt_tree;
    int ret = RMI_SUCCESS;

    while (base != top) {
        ret = rmi_rtt_set_s2ap(rd, rec, base, top,
                              &out_top, &rtt_tree, &index);

        if (ret == RMI_ERROR_RTT || ret == RMI_ERROR_AUX_RTT) {
            create_rtt(ipa = out_top, level = index + 1, rtt_tree);
            continue;
        } else if (ret != RMI_SUCCESS) {
            break;
        }

        base = out_top;
    }

    return ret;
}

```

R0287

On REC entry following a REC exit due to S2AP change, `rec.s2ap_response` is set to the value of `enter.flags.s2ap_response`.

I0288

If all of the following are true then the output value of RSI_MEM_SET_PERM_INDEX indicates “Host rejected the request”:

- `rec.s2ap_addr` is not equal to `rec.s2ap_top`.

- `rec.s2ap_response` is REJECT.

Otherwise, the output value of `RSI_MEM_SET_PERM_INDEX` indicates “Host accepted the request”.

See also:

- [A4.3.13 REC exit due to S2AP change pending](#)
- [A5.4 RIPAS change](#)
- [B3.64 RecS2APResponseToRsi function](#)
- [B4.3.30 RMI_REC_ENTER command](#)
- [B4.3.45 RMI_RTT_SET_S2AP command](#)
- [B5.3.10 RSI_MEM_SET_PERM_INDEX command](#)
- [D1.5.4 S2AP change flow](#)

A10.3.2.5 Stage 2 access permissions reset due to Host action

I₀₂₈₉ Execution of `RMI_RTT_DESTROY` or `RMI_RTT_AUX_DESTROY` sets the S2AP overlay index of the target RTTE to 0.

X₀₂₉₀ The Host is permitted at any time to destroy an RTT. Destruction of an RTT causes the RMM to lose information about the S2AP for the IPA range described by that RTT. The S2AP are observable by the Realm only while the RIPAS is RAM. Therefore, this specification defines the value to which the S2AP overlay index must be reset on a RIPAS transition to RAM.

See also:

- [A5.4 RIPAS change](#)

DRAFT

A10.4 Planes interrupts

Provisional

Decide whether this section should be integrated into the main [Realm interrupts](#) section.

- I₀₂₉₁ On REC creation, the GIC owner for the REC is P0.
- I₀₂₉₂ On Plane entry, P0 can transfer GIC ownership to the target Pn.
- I₀₂₉₃ On Plane entry, if P0 transfer GIC ownership to the target Pn then the GIC state in RsiPlaneEnter is ignored. This means that the GIC state of the owner is preserved and is shared across GIC ownership changes between planes.
- S₀₂₉₄ Allowing P0 to control which Plane is the GIC owner supports software usage models including the following:
- P0 is the GIC owner, and P0 emulates a vGIC for Pn, similar to how the Host emulates a vGIC for the Realm.
 - P0 is a lightweight “Pn switcher”, which does not emulate a vGIC. GIC ownership is transferred to the Pn which contains the main Realm guest OS.
- I₀₂₉₅ P0 can read the value of ICH_VTR_EL2 using the RSI_REALM_CONFIG command.
- I₀₂₉₆ On Plane exit, P0 is the GIC owner. If GIC ownership was transferred to Pn on Plane entry, it returns to P0 on Plane exit.
- I₀₂₉₇ On Plane exit, Pn’s GIC state is exposed to P0 via the RsiPlaneExit object.
- R₀₂₉₈ On REC entry the GIC state provided by the Host is assigned to the GIC owner.
- R₀₂₉₉ On REC entry, if the values of `enter.gicv3_lrs` describe one or more Pending interrupts and the most recent REC exit was from a Plane which is not the GIC owner then control returns to P0. This results in a Plane exit due to synchronous exception.
- R₀₃₀₀ On REC entry, if all of the following is true then control returns to P0, resulting in a Plane exit due to synchronous exception:
- The most recent REC exit was from Pn
 - The most recent REC exit was not from the GIC owner
 - The value of ICH_MISR_EL2 at the time of the REC exit was not zero.
- I₀₃₀₁ On REC exit, the Realm GIC state of the GIC owner Plane is reported to the Host.
- See also:
- [A2.3.2 REC attributes](#)
 - [A4.1 Realm exception model overview](#)
 - [A6.1 Realm interrupts](#)
 - [A10.2.3.1 Plane exit due to synchronous exception](#)
 - [B5.3.12 RSI_PLANE_ENTER command](#)
 - [B5.3.24 RSI_REALM_CONFIG command](#)

A10.5 Planes timers

Provisional

Decide whether this section should be integrated into the main [Realm timers](#) section.

- R₀₃₀₂ On REC exit from P0, the Realm EL1 timer state reported to the Host is P0's EL1 timer state.
- D₀₃₀₃ A Realm EL1 timer is *active* if it is enabled and unmasked.
- R₀₃₀₄ On REC exit from Pn, for each of the EL1 virtual and physical timers, if any of the following is true then the timer state reported to the Host is Pn's EL1 timer state:
- The Pn timer is active and the P0 timer is not active.
 - Both Pn and P0 timers are active and the Pn timer deadline is earlier than the P0 timer deadline.
- Otherwise, the timer state reported to the Host is P0's EL1 timer state.
- I₀₃₀₅ The following table summarises the timer state which is reported to the Host on REC exit.

P0 active	Pn active	Earliest CVAL	Reported to Host
0	0	P0	P0
0	0	Pn	P0
0	1	P0	Pn
0	1	Pn	Pn
1	0	P0	P0
1	0	Pn	P0
1	1	P0	P0
1	1	Pn	Pn

- I₀₃₀₆ On Plane exit, Pn's EL1 timer state is exposed to P0 via the RsiPlaneExit object.
- S₀₃₀₇ P0 software should check the Realm EL1 timer state on every return from RSI_PLANE_ENTER and update virtual interrupt state accordingly. This is true regardless of the value of `exit.exit_reason`: even if the return occurred for a reason unrelated to timers (for example, a Plane exit due to Data Abort), the Realm EL1 timer state should be checked.
- I₀₃₀₈ On Plane entry, the RMM may mask the hardware timer signal, following the same logic as for REC entry.
- U₀₃₀₉ Management of EL1 timer state for a Realm with multiple Planes can be implemented by multiplexing the following into the EL2 hardware timers:
- P0's EL1 timers
 - The Host's EL2 timers
- See also:
- [A6.2 Realm timers](#)
 - [A10.2.3 Plane exit](#)
 - [B5.4.16 RsiPlaneExit type](#)

Chapter A11

Realm memory encryption

This section describes encryption of physical memory which is accessible via Realm PAS. This encryption is transparent to Realm software, but has an impact on the security posture of a Realm.

- D₀₃₁₀ A Memory Encryption Context (MEC) is an encryption regime used to protect the memory owned by a Realm.
The memory protected by a Realm's MEC includes all memory which can be accessed by the Realm, and the RTTs which are owned by the Realm.
- U₀₃₁₁ Other Granules owned by the Realm, such as REC and RD, are protected with the RMM's MEC.
- D₀₃₁₂ A Memory Encryption Context Identifier (MECID) is a handle which is used to identify a MEC.
- I₀₃₁₃ The highest MECID value which the Host is permitted to pass via an RMI command is reported by the RMI_FEATURES command in RmiFeatureRegister1::MAX_MECID.
- D₀₃₁₄ A *valid MECID* is a MECID in the range $[0, \text{MAX_MECID}]$.
- R₀₃₁₅ On a platform which does not implement FEAT_MEC, MAX_MECID is zero.
- U₀₃₁₆ On a platform which implements FEAT_MEC, MAX_MECID is expected to be computed as follows:
- Determine the minimum MECID width supported across all system components capable of initiating Realm PAS transactions.
 - Determine the number of MECIDs which the platform needs to reserve for its own use. This is expected to be at least one, for protection of the RMM's memory.
 - Return $\text{MAX_MECID} = (2^{\text{MECID_WIDTH}}) - (\text{NUM_RESERVED_MECIDS} + 1)$

D₀₃₁₇ A MEC has a *MEC policy*.

The MEC policy values are shown in the following table.

Name	Description
MEC_POLICY_PRIVATE	The MEC protects memory owned by a single Realm. A MEC with this policy may be referred to as a <i>Private MEC</i> .
MEC_POLICY_SHARED	The MEC protects memory owned by multiple Realms. A MEC with this policy may be referred to as a <i>Shared MEC</i> .

D₀₃₁₈ A MEC has a *MEC state*.

The MEC state values are shown in the following table.

Name	Description
MEC_STATE_PRIVATE_ASSIGNED	A Private MEC which is assigned to a Realm.
MEC_STATE_PRIVATE_UNASSIGNED	A Private MEC which is not assigned to a Realm.
MEC_STATE_SHARED	A Shared MEC.

- D₀₃₁₉ A Shared MEC has a set of zero or more member Realms.
- R₀₃₂₀ At platform boot, the state of MEC zero is MEC_STATE_SHARED.
- R₀₃₂₁ At platform boot, the state of every MEC in the range $[1, \text{MAX_MECID}]$ is MEC_STATE_PRIVATE_UNASSIGNED.
- I₀₃₂₂ The input values of RMI_REALM_CREATE include a MECID.
- I₀₃₂₃ RMI_REALM_CREATE fails if any of the following is true:
- The MECID is not valid.
 - The state of the MEC is MEC_STATE_PRIVATE_ASSIGNED.

- I₀₃₂₄ On successful execution of RMI_REALM_CREATE:
- If the state of the MEC is MEC_STATE_PRIVATE_UNASSIGNED then the state becomes MEC_STATE_PRIVATE_ASSIGNED.
 - If the state of the MEC is MEC_STATE_SHARED then the new Realm is added to the Shared MEC.
- I₀₃₂₅ The RMI_MEC_SET_SHARED command changes the state of a MEC from MEC_STATE_PRIVATE_UNASSIGNED to MEC_STATE_SHARED.
- I₀₃₂₆ Execution of RMI_MEC_SET_SHARED fails if any of the following is true:
- The MECID provided by the Host is not valid.
 - The MECID provided by the Host identifies a MEC whose current state is not MEC_STATE_PRIVATE_UNASSIGNED.
 - Any valid MECID identifies a Shared MEC. This means that there can be at most a single Shared MEC in existence at a time.
- I₀₃₂₇ The RMI_MEC_SET_PRIVATE command changes the state of a MEC from MEC_STATE_SHARED to MEC_STATE_PRIVATE_UNASSIGNED.
- I₀₃₂₈ Execution of RMI_MEC_SET_PRIVATE fails if any of the following is true:
- MAX_MECID is zero.
 - The MECID provided by the Host is not valid.
 - The MECID provided by the Host identifies a MEC whose current state is not MEC_STATE_SHARED.
 - The MECID provided by the Host identifies a Shared MEC which contains a non-zero number of Realms.
- I₀₃₂₉ On a platform which reports MAX_MECID to be zero, all Realms use the Shared MEC identified by MECID zero.
- I₀₃₃₀ On a platform which reports MAX_MECID to be non-zero, the Host can choose between the following approaches:
- Use the Shared MEC for all Realms.
 - Assign a Private MEC to each Realm, with the total number of Realms not exceeding MAX_MECID.
 - Use the Shared MEC for some Realms; for the remaining Realms, assign a Private MEC to each, with the total number of this latter set of Realms not exceeding MAX_MECID.
- I₀₃₃₁ The MEC policy of a Realm is reflected in the Realm attestation token.
- R₀₃₃₂ On platform boot, the encryption context associated with every MEC changes.
- R₀₃₃₃ When the state of a MEC changes, the encryption context associated with that MEC changes.
- I₀₃₃₄ To illustrate the points at which encryption contexts must change, consider the following sequence:

Step	Reason for encryption context change
Platform boot	Platform boot
RMI_REALM_CREATE(rd=a, mecid=0)	
RMI_REALM_CREATE(rd=b, mecid=0)	
RMI_REALM_DESTROY(rd_a)	
RMI_REALM_DESTROY(rd_b)	
RMI_MEC_SET_PRIVATE(mecid=0)	MEC_STATE_SHARED to MEC_STATE_PRIVATE_UNASSIGNED
RMI_REALM_CREATE(rd=c, mecid=0)	MEC_STATE_PRIVATE_UNASSIGNED to MEC_STATE_PRIVATE_ASSIGNED
RMI_REALM_DESTROY(rd_c)	MEC_STATE_PRIVATE_ASSIGNED to MEC_STATE_PRIVATE_UNASSIGNED

Step	Reason for encryption context change
RMI_REALM_CREATE(rd=d, mecid=0)	MEC_STATE_PRIVATE_UNASSIGNED to MEC_STATE_PRIVATE_ASSIGNED

See also:

- [Arm Architecture Reference Manual Supplement, The Realm Management Extension \(RME\), for Armv9-A \[2\]](#)
- [A3.14 Support for Realm memory encryption](#)
- [A7.2.3.1.7 Realm MEC policy claim](#)
- [B4.3.11 RMI_MEC_SET_PRIVATE command](#)
- [B4.3.12 RMI_MEC_SET_SHARED command](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [B4.3.26 RMI_REALM_DESTROY command](#)
- [B4.4.17 RmiFeatureRegister1 type](#)

DRAFT

DRAFT

Part B

Interface

Chapter B1

Commands

This chapter describes how RMM commands are defined in this specification.

B1.1 Overview

R_VZRKZ	<p>The RMM exposes the following interfaces to the Host:</p> <ul style="list-style-type: none"> • The <i>Realm Management Interface</i> (RMI)
R_NPLKX	<p>The RMM exposes the following interfaces to a Realm:</p> <ul style="list-style-type: none"> • The <i>Realm Services Interface</i> (RSI) • The <i>Power State Coordination Interface</i> (PSCI) <p>Any other SMC executed by a Realm returns SMCCC_NOT_SUPPORTED.</p>
I_TKQXF	An RMM interface consists of a set of RMM commands.
I_RTRYT	An RMM interface is compliant with the SMC Calling Convention (SMCCC).
R_NNFPH	SMCCC version ≥ 1.2 is required.
X_FDXJG	SMCCC version 1.2 increases the number of SMC64 arguments and return values from 4 to 17. Some RMM commands use more than 4 input or output values.
R_VXJJQ	On a CCA platform which implements FEAT_SVE, SMCCC version ≥ 1.3 is required.
X_KCMSY	SMCCC version 1.3 introduces a bit in the FID which a caller can use to indicate that SVE state does not need to be preserved across the SMC call.
R_JNVJQ	On a CCA platform which implements FEAT_SME, SMCCC version ≥ 1.4 is required.
X_QXMZL	SMCCC version 1.4 adds support for preservation of SME state across an SMC call.
R_KWMVX	An RMM command uses the SMC64 calling convention.
S_DFNMZ	<p>To determine whether an RMM interface is implemented, software should use the following flow:</p> <ol style="list-style-type: none"> 1. Determine whether the SMCCC_VERSION command is implemented, following the procedure described in Arm SMC Calling Convention [16]. 2. Check that the SMCCC version is ≥ 1.1. 3. Execute the <Interface>.Version command, which returns: <ul style="list-style-type: none"> • SMCCC_NOT_SUPPORTED (-1) if <Interface> is not implemented. • A version number (>0) if <Interface> is implemented.
R_YBXKR	<p>All data types defined in this specification are little-endian.</p> <p>See also:</p> <ul style="list-style-type: none"> • Chapter B4 Realm Management Interface • Chapter B5 Realm Services Interface • Chapter B6 Power State Control Interface

B1.2 Command definition

I_{WBMVP}

The definition of an RMM command consists of:

- A *function identifier* (FID)
- A set of *input values* (referred to as “arguments” in SMCCC)
- A set of *output values* (referred to as “results” in SMCCC)
- A set of *context values*
- A partially-ordered set of *failure conditions*
- A set of *success conditions*
- A set of *footprint items*

I_{GCVWC}

Each failure condition, success condition and footprint item has an associated identifier. Identifiers are unique within each of the above groups, within each command.

An identifier has no meaning. It is only a label by which a given condition or footprint item can be referred to.

R_{STJHR}

On calling an RMI or RSI command, any of X1 - X16 which are not specified as input values in the command definition SBZ.

R_{KBWJD}

On return from an RMI or RSI command, any of X0 - X16 which are not specified as output values in the command definition MBZ.

See also:

- SMCCC [Arm SMC Calling Convention](#) [16]

B1.2.1 Example command

I_{NFVGF}

The following command, EXAMPLE_ADD, is an example of how the components of an RMM command definition are presented in this document.

This command takes as an input value the address `params_ptr` of an NS Granule which contains two integer values `x` and `y`. On successful execution of the command:

- The output value `sum` contains the sum of `x` and `y`
- The output value `zero` indicates whether either of `x` or `y` is zero

EXAMPLE_ADD is defined as follows:

Interface

FID

0x042

Input values

Name	Register	Field	Type	Description
<code>fid</code>	X0	[63:0]	UInt64	Command FID
<code>params_ptr</code>	X1	[63:0]	Address	PA of parameters

Context

The EXAMPLE_ADD command operates on the following context.

Name	Type	Value	Before	Description
params	ExampleParams	Params(params_ptr)	false	Parameters

Output values

Name	Register	Field	Type	Description
result	x0	[15:0]	CommandReturnCode	Command return status
sum	x1	[63:0]	UInt64	Sum of x and y
zero	x2	[63:0]	UInt64	Whether either x or y was zero

Failure conditions

ID	Condition
params_align	pre: !AddrIsGranuleAligned(params_ptr) post: ResultEqual(result, ERROR_INPUT)
params_gpt	pre: Granule(params_ptr).gpt != GPT_NS post: ResultEqual(result, ERROR_MEMORY)

Success conditions

ID	Post-condition
sum	sum == params.x + params.y
zero	zero == (params.x == 0) (params.y == 0)

B1.3 Command registers

D _{ZDGNM}	An <i>FID</i> is a value which identifies a particular RMM command.
I _{MJQ GK}	The FID of an RMM command is unique among the RMM commands in an RMM interface.
I _{RVPGY}	An FID is read from general-purpose register X0.
D _{XL SFS}	An <i>input value</i> is a value read by an RMM command from general-purpose registers.
D _{VCDCW}	An <i>output value</i> is a value written by an RMM command to general-purpose registers.
D _{CZLVJ}	A <i>command return code</i> is a value which specifies whether an RMM command succeeded or failed.
I _{FRZFT}	A command return code is written to general-purpose register X0.

B1.4 Command condition expressions

D _{CHRYB}	A <i>condition expression</i> is an expression which evaluates to a boolean value.
--------------------	------------------------------------------------------------------------------------

I_{BPNKQ} Following expansion of macros, a *condition expression* is a valid expression in Arm Specification Language (ASL).
See also:

- [Arm Specification Language Reference Manual \[17\]](#)
- [Chapter B3 Command condition functions](#)

B1.5 Command context values

D_{DLBYC} A *context value* is a value which is derived from the value of a command input register and which is used by a command condition expression.

I_{VKKKY} A context value can be thought of as a local variable for use by command condition expressions.
For example, consider the following example command condition expression:

```
!AddrIsGranuleAligned(RealmParams(params_ptr).rtt_base)
```

By introducing a context value `params` with the value `RealmParams(params_ptr)`, this command condition expression can be re-written as:

```
!AddrIsGranuleAligned(params.rtt_base)
```

D_{QDFNW} The *before* property of a context value indicates whether its expression is re-evaluated after the command has executed.

- `before = true`: the expression is not re-evaluated after the command has executed
- `before = false`: the expression is re-evaluated after the command has executed

I_{LTLQN} Specifying `before = true` for a context value allows system state to be sampled before command execution, and then used after command execution in a command success condition.

For example, the `RMI_REALM_DESTROY` command takes as an input value the address `rd` of a Realm Descriptor. Successful execution of the command results observable effects including the following:

- The state of the RD Granule changes from `RD` to `DELEGATED`
- The state of the RTT base Granule, whose address was previously held in the RD, changes from `RTT` to `DELEGATED`

The address of the RTT base Granule is not included in the input values of the command.

A context value is defined as follows:

Name	Type	Value	Before	Description
<code>rtt_base</code>	Address	<code>Realm(rd).rtt_base</code>	<code>true</code>	RTT base address

The state change of the RTT Granule can then be expressed as:

```
Granule(rtt_base).state == DELEGATED
```

I_{YNDGD} The *before* property of a context value has no effect if the value is only used in command failure conditions.

D_{XBHPB} An *in-memory value* is a value passed to a command via an in-memory data structure, the address of which is passed in an input register.

I_{ZTYSS} An in-memory value is a context value.

See also:

- [B4.3.25 RMI_REALM_CREATE command](#)

B1.6 Command failure conditions

D_{DNQOC} An RMM command *failure condition* defines a way in which the command can fail.

I_{GVBBZ} A failure condition consists of a *pre-condition* and a *post-condition*.

I_{WTSZH} A failure pre-condition can be thought of as the “trigger” of the failure: if the pre-condition is true then the command fails.

I_{KJHNX} A failure post-condition can be thought of as the “effect” of the failure: if the command failed due to a particular trigger, then the post-condition defines the error code which is returned.

I_{CVTGY} A failure pre-condition is a condition expression whose terms can include input values and context values.

I_{HNDNN} A failure post-condition is a condition expression whose terms can include input values and context values.

I_{KHJDY} Observability of the checking of command failure conditions is subject to a partial order.

An ordering relation “*A* precedes *B*” means either of the following:

- The pre-condition of *B* is well-formed only if the pre-condition of *A* is false. This is referred to as a *well-formedness ordering*.
- If the pre-conditions of *A* and *B* are both true, then the post-condition of *A* is observed. This is referred to as a *behavioral ordering*.

The absence of an ordering relation “*A* precedes *B*” means that, if the pre-conditions of *A* and *B* are both true then either the post-condition of *A* is observed or the post-condition of *B* is observed.

Orderings are specified between groups of failure conditions. For example, the expression $[A, B] < [C, D]$ means that both conditions *A* and *B* precede both conditions *C* and *D*.

The same information is also presented graphically, with failure conditions represented as nodes and ordering relations represented as edges.




Figure B1.1

The specification does not state whether an individual ordering relation is a well-formedness ordering or a behavioral ordering.

I _{JMTTY}	A given implementation of the RMM is expected to have deterministic behavior. That is, for a runtime instance of the RMM in a particular state, two executions of a command without an interleaving of other commands, with the same input values, results in the same outcome (either success, or the same failure condition.)
R _{WXZJJ}	If a failure pre-condition evaluates to true then the corresponding failure post-condition evaluates to true.
R _{DDGDW}	If a failure pre-condition evaluates to true then the command is aborted.
R _{TFZMS}	If a command fails then all output values except for X0 are UNDEFINED, unless stated otherwise.
R _{VHFHD}	If no failure pre-condition evaluates to true then the command succeeds.

B1.7 Command success conditions

D _{SZGNZ}	An RMM command <i>success condition</i> defines an observable effect of a successful execution of the command.
I _{LZXHB}	A success condition is a condition expression whose terms can include input values, context values and output values.
I _{NMCSF}	The order in which success conditions are listed has no architectural significance.
I _{NJQFG}	If an RMM command succeeds then the return code is <Interface>_SUCCESS.
R _{MKRVV}	If an RMM command succeeds then all of its success conditions evaluate to true.

B1.8 Concrete and abstract types

D _{NXQWV}	<p>A <i>concrete type</i> is a type which has a defined encoding.</p> <p>Examples of concrete types include:</p> <ul style="list-style-type: none"> • An integer which has a defined bit width. • An enumeration within which each label is associated with a unique binary value. • A struct which has a defined width, and within which each member has a defined position. The type of each member of a concrete struct is a concrete type.
I _{WDGMW}	Concrete types are used to define command input values and output values.
D _{WTCVJ}	<p>An <i>abstract type</i> is a type which does not have a defined encoding.</p> <p>Examples of concrete types include:</p> <ul style="list-style-type: none"> • An integer which does not have a defined bit width. • An enumeration which has a set of labels, but which does not define a binary value for each label. • A struct which has a set of members, but which does not define a struct width nor a position for each member. The type of each member of an abstract struct is an abstract type.
I _{QZRGY}	Abstract types are used to model the internal state of the RMM.
I _{LMKGP}	<p>A command failure condition or success condition may need to test for logical equality between a concrete type and a corresponding abstract type. For example, the command may set the value of an internal RMM variable to match the value of a command input. To enable such comparisons, the specification defines an <code>Equal()</code> function for each pair of corresponding concrete and abstract types.</p> <p>See also:</p> <ul style="list-style-type: none"> • B3.26 Equal function

B1.9 Command footprint

D _{ZDJDB}	The <i>footprint</i> of an RMM command defines the set of state items which successful execution of the command can modify.
I _{XMZYS}	The footprint of an RMM command may include state items which are not modified by successful execution of the command.
I _{RWQMJ}	If an RMM command changes the state of a Granule then the footprint typically does not include all attributes of the object which is created or destroyed. For example, the footprint of RMI_REALM_CREATE includes the state of the RD Granule, but does not include attributes of the newly-created Realm.
R _{WZYBV}	Except for items in the footprint of an RMM command and registers in the output values of the RMM command, execution of the command does not have any observable effects.

B1.10 Command testing

I _{MBNZM}	Command definitions can be used to generate testbenches which check whether an implementation complies with the specified failure and success conditions.
--------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------

I _{JGGJN}	A testbench for the EXAMPLE_ADD command presented above would look similar to the following:
--------------------	----------------------------------------------------------------------------------------------

```
// Test EXAMPLE_ADD command
Test_ExampleAdd(Registers regs_in)

    // Unpack input values
    RmmPa params_ptr = regs_in.X1;

    // Evaluate context values
    ExampleParams params = ExampleParams(params_ptr);

    // Evaluate failure pre-conditions
    boolean params_align_pre = !AddrIsGranuleAligned(params_ptr);
    boolean params_gpt_pre = Granule(params_ptr).gpt != GPT_NS;

    // Execute command
    regs_out = RmiExampleAdd(regs_in);

    // Pack output values
    CommandReturnCode result = regs_out.X0;
    integer sum = regs_out.X1;
    integer zero = regs_out.X2;

    // Check return code
    boolean success = (result == Success);

    // Evaluate failure post-conditions
    boolean params_align_post = result == Status(ErrorInput, 1);
    boolean params_gpt_post = result == Status(ErrorInputMemory, 0);

    // Evaluate success conditions
    boolean sum_post = sum == params.x + params.y;
    boolean zero_post = zero == (params.x == 0) || (params.z == 0);

    // Check failure conditions, in order specified
    assert params_align_pre IMPLIES params_align_post;
    assert (!params_align_pre && params_gpt_pre) IMPLIES params_gpt_post;

    // Check that, if no failure pre-condition was violated, command succeeded
    assert (!params_align_pre && !params_gpt_pre) IMPLIES success;
```

```
// Check success conditions, without any ordering
assert success IMPLIES S01_post;
assert success IMPLIES S02_post;
```

Note that the syntax $x \text{ IMPLIES } y$, which is logically equivalent to $!x \vee y$, is not yet defined in ASL.

DRAFT

Chapter B2

Interface versioning

This section describes how the RMI and RSI interfaces are versioned, and how the caller of each can determine whether there exists a mutually acceptable revision of the interface via which it can communicate with the RMM.

Other interfaces exposed by the RMM, such as PSCI, may define their own versioning schemes which differ from that used by RMI and RSI. For details, refer to the specification of the interface concerned.

\mathcal{I}_{LZVQR}

Revisions of the RMI and the RSI are identified by a (major, minor) version tuple.

The semantics of this version tuple are as follows. For two revisions of the interface $P = (maj_P, min_P)$ and $Q = (maj_Q, min_Q)$:

- If $maj_P \neq maj_Q$ then the two interfaces may contain incompatible commands.
- If $maj_P == maj_Q$ and $min_P < min_Q$ then:
 - Every command defined in P has the same behavior in Q, when called with input values that are specified as valid in P.
 - A command defined in P may accept additional input values in Q. These could be provided via any of:
 - * Input registers which were unused in P.
 - * Input memory locations which were specified as SBZ in P.
 - * Encodings which were specified as reserved in P.
 - A command defined in P may return additional output values in Q. These could be returned via any of:
 - * Output registers which were unused in P.
 - * Output memory locations which were specified as MBZ in P.
 - * Encodings which were specified as reserved in P.
 - Q may contain additional commands which are not present in P.
- P is *less than* Q if one of the following conditions is true:
 - $maj_P < maj_Q$
 - $maj_P == maj_Q$ and $min_P < min_Q$

 \mathcal{I}_{ZCPBC}

For each interface, an RMM implementation supports a set of revisions. The size of this set is at least one.

 \mathcal{I}_{RMSLZ}

If an RMM implementation supports a given interface revision (x, y) then Arm expects that it will also supports all earlier revisions with the same major version number. That is:

$(x, 0), (x, 1) \dots (x, y-1), (x, y)$.

A possible exception to this may occur if a security vulnerability is discovered in a particular revision of the interface. For example, if interface revision (x, bad) is found to contain a vulnerability then an RMM implementation may choose to support the following set of revisions:

$(x, 0), (x, 1) \dots (x, bad-1), (x, bad+1) \dots (x, y-1), (x, y)$.

 \mathcal{I}_{GLDQG}

The set of interface revisions supported by an RMM implementation may include revisions with different major version numbers, for example:

$(1, 0), (1, 1) \dots (1, m)$

$(2, 0), (2, 1) \dots (2, n)$

 \mathcal{I}_{JNVXJ}

The RMI_VERSION and RSI_VERSION commands allow the caller and the RMM to determine whether there exists a mutually acceptable revision of the interface via which the two components can communicate.

In each case:

- The caller provides a requested interface revision.
- The output values include a status code and two revisions which are supported by the RMM: a *lower revision* and a *higher revision*.
- The *higher revision* value is the highest interface revision which is supported by the RMM.
- The *lower revision* is less than or equal to the *higher revision*.

The status code and *lower revision* output values indicate which of the following is true, in order of precedence:

- a) The RMM supports an interface revision which is compatible with the requested revision.
 - The status code is “success”.

- The *lower revision* is equal to the requested revision.
- b) The RMM does not support an interface revision which is compatible with the requested revision The RMM supports an interface revision which is incompatible with and less than the requested revision.
 - The status code is “failure”.
 - The *lower revision* is the highest interface revision which is both less than the requested revision and supported by the RMM.
- c) The RMM does not support an interface revision which is compatible with the requested revision The RMM supports an interface revision which is incompatible with and greater than the requested revision.
 - The status code is “failure”.
 - The *lower revision* is equal to the *higher revision*.

The following table shows how each of a set of example scenarios maps onto the above outcomes.

Scenario	Revisions supported by RMM	Revision requested by caller	Outcome	“Lower revision” output value	“Higher revision” output value
1	(1, 0)	(1, 0)	Success (a)	(1, 0)	(1, 0)
2	(1, 0), (1, 1)	(1, 0)	Success (a)	(1, 0)	(1, 1)
3	(1, 0), (2, 0)	(1, 0)	Success (a)	(1, 0)	(2, 0)
4	(1, 0)	(1, 1)	Failure (b)	(1, 0)	(1, 0)
5	(1, 0), (1, 1)	(1, 2)	Failure (b)	(1, 1)	(1, 1)
6	(1, 0), (1, 1)	(2, 0)	Failure (b)	(1, 1)	(1, 1)
7	(1, 0), (1, 1), (1, 3)	(1, 2)	Failure (b)	(1, 1)	(1, 3)
8	(1, 0)	(2, 0)	Failure (b)	(1, 0)	(1, 0)
9	(1, 0)	(2, 1)	Failure (b)	(1, 0)	(1, 0)
10	(1, 0), (1, 1)	(2, 0)	Failure (b)	(1, 1)	(1, 1)
11	(1, 0), (1, 1)	(2, 1)	Failure (b)	(1, 1)	(1, 1)
12	(1, 0), (1, 1), (2, 0)	(2, 1)	Failure (b)	(2, 0)	(2, 0)
13	(2, 0)	(1, 0)	Failure (c)	(2, 0)	(2, 0)
14	(2, 0)	(1, 1)	Failure (c)	(2, 0)	(2, 0)
15	(2, 0), (2, 1)	(1, 0)	Failure (c)	(2, 1)	(2, 1)

See also:

- [B4.1 RMI version](#)
- [B4.3.55 RMI_VERSION command](#)
- [B5.1 RSI version](#)
- [B5.3.25 RSI_VERSION command](#)

Chapter B3

Command condition functions

This chapter describes functions which are used in command condition expressions.

See also:

- [B1.4 Command condition expressions](#)

B3.1 AddrInRange function

Returns TRUE if `addr` is within `[base, base+size]`.

```
func AddrInRange(  
    addr : Address,  
    base : Address,  
    size : integer) => boolean  
begin  
    return ((UInt(addr) >= UInt(base))  
        && (UInt(addr) <= UInt(base) + size));  
end
```

B3.2 AddrIsAligned function

Returns TRUE if address `addr` is aligned to an `n` byte boundary.

```
func AddrIsAligned(  
    addr : Address,  
    n : integer) => boolean
```

B3.3 AddrIsAuxLive function

Returns TRUE if IPA `addr` is *auxiliary-live*, that is live in any auxiliary RTT.

```
func AddrIsAuxLive(  
    addr : Address,  
    realm : RmmRealm) => boolean
```

B3.4 AddrIsGranuleAligned function

Returns TRUE if address `addr` is aligned to the size of a Granule.

```
func AddrIsGranuleAligned(  
    addr : Address) => boolean
```

```
func AddrIsGranuleAligned(  
    addr : integer) => boolean
```

See also:

- [A2.2 Granule](#)

B3.5 AddrIsProtected function

Returns TRUE if address `addr` is a Protected IPA for `realm`.

```
func AddrIsProtected(  
    addr : Address,  
    realm : RmmRealm) => boolean  
begin  
    return UInt(addr) < 2^(realm.ipa_width - 1);  
end
```

B3.6 AddrIsRttLevelAligned function

Returns TRUE if Address `addr` is aligned to the size of the address range described by an RTTE in a level `level` RTT.

Returns FALSE if `level` is invalid.

```
func AddrIsRttLevelAligned(  
    addr : Address,  
    level : integer) => boolean
```

B3.7 AddrIsWithin function

Returns TRUE if address `addr` is within the outer range [`base`, `top`).

```
func AddrIsWithin(  
    addr : Address,  
    base : Address,  
    top : Address) => boolean  
begin  
    var addr_int : integer = UInt(addr);  
    var top_int : integer = UInt(top);  
    var base_int : integer = UInt(base);  
    return ((UInt(addr) >= UInt(base))  
        && (UInt(addr) < UInt(top)));  
end
```

B3.8 AddrRangelsAuxLive function

Returns TRUE if any IPA in range [base, top) is *auxiliary-live*, that is live in any auxiliary RTT.

```
func AddrRangeIsAuxLive(  
    base : Address,  
    top : Address,  
    realm : RmmRealm) => boolean
```

B3.9 AddrRangelsProtected function

Returns TRUE if all addresses in range [base, top) are Protected IPAs for realm.

```
func AddrRangeIsProtected(  
    base : Address,  
    top : Address,  
    realm : RmmRealm) => boolean  
begin  
    var size = UInt(top) - UInt(base);  
    return (AddrIsProtected(base, realm)  
        && size > 0  
        && size < 2^realm.ipa_width  
        && AddrIsProtected(ToAddress(UInt(top) - 1), realm));  
end
```

B3.10 AddrRangelsWithin function

Returns TRUE if all addresses in the inner range [inner_base, inner_top) are within the outer range [outer_base, outer_top).

```
func AddrRangeIsWithin(  
    inner_base : Address,  
    inner_top : Address,  
    outer_base : Address,  
    outer_top : Address) => boolean  
begin  
    return (AddrIsWithin(inner_base, outer_base, outer_top)  
        && AddrIsWithin(inner_top, outer_base, outer_top));  
end
```

B3.11 AlignDownToRttLevel function

Round down addr to align to the size of the address range described by an RTTE in a level level RTT.

```
func AlignDownToRttLevel(  
    addr : Address,  
    level : integer) => Address
```

B3.12 AlignUpToRttLevel function

Round up addr to align to the size of the address range described by an RTTE in a level level RTT.

```
func AlignUpToRttLevel(  
    addr : Address,  
    level : integer) => Address
```

B3.13 AuxAlias16 function

Returns TRUE if any of the first `count` entries in a list of auxiliary Granule addresses are aliased - either among themselves, or with the address of another RMM object.

```
func AuxAlias16(  
    obj : Address,  
    aux : array [16] of Address,  
    count : integer) => boolean  
begin  
    assert 0 <= count && count <= 16;  
    var sorted = AuxSort(aux, count);  
  
    for i = 0 to count - 1 do  
        if sorted[i] == obj then  
            return TRUE;  
        end  
        if i >= 1 && sorted[i] == sorted[i - 1] then  
            return TRUE;  
        end  
    end  
    return FALSE;  
end
```

B3.14 AuxAlias32 function

Returns TRUE if any of the first `count` entries in a list of auxiliary Granule addresses are aliased - either among themselves, or with the address of another RMM object.

```
func AuxAlias32(  
    obj : Address,  
    aux : array [32] of Address,  
    count : integer) => boolean  
begin  
    assert 0 <= count && count <= 32;  
    var sorted = AuxSort(aux, count);  
  
    for i = 0 to count - 1 do  
        if sorted[i] == obj then  
            return TRUE;  
        end  
        if i >= 1 && sorted[i] == sorted[i - 1] then  
            return TRUE;  
        end  
    end  
    return FALSE;  
end
```

B3.15 AuxAligned16 function

Returns TRUE if the first `count` entries in a list of auxiliary Granule addresses are aligned to the size of a Granule.

```
func AuxAligned16(  
    aux : array [16] of Address,  
    count : integer) => boolean  
begin  
    assert 0 <= count && count <= 16;  
    for i = 0 to count - 1 do  
        if !AddrIsGranuleAligned(aux[i]) then
```

```
        return FALSE;
    end
end
return TRUE;
end
```

B3.16 AuxAligned32 function

Returns TRUE if the first `count` entries in a list of auxiliary Granule addresses are aligned to the size of a Granule.

```
func AuxAligned32(
    aux : array [32] of Address,
    count : integer) => boolean
begin
    assert 0 <= count && count <= 32;
    for i = 0 to count - 1 do
        if !AddrIsGranuleAligned(aux[i]) then
            return FALSE;
        end
    end
    return TRUE;
end
```

B3.17 AuxEqual16 function

Returns TRUE if the first `count` entries in two lists of auxiliary Granule addresses are equal.

```
func AuxEqual16(
    aux1 : array [16] of Address,
    aux2 : array [16] of Address,
    count : integer) => boolean
begin
    assert 0 <= count && count <= 16;
    for i = 0 to count - 1 do
        if aux1[i] != aux2[i] then
            return FALSE;
        end
    end
    return TRUE;
end
```

B3.18 AuxEqual32 function

Returns TRUE if the first `count` entries in two lists of auxiliary Granule addresses are equal.

```
func AuxEqual32(
    aux1 : array [32] of Address,
    aux2 : array [32] of Address,
    count : integer) => boolean
begin
    assert 0 <= count && count <= 32;
    for i = 0 to count - 1 do
        if aux1[i] != aux2[i] then
            return FALSE;
        end
    end
    return TRUE;
end
```

B3.19 AuxSort function

Sort first `count` entries in array of auxiliary Granule addresses.

```
func AuxSort(  
    addrs : array [16] of Address,  
    count : integer) => array [16] of Address
```

```
func AuxSort(  
    addrs : array [32] of Address,  
    count : integer) => array [32] of Address
```

B3.20 AuxStateEqual16 function

Returns TRUE if the state of the first `count` entries in a list of auxiliary Granule addresses is equal to `state`.

```
func AuxStateEqual16(  
    aux : array [16] of Address,  
    count : integer,  
    state : RmmGranuleState) => boolean  
begin  
    assert 0 <= count && count <= 16;  
    for i = 0 to count - 1 do  
        if (!PaIsDelegable(aux[i])  
            || GranuleAt(aux[i]).state != state) then  
            return FALSE;  
        end  
    end  
    return TRUE;  
end
```

B3.21 AuxStateEqual32 function

Returns TRUE if the state of the first `count` entries in a list of auxiliary Granule addresses is equal to `state`.

```
func AuxStateEqual32(  
    aux : array [32] of Address,  
    count : integer,  
    state : RmmGranuleState) => boolean  
begin  
    assert 0 <= count && count <= 32;  
    for i = 0 to count - 1 do  
        if (!PaIsDelegable(aux[i])  
            || GranuleAt(aux[i]).state != state) then  
            return FALSE;  
        end  
    end  
    return TRUE;  
end
```

B3.22 AuxStates function

Inductive function which identifies the states of the first `count` entries in a list of auxiliary Granules.

This function is used in the definition of command footprint.

```
func AuxStates(  
    aux : array [16] of Address,  
    count : integer)
```

```
func AuxStates(  
    aux : array [32] of Address,  
    count : integer)
```

B3.23 CurrentRealm function

Returns the current Realm.

```
func CurrentRealm() => RmmRealm
```

B3.24 CurrentRec function

Returns the current REC.

```
func CurrentRec() => RmmRec
```

B3.25 DeviceCommunicate function

Process device communication data and return the new state of the device transaction.

```
func DeviceCommunicate(  
    pdev : RmmPdev,  
    data : RmiDevCommData) => RmmDevCommState
```

```
func DeviceCommunicate(  
    vdev : RmmVdev,  
    data : RmiDevCommData) => RmmDevCommState
```

```
func DeviceCommunicate(  
    vdev : RmmRdev) => RmmDevCommState
```

B3.26 Equal function

Check whether concrete and abstract values are equal

```
func Equal(  
    abstract : RmmFeature,  
    concrete : RmiFeature) => boolean
```

```
func Equal(  
    concrete : RmiFeature,  
    abstract : RmmFeature) => boolean
```

```
func Equal(  
    abstract : RmmHashAlgorithm,  
    concrete : RmiHashAlgorithm) => boolean
```

```
func Equal(  
    concrete : RmiHashAlgorithm,  
    abstract : RmmHashAlgorithm) => boolean
```

```
func Equal(  
    abstract : RmmLfaPolicy,  
    concrete : RmiLfaPolicy) => boolean
```

```
func Equal(  
    concrete : RmiLfaPolicy,  
    abstract : RmmLfaPolicy) => boolean
```

```
func Equal (
    abstract : RmmPdevProtConfig,
    concrete : RmiPdevProtConfig) => boolean
```

```
func Equal (
    concrete : RmiPdevProtConfig,
    abstract : RmmPdevProtConfig) => boolean
```

```
func Equal (
    abstract : RmmPdevState,
    concrete : RmiPdevState) => boolean
```

```
func Equal (
    concrete : RmiPdevState,
    abstract : RmmPdevState) => boolean
```

```
func Equal (
    abstract : RmmPlaneRttFeature,
    concrete : RmiPlaneRttFeature) => boolean
```

```
func Equal (
    concrete : RmiPlaneRttFeature,
    abstract : RmmPlaneRttFeature) => boolean
```

```
func Equal (
    abstract : RmmRecRunnable,
    concrete : RmiRecRunnable) => boolean
```

```
func Equal (
    concrete : RmiRecRunnable,
    abstract : RmmRecRunnable) => boolean
```

```
func Equal (
    abstract : RmmRipas,
    concrete : RmiRipas) => boolean
```

```
func Equal (
    concrete : RmiRipas,
    abstract : RmmRipas) => boolean
```

```
func Equal (
    abstract : RmmVdevState,
    concrete : RmiVdevState) => boolean
```

```
func Equal (
    concrete : RmiVdevState,
    abstract : RmmVdevState) => boolean
```

```
func Equal (
    abstract : RmmRdevState,
    concrete : RsiDeviceState) => boolean
```

```
func Equal (
    concrete : RsiDeviceState,
    abstract : RmmRdevState) => boolean
```

```
func Equal (
    abstract : RmmFeature,
    concrete : RsiFeature) => boolean
```

```
func Equal (
```

```

        concrete : RsiFeature,
        abstract : RmmFeature) => boolean
    
```

```

func Equal (
    abstract : RmmHashAlgorithm,
    concrete : RsiHashAlgorithm) => boolean
    
```

```

func Equal (
    concrete : RsiHashAlgorithm,
    abstract : RmmHashAlgorithm) => boolean
    
```

```

func Equal (
    abstract : RmmRipas,
    concrete : RsiRipas) => boolean
    
```

```

func Equal (
    concrete : RsiRipas,
    abstract : RmmRipas) => boolean
    
```

```

func Equal (
    abstract : RmmRipasChangeDestroyed,
    concrete : RsiRipasChangeDestroyed) => boolean
    
```

```

func Equal (
    concrete : RsiRipasChangeDestroyed,
    abstract : RmmRipasChangeDestroyed) => boolean
    
```

See also:

- [B1.8 Concrete and abstract types](#)

B3.27 FeatureToRmi function

Convert feature bit to RMI type.

```

func FeatureToRmi (
    value : RmmFeature) => RmiFeature
begin
    case value of
        when FEATURE_FALSE => return RMI_FEATURE_FALSE;
        when FEATURE_TRUE  => return RMI_FEATURE_TRUE;
    end
end
end
    
```

B3.28 FeatureToRsi function

Convert feature bit to RSI type.

```

func FeatureToRsi (
    value : RmmFeature) => RsiFeature
begin
    case value of
        when FEATURE_FALSE => return RSI_FEATURE_FALSE;
        when FEATURE_TRUE  => return RSI_FEATURE_TRUE;
    end
end
end
    
```

B3.29 Gicv3ConfigIsValid function

Returns TRUE if the values of all gicv3_* attributes are valid.

```
func Gicv3ConfigIsValid(  
    gicv3_hcr : bits(64),  
    gicv3_lrs : array [16] of bits(64)) => boolean
```

See also:

- [A6.1 Realm interrupts](#)
- [B4.4.33 RmiRecEnter type](#)

B3.30 GranuleAccessPermitted function

Returns TRUE if the Granule located at physical address addr is accessible via pas.

```
func GranuleAccessPermitted(  
    addr : Address,  
    pas : RmmPhysicalAddressSpace) => boolean  
begin  
    case GranuleAt(addr).gpt of  
        when GPT_NS      => return (pas == PAS_NS);  
        when GPT_REALM   => return (pas == PAS_REALM);  
        when GPT_SECURE  => return (pas == PAS_SECURE);  
        when GPT_ROOT    => return (pas == PAS_ROOT);  
        when GPT_AAP     => return TRUE;  
    end  
end
```

B3.31 GranuleAt function

Returns the Granule located at physical address addr.

```
func GranuleAt(  
    addr : Address) => RmmGranule
```

See also:

- [A2.2 Granule](#)

B3.32 ImplFeatures function

Returns features supported by the implementation.

```
func ImplFeatures() => RmmFeatures
```

See also:

- [Chapter A3 Feature discovery and configuration](#)

B3.33 MecMembers function

Returns number of Realms which are members of a given MEC.

```
func MecMembers(  
    mecid : bits(64)) => integer
```

See also:

- [Chapter A11 Realm memory encryption](#)

B3.34 MecPolicy function

Returns policy associated with a given MEC.

```
func MecPolicy(  
    mecid : bits(64)) => RmmMecPolicy  
begin  
    case MecState(mecid) of  
        when MEC_STATE_SHARED           => return MEC_POLICY_SHARED;  
        when MEC_STATE_PRIVATE_ASSIGNED => return MEC_POLICY_PRIVATE;  
        when MEC_STATE_PRIVATE_UNASSIGNED => return MEC_POLICY_PRIVATE;  
    end  
end
```

See also:

- [Chapter A11 Realm memory encryption](#)

B3.35 MecState function

Returns state of a given MEC.

```
func MecState(  
    mecid : bits(64)) => RmmMecState
```

See also:

- [Chapter A11 Realm memory encryption](#)

B3.36 MemPermLabelSupported function

Returns TRUE if the specified value is a valid encoding for a memory permission label and the label is supported by the implementation.

```
func MemPermLabelSupported(  
    label : bits(64)) => boolean
```

B3.37 MinAddress function

Returns the smaller of two addresses.

```
func MinAddress(  
    addr1 : Address,  
    addr2 : Address) => Address  
begin  
    return ToAddress(Min(UInt(addr1), UInt(addr2)));  
end
```

B3.38 MpidrEqual function

Returns TRUE if the specified MPIDR values are logically equivalent.

```
func MpidrEqual(  
    rmm_mpidr : bits(64),  
    rmi_mpidr : RmiRecMpidr) => boolean  
begin
```



```
return (rmm_mpidr[ 3: 0] == rmi_mpidr.aff0
      && rmm_mpidr[15: 8] == rmi_mpidr.aff1
      && rmm_mpidr[23:16] == rmi_mpidr.aff2
      && rmm_mpidr[31:24] == rmi_mpidr.aff3);
end
```

B3.39 MpidrIsUsed function

Returns TRUE if the specified MPIDR value identifies a REC in the current Realm.

```
func MpidrIsUsed(
    mpidr : bits(64)) => boolean
```

B3.40 PalsDelegable function

Returns TRUE if the Granule located at physical address `addr` is delegable.

```
func PaIsDelegable(
    addr : Address) => boolean
```

B3.41 PdevAt function

Returns the PDEV object located at physical address `addr`.

```
func PdevAt (
    addr : Address) => RmmPdev
```

B3.42 PdevAuxCount function

Returns the number of auxiliary Granules required for a PDEV with the specified flags.

The return value is guaranteed not to be greater than 32.

For a given flags value, this function always returns the same value.

```
func PdevAuxCount (
    flags : RmiPdevFlags) => integer
```

B3.43 PdevFlags function

Get RmiPdevFlags value.

```
func PdevFlags(
    pdev : RmmPdev) => RmiPdevFlags
begin
    var flags : RmiPdevFlags;

    case pdev.prot_config of
        when PDEV_IOCOH_E2E_IDE => flags.prot_config = RMI_PDEV_IOCOH_E2E_IDE;
        when PDEV_IOCOH_E2E_SYS => flags.prot_config = RMI_PDEV_IOCOH_E2E_SYS;
        when PDEV_FCOH_E2E_IDE  => flags.prot_config = RMI_PDEV_FCOH_E2E_IDE;
        when PDEV_FCOH_E2E_SYS  => flags.prot_config = RMI_PDEV_FCOH_E2E_SYS;
    end

    return flags;
end
```

B3.44 PlaneRegIsValid function

Whether encoding identifies a Plane register.

```
func PlaneRegIsValid(  
    realm : RmmRealm,  
    encoding : bits(64)) => boolean
```

B3.45 PlaneRegValue function

Value of a Plane register.

```
func PlaneRegValue(  
    realm : RmmRealm,  
    plane_idx : integer,  
    encoding : bits(64)) => bits(64)
```

B3.46 PsciReturnCodeEncode function

Return encoding for a PsciReturnCode value.

```
func PsciReturnCodeEncode(  
    value : PsciReturnCode) => bits(64)
```

B3.47 PsciReturnCodePermitted function

Whether a PSCI return code is permitted.

```
func PsciReturnCodePermitted(  
    calling_rec : RmmRec,  
    target_rec : RmmRec,  
    value : PsciReturnCode) => boolean  
begin  
    if value == PSCI_SUCCESS then  
        return TRUE;  
    end  
  
    var fid : bits(64) = calling_rec.gprs[0];  
  
    // Host is permitted to deny a PSCI_CPU_ON request, if the target  
    // CPU is not already on.  
    if (fid == FID_PSCI_CPU_ON  
        && target_rec.flags.runnable != RUNNABLE  
        && value == PSCI_DENIED) then  
        return TRUE;  
    end  
  
    return FALSE;  
end
```

See also:

- [A4.3.7 REC exit due to PSCI](#)
- [B4.3.23 RMI_PSCI_COMPLETE command](#)

B3.48 RdevFromId function

Returns any RDEV identified by dev_id and assigned to realm.

```
func RdevFromId(  
    realm : RmmRealm,  
    vdev_id : bits(64)) => RmmRdev
```

B3.49 RdevFromIds function

Returns the RDEV identified by the tuple (dev_id, inst_id) and assigned to realm.

```
func RdevFromIds(  
    realm : RmmRealm,  
    vdev_id : bits(64),  
    inst_id : integer) => RmmRdev
```

B3.50 RdevIdsValid function

Returns TRUE if vdev_id identifies a Realm device which is assigned to realm.

```
func RdevIdsValid(  
    realm : RmmRealm,  
    vdev_id : bits(64)) => boolean
```

B3.51 RdevIdsAreValid function

Returns TRUE if the tuple (vdev_id, inst_id) identifies a Realm device which is assigned to realm.

```
func RdevIdsAreValid(  
    realm : RmmRealm,  
    vdev_id : bits(64),  
    inst_id : integer) => boolean
```

B3.52 RdevMeasurementParamsValid function

Returns TRUE if the specified device measurement parameters are valid.

```
func RdevMeasurementParamsValid(  
    params : RsiDeviceMeasurementsParams) => boolean  
begin  
    var no_ids : boolean = TRUE;  
    for i = 0 to 255 do  
        if (params.meas_ids[i] == RSI_TRUE) then  
            no_ids = FALSE;  
        end  
    end  
  
    if (no_ids) then  
        return FALSE;  
    end  
  
    for i = 0 to 7 do  
        if (params.meas_ids[i] == RSI_FALSE  
            && params.meas_params[i] == RSI_TRUE) then  
            return FALSE;  
        end  
    end  
  
    return TRUE;  
end
```

B3.53 ReadMemory function

Read contents of memory at address range [addr + offset, addr + offset + size)

offset and size are both numbers of bytes.

```
func ReadMemory(  
    addr : bits(64),  
    offset : integer,  
    size : integer) => bits(size * 8)
```

B3.54 RealmAt function

Returns the Realm whose RD is located at physical address addr.

```
func RealmAt(  
    addr : Address) => RmmRealm
```

See also:

- [A2.1 Realm](#)

B3.55 RealmIsLive function

Returns TRUE if the Realm whose RD is located at physical address addr is live.

```
func RealmIsLive(  
    addr : Address) => boolean
```

See also:

- [A2.1.4 Realm liveness](#)

B3.56 RealmParamsSupported function

Returns TRUE if the Realm parameters are supported by the implementation.

```
func RealmParamsSupported(  
    params : RmiRealmParams) => boolean  
begin  
    var impl : RmmFeatures = ImplFeatures();  
  
    if (params.flags0.lpa2 == RMI_FEATURE_TRUE  
        && impl.feats_lpa2 != FEATURE_TRUE) then  
        return FALSE;  
    end  
  
    if (params.flags0.sve == RMI_FEATURE_TRUE  
        && impl.feats_sve != FEATURE_TRUE) then  
        return FALSE;  
    end  
  
    if (params.flags0.pmu == RMI_FEATURE_TRUE  
        && impl.feats_pmu != FEATURE_TRUE) then  
        return FALSE;  
    end  
  
    if (params.flags0.da == RMI_FEATURE_TRUE  
        && impl.feats_da != FEATURE_TRUE) then
```

```
        return FALSE;
    end

    if (params.s2sz > impl.max_ipa_width) then
        return FALSE;
    end

    if (params.sve_vl > impl.max_sve_vl) then
        return FALSE;
    end

    if (params.num_bps == 0
        || params.num_bps + 1 > impl.num_bps) then
        return FALSE;
    end

    if (params.num_wps == 0
        || params.num_wps + 1 > impl.num_wps) then
        return FALSE;
    end

    if (params.pmu_num_ctrs > impl.pmu_num_ctrs) then
        return FALSE;
    end

    if (params.hash_algo == RMI_HASH_SHA_256
        && impl.feats_sha_256 != FEATURE_TRUE) then
        return FALSE;
    end

    if (params.hash_algo == RMI_HASH_SHA_512
        && impl.feats_sha_512 != FEATURE_TRUE) then
        return FALSE;
    end

    if (params.num_aux_planes > impl.max_num_aux_planes) then
        return FALSE;
    end

    if (params.num_aux_planes > 0) then
        if (params.flags1.rtt_tree_pp == RMI_FEATURE_FALSE
            && impl.plane_rtt == PLANE_RTT_AUX) then
            return FALSE;
        end

        if (params.flags1.rtt_tree_pp == RMI_FEATURE_TRUE
            && impl.plane_rtt == PLANE_RTT_SINGLE) then
            return FALSE;
        end
    end

    return TRUE;
end
```

See also:

- [A2.1.6 Realm parameters](#)
- [Chapter A3 Feature discovery and configuration](#)

B3.57 RealmRttBaseEqual function

Returns TRUE if RTT base values of `realm` match the provided values.

```
func RealmRttBaseEqual(  
    realm : RmmRealm,  
    rtt_base : Address,  
    aux_rtt_base : array[3] of Address) => boolean  
begin  
    if (realm.rtt_base[0] != rtt_base) then  
        return FALSE;  
    end  
  
    for i = 0 to 2 do  
        if (realm.rtt_base[i + 1] != aux_rtt_base[i]) then  
            return FALSE;  
        end  
    end  
  
    return TRUE;  
end
```

B3.58 RealmVmidEqual function

Returns TRUE if RTT base values of `realm` match the provided values.

```
func RealmVmidEqual(  
    realm : RmmRealm,  
    vmid : bits(16),  
    aux_vmid : array[3] of bits(16)) => boolean  
begin  
    if (realm.vmid[0] != vmid) then  
        return FALSE;  
    end  
  
    for i = 0 to 2 do  
        if (realm.vmid[i + 1] != aux_vmid[i]) then  
            return FALSE;  
        end  
    end  
  
    return TRUE;  
end
```

B3.59 RecAt function

Returns the REC object located at physical address `addr`.

```
func RecAt(  
    addr : Address) => RmmRec
```

See also:

- [A2.3 Realm Execution Context](#)

B3.60 RecAuxCount function

Returns the number of auxiliary Granules required for a REC in the Realm described by `rd`.

The return value is guaranteed not to be greater than 16.

For a given Realm, this function always returns the same value.

```
func RecAuxCount(  
    rd : Address) => integer
```

B3.61 RecFromMpidr function

Returns the REC object identified by the specified MPIDR value, in the current Realm.

```
func RecFromMpidr(  
    mpidr : bits(64)) => RmmRec
```

B3.62 RecIndex function

Returns the REC index which corresponds to mpidr.

```
func RecIndex(  
    mpidr : RmiRecMpidr) => integer  
begin  
    return (UInt(mpidr.aff0)  
        + 16 * UInt(mpidr.aff1)  
        + 16 * 256 * UInt(mpidr.aff2)  
        + 16 * 256 * 256 * UInt(mpidr.aff3));  
end
```

See also:

- [A2.3.3 REC index and MPIDR value](#)

B3.63 RecRipasResponseToRsi function

Returns response to RIPAS change request.

```
func RecRipasResponseToRsi(  
    rec : RmmRec) => RsiResponse  
begin  
    if ((rec.ripas_value == RAM)  
        && (rec.ripas_addr != rec.ripas_top)  
        && (rec.ripas_response == REJECT)) then  
        return RSI_REJECT;  
    end  
  
    return RSI_ACCEPT;  
end
```

See also:

- [A5.4 RIPAS change](#)

B3.64 RecS2APResponseToRsi function

Returns response to S2AP change request.

```
func RecS2APResponseToRsi(  
    rec : RmmRec) => RsiResponse  
begin  
    if ((rec.s2ap_addr != rec.s2ap_top)
```

```
        && (rec.s2ap_response == REJECT)) then
    return RSI_REJECT;
end

    return RSI_ACCEPT;
end
```

See also:

- [A10.3.2.4 Stage 2 access permissions change](#)

B3.65 RemExtend function

Extend REM, using size LSBs from new_value, with the remaining bits zero-padded to form a 512-bit value.

```
func RemExtend(
    hash_algo : RmmHashAlgorithm,
    old_value : RmmRealmMeasurement,
    new_value : RmmRealmMeasurement,
    size : integer) => RmmRealmMeasurement
```

See also:

- [A7.1.2 Realm Extensible Measurement](#)

B3.66 ResultEqual function

Returns TRUE if command result matches the stated value.

```
func ResultEqual(
    result : RmiCommandReturnCode,
    status : RmiStatusCode) => boolean
```

```
func ResultEqual(
    result : RmiCommandReturnCode,
    status : RmiStatusCode,
    index : integer) => boolean
```

B3.67 RimExtendData function

Extend RIM with contribution from DATA creation.

```
func RimExtendData(
    realm : RmmRealm,
    ipa : Address,
    data : Address,
    flags : RmiDataFlags) => RmmRealmMeasurement
```

See also:

- [B4.3.1.4 RMI_DATA_CREATE extension of RIM](#)

B3.68 RimExtendRec function

Extend RIM with contribution from REC creation.

```
func RimExtendRec(
    realm : RmmRealm,
    params : RmiRecParams) => RmmRealmMeasurement
```

See also:

- [B4.3.28.4 RMI_REC_CREATE extension of RIM](#)

B3.69 RimExtendRipas function

Extend RIM with contribution from RIPAS change for an IPA range.

```
func RimExtendRipas(  
    realm : RmmRealm,  
    base : Address,  
    top : Address,  
    level : integer) => RmmRealmMeasurement  
begin  
    var rim = realm.measurements[0];  
    var size = RttLevelSize(level);  
    var addr = base;  
  
    while (UInt(addr) < UInt(top)) do  
        rim = RimExtendRipasForEntry(rim, addr, level);  
        addr = ToAddress(UInt(addr) + size);  
    end  
  
    return rim;  
end
```

See also:

- [B4.3.41.4 RMI_RTT_INIT_RIPAS extension of RIM](#)

B3.70 RimExtendRipasForEntry function

Extend RIM with contribution from RIPAS change for a single RTT entry.

```
func RimExtendRipasForEntry(  
    rim : RmmRealmMeasurement,  
    ipa : Address,  
    level : integer) => RmmRealmMeasurement
```

B3.71 RimInit function

Initialize RIM.

```
func RimInit(  
    hash_algo : RmmHashAlgorithm,  
    params : RmiRealmParams) => RmmRealmMeasurement
```

See also:

- [B4.3.25.4 RMI_REALM_CREATE initialization of RIM](#)

B3.72 RipasToRmi function

Encodes a RIPAS value.

```
func RipasToRmi(  
    ripas : RmmRipas) => RmiRipas  
begin  
    case ripas of
```

```
        when EMPTY      => return RMI_EMPTY;
        when RAM         => return RMI_RAM;
        when DESTROYED  => return RMI_DESTROYED;
        when DEV         => return RMI_DEV;
    end
end
```

B3.73 RmiAddressRangesEqual16 function

Returns TRUE if the first count entries in two arrays of address ranges are equal.

```
func RmiAddressRangesEqual16(
    ranges1 : array [16] of RmmAddressRange,
    ranges2 : array [16] of RmiAddressRange,
    count : integer) => boolean
begin
    assert 0 <= count && count <= 16;
    for i = 0 to count - 1 do
        if ranges1[i].base != ranges2[i].base then
            return FALSE;
        end
        if ranges1[i].top != ranges2[i].top then
            return FALSE;
        end
    end
    return TRUE;
end
```

B3.74 RmiAddressRangesEqual4 function

Returns TRUE if the first count entries in two arrays of address ranges are equal.

```
func RmiAddressRangesEqual4(
    ranges1 : array [4] of RmmAddressRange,
    ranges2 : array [4] of RmiAddressRange,
    count : integer) => boolean
begin
    assert 0 <= count && count <= 4;
    for i = 0 to count - 1 do
        if ranges1[i].base != ranges2[i].base then
            return FALSE;
        end
        if ranges1[i].top != ranges2[i].top then
            return FALSE;
        end
    end
    return TRUE;
end
```

B3.75 RmiDevCommDataAt function

Returns device communication data structure stored at physical address addr.

If the PAS of addr is not NS, the return value is UNKNOWN.

```
func RmiDevCommDataAt (
    addr : Address) => RmiDevCommData
```

B3.76 RmiFeatureRegister0Decode function

Decode RmiFeatureRegister0 value.

```
func RmiFeatureRegister0Decode(  
    value : bits(64)) => RmiFeatureRegister0
```

B3.77 RmiFeatureRegisterEncode function

Encode feature register.

```
func RmiFeatureRegisterEncode(  
    index : integer) => bits(64)  
begin  
    var impl : RmmFeatures = ImplFeatures();  
    var result : bits(64) = Zeros();  
  
    if (index == 0) then  
        var reg : RmiFeatureRegister0;  
  
        reg.S2SZ = impl.max_ipa_width;  
        reg.LPA2 = FeatureToRmi(impl.feats_lpa2);  
        reg.SVE = FeatureToRmi(impl.feats_sve);  
        reg.SVE_VL = impl.max_sve_vl;  
  
        assert impl.num_bps >= 2 && impl.num_bps <= 2^6;  
        reg.NUM_BPS = impl.num_bps - 1;  
  
        assert impl.num_wps >= 2 && impl.num_wps <= 2^6;  
        reg.NUM_WPS = impl.num_wps - 1;  
  
        reg.PMU = FeatureToRmi(impl.feats_pmu);  
        reg.PMU_NUM_CTRS = impl.pmu_num_ctrs;  
        reg.HASH_SHA_256 = FeatureToRmi(impl.feats_sha_256);  
        reg.HASH_SHA_512 = FeatureToRmi(impl.feats_sha_512);  
        reg.DA = FeatureToRmi(impl.feats_da);  
  
        case impl.plane_rtt of  
            when PLANE_RTT_AUX => reg.PLANE_RTT = RMI_PLANE_RTT_AUX;  
            when PLANE_RTT_AUX_SINGLE => reg.PLANE_RTT = RMI_PLANE_RTT_AUX_SINGLE;  
            when PLANE_RTT_SINGLE => reg.PLANE_RTT = RMI_PLANE_RTT_SINGLE;  
        end  
  
        reg.MAX_NUM_AUX_PLANES = impl.max_num_aux_planes;  
        reg.MAX_RECS_ORDER = impl.max_recs_order;  
  
        assert impl.gicv3_num_lrs >= 1 && impl.gicv3_num_lrs <= 2^4;  
        reg.GICV3_NUM_LRS = impl.gicv3_num_lrs - 1;  
  
        // Omitted: encode reg into bits(64) value  
    end  
  
    if (index == 1) then  
        var reg : RmiFeatureRegister1;  
  
        reg.MAX_MECID = impl.max_mecid;  
  
        // Omitted: encode reg into bits(64) value  
    end  
end
```

```
        return result;  
    end
```

B3.78 RmiPdevEventIsValid function

Returns TRUE if `ev` is a valid encoding, and the event is supported by `pdev`.

```
func RmiPdevEventIsValid(  
    ev : RmiPdevEvent) => boolean  
begin  
    return (ev == RMI_IDE_KEY_REFRESH);  
end
```

B3.79 RmiPdevFlagsDecode function

Decode `RmiPdevFlags` value.

```
func RmiPdevFlagsDecode(  
    value : bits(64)) => RmiPdevFlags
```

B3.80 RmiPdevParamsAt function

Returns PDEV parameters stored at physical address `addr`.

If the PAS of `addr` is not NS, the return value is UNKNOWN.

```
func RmiPdevParamsAt(  
    addr : Address) => RmiPdevParams
```

B3.81 RmiPdevParamsIsValid function

Returns TRUE if the memory location contains a valid encoding of the `RmiPdevParams` type and all the following are true:

- The device identifier is valid
- The device identifier is not equal to the device identifier of another PDEV
- The Root Port identifier is valid
- The IDE stream identifier is valid
- The RID range is valid
- The RID range does not overlap the RID range of another PDEV
- The base and top of every address range is aligned to the size of a Granule
- Every address range falls within a memory range permitted by the system
- None of the address ranges overlaps another address range for this PDEV
- None of the address ranges overlaps an address range for another PDEV

```
func RmiPdevParamsIsValid(  
    addr : Address) => boolean
```

B3.82 RmiRealmParamsAt function

Returns Realm parameters stored at physical address `addr`.

If the PAS of `addr` is not NS, the return value is UNKNOWN.

```
func RmiRealmParamsAt(  
    addr : Address) => RmiRealmParams
```

See also:

- [A2.1.6 Realm parameters](#)

B3.83 RmiRealmParamsIsValid function

Returns TRUE if the memory location contains a valid encoding of the RmiRealmParams type.

```
func RmiRealmParamsIsValid(  
    addr : Address) => boolean
```

B3.84 RmiRecParamsAt function

Returns REC parameters stored at physical address `addr`.

If the PAS of `addr` is not NS, the return value is UNKNOWN.

```
func RmiRecParamsAt(  
    addr : Address) => RmiRecParams
```

B3.85 RmiRecRunAt function

Returns the RecRun object stored at physical address `addr`.

```
func RmiRecRunAt(  
    addr : Address) => RmiRecRun
```

See also:

- [A4.2 REC entry](#)
- [A4.3 REC exit](#)

B3.86 RmiVdevFlagsDecode function

Decode RmiVdevFlags value.

```
func RmiVdevFlagsDecode(  
    value : bits(64)) => RmiVdevFlags
```

B3.87 RmiVdevParamsAt function

Returns VDEV parameters stored at physical address `addr`.

If the PAS of `addr` is not NS, the return value is UNKNOWN.

```
func RmiVdevParamsAt(  
    addr : Address) => RmiVdevParams
```

B3.88 RmiVdevParamsIsValid function

Returns TRUE if the memory location contains a valid encoding of the RmiPdevParams type.

```
func RmiVdevParamsIsValid(  
    addr : Address) => boolean
```

B3.89 RsiDeviceInfoAt function

Returns device configuration stored at IPA `addr`, mapped in the current Realm.

```
func RsiDeviceInfoAt(  
    addr : Address) => RsiDeviceInfo
```

B3.90 RsiDeviceMeasParamsAt function

Returns RDEV measurement parameters stored at IPA addr.

```
func RsiDeviceMeasParamsAt(  
    addr : Address) => RsiDeviceMeasurementsParams
```

B3.91 RsiFeatureRegisterEncode function

Encode feature register.

```
func RsiFeatureRegisterEncode(  
    index : integer) => bits(64)  
begin  
    var impl : RmmFeatures = ImplFeatures();  
    var result : bits(64) = Zeros();  
  
    if (index == 0) then  
        var reg : RsiFeatureRegister0;  
  
        reg.DA = FeatureToRsi(impl.feats_da);  
  
        // Omitted: set reg.MRO depending on whether platform  
        // implements FEAT_S2PIE  
  
        // Omitted: encode reg into bits(64) value  
    end  
  
    return result;  
end
```

B3.92 RsiHostCallAt function

Returns Host call data stored at IPA addr, mapped in the current Realm.

```
func RsiHostCallAt(  
    addr : Address) => RsiHostCall
```

B3.93 RsiPlaneRunAt function

Returns the PlaneRun object stored at IPA addr.

```
func RsiPlaneRunAt(  
    realm : RmmRealm,  
    addr : Address) => RsiPlaneRun
```

B3.94 RsiRealmConfigAt function

Returns Realm configuration stored at IPA addr, mapped in the current Realm.

```
func RsiRealmConfigAt(  
    addr : Address) => RsiRealmConfig
```

B3.95 RttAllEntriesContiguous function

Returns TRUE if all entries in the RTT at address `rtt` at level `level` have contiguous output addresses, starting with `addr`.

```
func RttAllEntriesContiguous(  
    rtt : RmmRtt,  
    addr : Address,  
    level : integer) => boolean
```

See also:

- [A5.5 Realm Translation Table](#)

B3.96 RttAllEntriesRipas function

Returns TRUE if all entries in the RTT at address `rtt` have RIPAS `ripas`.

```
func RttAllEntriesRipas(  
    rtt : RmmRtt,  
    ripas : RmmRipas) => boolean
```

B3.97 RttAllEntriesState function

Returns TRUE if all entries in the RTT at address `rtt` have state `state`.

```
func RttAllEntriesState(  
    rtt : RmmRtt,  
    state : RmmRttEntryState) => boolean
```

See also:

- [A5.5 Realm Translation Table](#)

B3.98 RttAt function

Returns the RTT at address `rtt`.

```
func RttAt(  
    addr : Address) => RmmRtt
```

B3.99 RttConfigsValid function

Returns TRUE if the RTT configuration values provided are self-consistent and are supported by the platform.

```
func RttConfigIsValid(  
    ipa_width : integer,  
    rtt_level_start : integer,  
    rtt_num_start : integer) => boolean
```

See also:

- [A5.5 Realm Translation Table](#)

B3.100 RttDescriptorDecode function

Decode an RTT descriptor.

```
func RttDescriptorDecode(  
    desc : bits(64)) => RmmRttEntry
```

B3.101 RttDescriptorIsValidForUnprotected function

Returns TRUE if, within the descriptor `desc`, all of the following are true:

- All fields which are *Host-controlled Unprotected RTT attributes* are set to architecturally valid values.
- All fields which are not *Host-controlled Unprotected RTT attributes* are set to zero.

```
func RttDescriptorIsValidForUnprotected(  
    desc : bits(64)) => boolean
```

See also:

- [A5.5.11.3 RTT entry attributes for ASSIGNED_NS mappings](#)

B3.102 RttEntriesInRangeRipas function

Returns TRUE if all entries in the RTT at address `rtt` at level `level`, within IPA range [`base`, `top`), have RIPAS `ripas`.

```
func RttEntriesInRangeRipas(  
    rtt : RmmRtt,  
    level : integer,  
    base : Address,  
    top : Address,  
    ripas : RmmRipas) => boolean
```

B3.103 RttEntryAt function

Returns the `i`th entry in the RTT at address `rtt`.

```
func RttEntryAt(  
    rtt : Address,  
    i : integer) => RmmRttEntry
```

See also:

- [A5.5 Realm Translation Table](#)

B3.104 RttEntryIndex function

Returns the index of the entry in a level `level` RTT which is identified by `addr`.

```
func RttEntryIndex(  
    addr : Address,  
    level : integer) => integer
```

See also:

- [A5.5 Realm Translation Table](#)

B3.105 RttEntryStateToRmi function

Encodes the state of an RTTE.

```
func RttEntryStateToRmi(  
    state : RmmRttEntryState) => RmiRttEntryState
```

```
begin
  case state of
    when UNASSIGNED      => return RMI_UNASSIGNED;
    when ASSIGNED        => return RMI_ASSIGNED;
    when UNASSIGNED_NS   => return RMI_UNASSIGNED;
    when ASSIGNED_NS     => return RMI_ASSIGNED;
    when TABLE          => return RMI_TABLE;
    when ASSIGNED_DEV_PRIVATE => return RMI_ASSIGNED_DEV_PRIVATE;
    when ASSIGNED_DEV_SHARED => return RMI_ASSIGNED_DEV_SHARED;
    when AUX_DESTROYED   => return RMI_AUX_DESTROYED;
  end
end
```

B3.106 RttFold function

Returns the RTTE which results from folding the homogeneous RTT at address `rtt`.

```
func RttFold(
  rtt : RmmRtt) => RmmRttEntry
```

See also:

- [A5.5.6 RTT folding](#)

B3.107 RttIsHomogeneous function

Returns TRUE if the RTT at address `rtt` is homogeneous.

```
func RttIsHomogeneous(
  rtt : RmmRtt) => boolean
```

See also:

- [A5.5.6 RTT folding](#)

B3.108 RttIsLive function

Returns TRUE if the RTT at address `rtt` is live.

```
func RttIsLive(
  rtt : RmmRtt) => boolean
```

See also:

- [A5.5.8 RTTE liveness and RTT liveness](#)
- [A5.5.9 RTT destruction](#)

B3.109 RttLevelIsBlockOrPage function

Returns TRUE if `level` is either a block or page RTT level for the Realm described by `rd`.

```
func RttLevelIsBlockOrPage(
  rd : Address,
  level : integer) => boolean
```

See also:

- [A5.5 Realm Translation Table](#)

B3.110 RttLevelsStarting function

Returns TRUE if `level` is the starting level of the RTT for the Realm described by `rd`.

```
func RttLevelsStarting(  
    rd : Address,  
    level : integer) => boolean
```

See also:

- [A5.5 Realm Translation Table](#)

B3.111 RttLevelsValid function

Returns TRUE if `level` is a valid RTT level for the Realm described by `rd`.

```
func RttLevelsValid(  
    rd : Address,  
    level : integer) => boolean
```

See also:

- [A5.5 Realm Translation Table](#)

B3.112 RttLevelSize function

Returns the size of the address space described by each entry in an RTT at `level`.

If `level` is invalid, the return value is UNKNOWN.

```
func RttLevelSize(  
    level : integer) => integer
```

See also:

- [A5.5 Realm Translation Table](#)

B3.113 RttsAllProtectedEntriesRipas function

Returns TRUE if the RIPAS of all entries identified by Protected IPAs in all of the starting-level RTT Granules is equal to `ripas`.

```
func RttsAllProtectedEntriesRipas(  
    rtt_base : Address,  
    rtt_num_start : integer,  
    ripas : RmmRipas) => boolean
```

B3.114 RttsAllProtectedEntriesState function

Returns TRUE if the state of all entries identified by Protected IPAs in all of the starting-level RTT Granules is equal to `state`.

```
func RttsAllProtectedEntriesState(  
    rtt_base : Address,  
    rtt_num_start : integer,  
    state : RmmRttEntryState) => boolean
```

B3.115 *RttsAllUnprotectedEntriesState* function

Returns TRUE if the state of all entries identified by Unprotected IPAs in all of the starting-level RTT Granules is equal to *state*.

```
func RttsAllUnprotectedEntriesState(  
    rtt_base : Address,  
    rtt_num_start : integer,  
    state : RmmRttEntryState) => boolean
```

B3.116 *RttsGranuleState* function

Inductive function which identifies the states of the starting-level RTT Granules.

This function is used in the definition of command footprint.

```
func RttsGranuleState(  
    rtt_base : Address,  
    rtt_num_start : integer)
```

B3.117 *RttSkipEntriesUnlessRipas* function

Scanning *rtt* starting from *ipa*, returns the IPA of the first entry whose RIPAS is *ripas*.

If no entry is found whose RIPAS is *ripas*, returns the next IPA after the last entry in *rtt*.

The return value is aligned to the size of the address range described by an entry at RTT level.

```
func RttSkipEntriesUnlessRipas(  
    rtt : RmmRtt,  
    level : integer,  
    ipa : Address,  
    ripas : RmmRipas) => Address
```

B3.118 *RttSkipEntriesUnlessState* function

Scanning *rtt* starting from *ipa*, returns the IPA of the first entry whose state is *state*.

If no entry is found whose state is *state*, returns the next IPA after the last entry in *rtt*.

The return value is aligned to the size of the address range described by an entry at RTT level.

```
func RttSkipEntriesUnlessState(  
    rtt : RmmRtt,  
    level : integer,  
    ipa : Address,  
    state : RmmRttEntryState) => Address
```

B3.119 *RttSkipEntriesWithRipas* function

Scan *rtt* starting from *base* and terminating at *top*.

- If *stop_at_destroyed* is FALSE then return IPA of the first entry whose state is TABLE.
- If *stop_at_destroyed* is TRUE then return IPA of the first entry whose state is TABLE or whose RIPAS is DESTROYED.

If no such entry is found, returns the smaller of:

- The next IPA after the last entry in *rtt*
- The *top* argument.

The return value is aligned to the size of the address range described by an entry at RTT level.

```
func RttSkipEntriesWithRipas(  
    rtt : RmmRtt,  
    level : integer,  
    base : Address,  
    top : Address,  
    stop_at_destroyed : boolean) => Address  
begin  
    var result : Address = RttSkipEntriesUnlessState(  
        rtt, level, base, TABLE);  
  
    if stop_at_destroyed then  
        result = MinAddress(result,  
            RttSkipEntriesUnlessRipas(  
                rtt, level, base, DESTROYED));  
    end  
  
    result = MinAddress(result, top);  
  
    return AlignDownToRttLevel(result, level);  
end
```

B3.120 RttSkipNonLiveEntries function

Scanning rtt starting from ipa, returns the IPA of the first live entry.

If no live entry is found, returns the next IPA after the last entry in rtt.

The return value is aligned to the size of the address range described by an entry at RTT level.

```
func RttSkipNonLiveEntries(  
    rtt : RmmRtt,  
    level : integer,  
    ipa : Address) => Address  
begin  
    var result : Address = RttSkipEntriesUnlessState(  
        rtt, level, ipa, ASSIGNED);  
  
    result = MinAddress(result,  
        RttSkipEntriesUnlessState(  
            rtt, level, ipa, ASSIGNED_NS));  
  
    result = MinAddress(result,  
        RttSkipEntriesUnlessState(  
            rtt, level, ipa, TABLE));  
  
    result = MinAddress(result,  
        RttSkipEntriesUnlessState(  
            rtt, level, ipa, ASSIGNED_DEV_PRIVATE));  
  
    result = MinAddress(result,  
        RttSkipEntriesUnlessState(  
            rtt, level, ipa, ASSIGNED_DEV_SHARED));  
  
    return AlignDownToRttLevel(result, level);  
end
```

See also:

- [A5.5.8 RTTE liveness and RTT liveness](#)

B3.121 *RttsStateEqual* function

Returns TRUE if the state of all of the starting-level RTT Granules is equal to *state*.

```
func RttsStateEqual(  
    rtt_base : Address,  
    rtt_num_start : integer,  
    state : RmmGranuleState) => boolean  
begin  
    for i = 0 to rtt_num_start - 1 do  
        var addr = (UInt(rtt_base) + i * RMM_GRANULE_SIZE) [(ADDRESS_WIDTH-1):0];  
        if (!PaIsDelegable(addr)  
            || GranuleAt(addr).state != state) then  
            return FALSE;  
        end  
    end  
    return TRUE;  
end
```

B3.122 *RttWalk* function

Returns the result of an RTT walk from the base of RTT tree index owned by *rd*, to address *addr*.

The walk does not progress beyond *level*.

```
func RttWalk(  
    rd : Address,  
    addr : Address,  
    level : integer,  
    index : integer) => RmmRttWalkResult
```

See also:

- [A5.5.10 RTT walk](#)

B3.123 *RttWalkAnyNotAligned* function

Performs one or more RTT walks within the IPA range [*base*, *top*), on one or more RTT trees owned by *rd*.

For a Realm which is configured to use an RTT per Plane, it is permitted for an implementation to either walk just one of the RTTs, or to walk all of them.

It is permitted for an implementation to perform multiple walks, at successive IPAs, on the same RTT.

If one of the walks performed terminates earlier than *level* then the return value indicates the RTT index and the IPA at which the walk was performed. In this case, the result's "valid" value is TRUE.

If none of the walks performed terminates earlier than *level* then the result's "valid" value is FALSE.

```
func RttWalkAnyNotAligned(  
    rd : Address,  
    base : Address,  
    top : Address,  
    level : integer) => RmmRttWalkNotAligned
```

B3.124 TdildlsFree function

Returns TRUE if `tdi_id` is unused within the segment identified by `segment_id`.

```
func TdiIdIsFree(  
    tdi_id : bits(64),  
    segment_id : bits(16)) => boolean
```

B3.125 ToAddress function

Convert integer to Address.

```
func ToAddress(value : integer) => Address  
begin  
    return value[(ADDRESS_WIDTH-1):0];  
end
```

B3.126 ToBits64 function

Convert integer to Bits64.

```
func ToBits64(value : integer) => bits(64)  
begin  
    return value[63:0];  
end
```

B3.127 VdevAt function

Returns the VDEV object located at physical address `addr`.

```
func VdevAt(  
    addr : Address) => RmmVdev
```

B3.128 VdevAuxCount function

Returns the number of auxiliary Granules required for a VDEV with the specified flags.

The return value is guaranteed not to be greater than 32.

For a given flags value, this function always returns the same value.

```
func VdevAuxCount(  
    pdev_flags : RmiPdevFlags,  
    vdev_flags : RmiVdevFlags) => integer
```

B3.129 VmidsAreFree function

Returns TRUE if `vmid` is unused.

```
func VmidsAreFree(  
    vmid : array[4] of bits(16)) => boolean
```

```
func VmidsAreFree(  
    vmid : bits(16),  
    aux_vmid : array[3] of bits(16)) => boolean
```

B3.130 *VmidsAreValid* function

Returns TRUE if `vmid` is valid on the platform.

```
func VmidsAreValid(  
    vmid : bits(16),  
    aux_vmid : array [3] of bits(16)) => boolean
```

If the underlying hardware platform does not implement FEAT_VMID16 then a VMID value with `vmid[15:8] != 0` is invalid.

See also:

- [A2.1.3 Realm attributes](#)
- [B4.3.25 RMI_REALM_CREATE command](#)

DRAFT

Chapter B4

Realm Management Interface

This chapter defines the interface used by the Host to manage Realms.

B4.1 RMI version

R_{NCFDX} This specification defines version 1.1 of the Realm Management Interface.

See also:

- [Chapter B2 Interface versioning](#)
- [B4.3.55 RMI_VERSION command](#)

B4.2 RMI command return codes

I_{JQMBN} The return code of an RMI command is a tuple which contains *status* and *index* fields.

I_{YCHQV} The *status* field of an RMI command return code indicates whether the command

- succeeded, or
- failed, and the reason for the failure.

I_{PPNST} If an RMI command succeeds then the status of its return code is RMI_SUCCESS.

I_{MBVPG} The *index* field of an RMI command return code can provide additional information about the reason for a command failure. The meaning of the index field depends on the status, and is described by the following table.

Status	Description	Meaning of index
RMI_SUCCESS	Command completed successfully	None: index is zero.
RMI_ERROR_INPUT	The value of a command input value caused the command to fail	None: index is zero.
RMI_ERROR_REALM	An attribute of a Realm does not match the expected value	Varies between usages. See individual commands for details.
RMI_ERROR_REC	An attribute of a REC does not match the expected value	None: index is zero.
RMI_ERROR_RTT	An RTT walk terminated before reaching the target RTT level, or reached an RTTE with an unexpected value	RTT level at which the walk terminated.
RMI_ERROR_RTT_AUX	RTTE in an auxiliary RTT contained an unexpected value	In some cases, indicates auxiliary RTT level at which the walk terminated. See individual commands for details.

I_{QQQNB} Multiple failure conditions in an RMI command may return the same error code - that is, the same status and index values.

R_{XRDYQ} If an input to an RMI command uses an invalid encoding then the command fails and returns RMI_ERROR_INPUT. Command inputs include registers and in-memory data structures.

Invalid encodings include:

- using a reserved encoding in an enumeration

See also:

- [B4.4.3 RmiCommandReturnCode type](#)

B4.3 RMI commands

The following table summarizes the FIDs of commands in the RMI interface.

FID	Command
0xC4000150	RMI_VERSION
0xC4000151	RMI_GRANULE_DELEGATE
0xC4000152	RMI_GRANULE_UNDELEGATE
0xC4000153	RMI_DATA_CREATE
0xC4000154	RMI_DATA_CREATE_UNKNOWN
0xC4000155	RMI_DATA_DESTROY
0xC4000156	RMI_PDEV_AUX_COUNT
0xC4000157	RMI_REALM_ACTIVATE
0xC4000158	RMI_REALM_CREATE
0xC4000159	RMI_REALM_DESTROY
0xC400015A	RMI_REC_CREATE
0xC400015B	RMI_REC_DESTROY
0xC400015C	RMI_REC_ENTER
0xC400015D	RMI_RTT_CREATE
0xC400015E	RMI_RTT_DESTROY
0xC400015F	RMI_RTT_MAP_UNPROTECTED
0xC4000160	RMI_VDEV_AUX_COUNT
0xC4000161	RMI_RTT_READ_ENTRY
0xC4000162	RMI_RTT_UNMAP_UNPROTECTED
...	
0xC4000164	RMI_PSCI_COMPLETE
0xC4000165	RMI_FEATURES
0xC4000166	RMI_RTT_FOLD
0xC4000167	RMI_REC_AUX_COUNT
0xC4000168	RMI_RTT_INIT_RIPAS
0xC4000169	RMI_RTT_SET_RIPAS
...	
0xC4000170	RMI_GRANULE_DEV_DELEGATE
0xC4000171	RMI_GRANULE_DEV_UNDELEGATE
0xC4000172	RMI_DEV_MEM_MAP
0xC4000173	RMI_DEV_MEM_UNMAP
0xC4000174	RMI_PDEV_ABORT

FID	Command
0xC4000175	RMI_PDEV_COMMUNICATE
0xC4000176	RMI_PDEV_CREATE
0xC4000177	RMI_PDEV_DESTROY
0xC4000178	RMI_PDEV_GET_STATE
0xC4000179	RMI_PDEV_IDE_RESET
0xC400017A	RMI_PDEV_NOTIFY
0xC400017B	RMI_PDEV_SET_PUBKEY
0xC400017C	RMI_PDEV_STOP
0xC400017D	RMI_RTT_AUX_CREATE
0xC400017E	RMI_RTT_AUX_DESTROY
0xC400017F	RMI_RTT_AUX_FOLD
0xC4000180	RMI_RTT_AUX_MAP_PROTECTED
0xC4000181	RMI_RTT_AUX_MAP_UNPROTECTED
...	
0xC4000183	RMI_RTT_AUX_UNMAP_PROTECTED
0xC4000184	RMI_RTT_AUX_UNMAP_UNPROTECTED
0xC4000185	RMI_VDEV_ABORT
0xC4000186	RMI_VDEV_COMMUNICATE
0xC4000187	RMI_VDEV_CREATE
0xC4000188	RMI_VDEV_DESTROY
0xC4000189	RMI_VDEV_GET_STATE
0xC400018A	RMI_VDEV_STOP
0xC400018B	RMI_RTT_SET_S2AP
0xC400018C	RMI_MEC_SET_SHARED
0xC400018D	RMI_MEC_SET_PRIVATE
0xC400018E	RMI_VDEV_COMPLETE

B4.3.1 RMI_DATA_CREATE command

Creates a Data Granule, copying contents from a Non-secure Granule provided by the caller.

See also:

- [Chapter A5 Realm memory management](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [D1.2.3 Initialize memory of New Realm flow](#)

B4.3.1.1 Interface

B4.3.1.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000153
rd	X1	63:0	Address	PA of the RD for the target Realm
data	X2	63:0	Address	PA of the target Data
ipa	X3	63:0	Address	IPA at which the Granule will be mapped in the target Realm
src	X4	63:0	Address	PA of the source Granule
flags	X5	63:0	RmiDataFlags	Flags

B4.3.1.1.2 Context

The RMI_DATA_CREATE command operates on the following context.

Name	Type	Value	Before	Description
realm_pre	RmmRealm	RealmAt (rd)	true	Realm
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL , RMM_RTT_TREE_PRIMARY)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index

B4.3.1.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.1.2 Failure conditions

ID	Condition
src_align	pre: !AddrIsGranuleAligned(src) post: ResultEqual(result, RMI_ERROR_INPUT)
src_bound	pre: !PaIsDelegable(src) post: ResultEqual(result, RMI_ERROR_INPUT)
src_pas	pre: !GranuleAccessPermitted(src, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
data_align	pre: !AddrIsGranuleAligned(data) post: ResultEqual(result, RMI_ERROR_INPUT)
data_bound	pre: !PaIsDelegable(data) post: ResultEqual(result, RMI_ERROR_INPUT)
data_state	pre: GranuleAt(data).state != DELEGATED post: ResultEqual(result, RMI_ERROR_INPUT)
data_bound2	pre: ((realm.feat_lpa2 == FEATURE_FALSE) && (UInt(data) >= 2^48)) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsGranuleAligned(ipa) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: !AddrIsProtected(ipa, realm_pre) post: ResultEqual(result, RMI_ERROR_INPUT)
realm_state	pre: realm_pre.state != REALM_NEW post: ResultEqual(result, RMI_ERROR_REALM)
rtt_walk	pre: walk.level < RMM_RTT_PAGE_LEVEL post: ResultEqual(result, RMI_ERROR_RTT, walk.level)
rtte_state	pre: walk.rtte.state != UNASSIGNED post: ResultEqual(result, RMI_ERROR_RTT, walk.level)

B4.3.1.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [realm_state]
[rd_bound, rd_state] < [rtt_walk, rtte_state]
[ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.1.3 Success conditions

ID	Condition
data_state	<code>GranuleAt(data).state == DATA</code>
rtte_state	<code>walk.rtte.state == ASSIGNED</code>
rtte_ripas	<code>walk.rtte.ripas == RAM</code>
rtte_addr	<code>walk.rtte.addr == data</code>
rim	<code>realm.measurements[0] == RimExtendData(realm_pre, ipa, data, flags)</code>

B4.3.1.4 RMI_DATA_CREATE extension of RIM

On successful execution of RMI_DATA_CREATE, the new RIM value of the target Realm is calculated by the RMM as follows:

1. If `flags.measure == RMI_MEASURE_CONTENT` then using the RHA of the target Realm, compute the hash of the contents of the DATA Granule.
2. Allocate an `RmmMeasurementDescriptorData` data structure.
3. Populate the measurement descriptor:
 - Set the `desc_type` field to the descriptor type.
 - Set the `len` field to the descriptor length.
 - Set the `rim` field to the current RIM value of the target Realm.
 - Set the `ipa` field to the IPA at which the DATA Granule is mapped in the target Realm.
 - Set the `flags` field to the flags provided by the Host.
 - If `flags.measure == RMI_MEASURE_CONTENT` then set the `content` field to the hash of the contents of the DATA Granule. Otherwise, set the `content` field to zero.
4. Using the RHA of the target Realm, compute the hash of the measurement descriptor. Set the RIM of the target Realm to this value, zero filling upper bytes if the RHA output is smaller than the size of the RIM.

See also:

- [A7.1.1 Realm Initial Measurement](#)
- [B3.67 RimExtendData function](#)
- [C2.15 RmmMeasurementDescriptorData type](#)

B4.3.1.5 Footprint

ID	Value
data_state	<code>GranuleAt(data).state</code>
rim	<code>realm.measurements[0]</code>
rtte	<code>RttEntryAt(walk.rtt_addr, entry_idx)</code>

B4.3.2 RMI_DATA_CREATE_UNKNOWN command

Creates a Data Granule with unknown contents.

See also:

- [A2.2.4 Granule wiping](#)
- [Chapter A5 Realm memory management](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [D1.5.1 Add memory to Active Realm flow](#)

B4.3.2.1 Interface

B4.3.2.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xc4000154
rd	X1	63:0	Address	PA of the RD for the target Realm
data	X2	63:0	Address	PA of the target Data
ipa	X3	63:0	Address	IPA at which the Granule will be mapped in the target Realm

B4.3.2.1.2 Context

The RMI_DATA_CREATE_UNKNOWN command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL , RMM_RTT_TREE_PRIMARY)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index

B4.3.2.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.2.2 Failure conditions

ID	Condition
data_align	pre: !AddrIsGranuleAligned (data) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
data_bound	pre: !PaIsDelegable(data) post: ResultEqual(result, RMI_ERROR_INPUT)
data_state	pre: GranuleAt(data).state != DELEGATED post: ResultEqual(result, RMI_ERROR_INPUT)
data_bound2	pre: ((realm.feat_lpa2 == FEATURE_FALSE) && (UInt(data) >= 2^48)) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsGranuleAligned(ipa) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: !AddrIsProtected(ipa, realm) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < RMM_RTT_PAGE_LEVEL post: ResultEqual(result, RMI_ERROR_RTT, walk.level)
rtte_state	pre: walk.rtte.state != UNASSIGNED post: ResultEqual(result, RMI_ERROR_RTT, walk.level)

B4.3.2.2.1 Failure condition ordering

[rd_bound, rd_state] < [rtt_walk, rtte_state]
[ipa_bound] < [rtt_walk, rtte_state]



B4.3.2.3 Success conditions

ID	Condition
data_state	GranuleAt(data).state == DATA
data_content	Contents of target Granule are wiped.
rtte_state	walk.rtte.state == ASSIGNED
rtte_addr	walk.rtte.addr == data

B4.3.2.4 Footprint

ID	Value
data_state	<code>GranuleAt(data).state</code>
rtte	<code>RttEntryAt(walk.rtt_addr, entry_idx)</code>

DRAFT

B4.3.3 RMI_DATA_DESTROY command

Destroys a Data Granule.

See also:

- [Chapter A5 Realm memory management](#)
- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [D1.2.5 Realm destruction flow](#)

B4.3.3.1 Interface

B4.3.3.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000155
rd	X1	63:0	Address	PA of the RD which owns the target Data
ipa	X2	63:0	Address	IPA at which the Granule is mapped in the target Realm

B4.3.3.1.2 Context

The RMI_DATA_DESTROY command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL , RMM_RTT_TREE_PRIMARY)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
walk_top	Address	RttSkipNonLiveEntries (RttAt (walk.rtt_addr), walk.level, ipa)	false	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

B4.3.3.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
data	X1	63:0	Address	PA of the Data Granule which was destroyed
top	X2	63:0	Address	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

The data output value is valid only when the command result is RMI_SUCCESS.

The values of the `result` and `top` output values for different command outcomes are summarized in the following table.

Scenario	result	top	walk.rtte.state
ipa is mapped as a page	RMI_SUCCESS	> ipa	Before execution: ASSIGNED After execution: UNASSIGNED and RIPAS is DESTROYED
ipa is not mapped	(RMI_ERROR_RTT, <= 3)	> ipa	UNASSIGNED
ipa is mapped as a block	(RMI_ERROR_RTT, 0 0 < level < 3)	== ipa	ASSIGNED
RTT walk was not performed, due to any other command failure	Another error code	0	Unknown

See also:

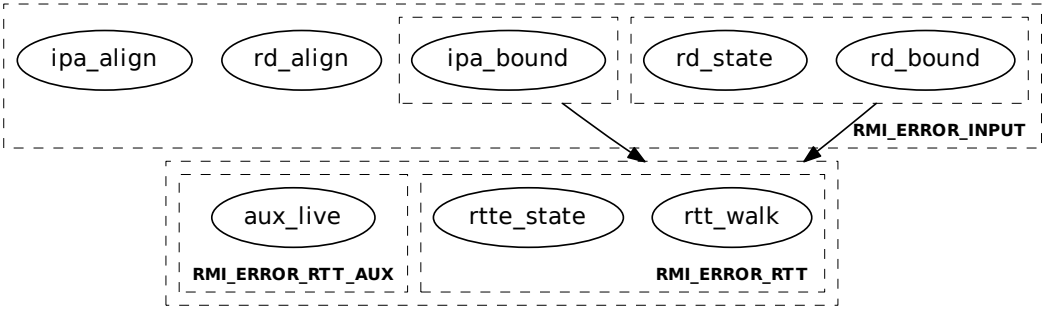
- [A5.5.8 RTTE liveness and RTT liveness](#)

B4.3.3.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsGranuleAligned(ipa) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: !AddrIsProtected(ipa, realm) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < RMM_RTT_PAGE_LEVEL post: (ResultEqual(result, RMI_ERROR_RTT, walk.level) && (top == walk_top))
rtte_state	pre: walk.rtte.state != ASSIGNED post: (ResultEqual(result, RMI_ERROR_RTT, walk.level) && (top == walk_top))
aux_live	pre: AddrIsAuxLive(ipa, realm) post: ResultEqual(result, RMI_ERROR_RTT_AUX, 0)

B4.3.3.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state, aux_live]
[ipa_bound] < [rtt_walk, rtte_state, aux_live]
```



B4.3.3.3 Success conditions

ID	Condition
data_state	<code>GranuleAt(walk.rtte.addr).state == DELEGATED</code>
rtte_state	<code>walk.rtte.state == UNASSIGNED</code>
ripas_ram	pre: <code>walk.rtte.ripas == RAM</code> post: <code>walk.rtte.ripas == DESTROYED</code>
data	<code>data == walk.rtte.addr</code>
top	<code>top == walk_top</code>

B4.3.3.4 Footprint

ID	Value
data_state	<code>GranuleAt(walk.rtte.addr).state</code>
rtte	<code>RttEntryAt(walk.rtt_addr, entry_idx)</code>

B4.3.4 RMI_DEV_MEM_MAP command

Maps device memory.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.4.1 Interface

B4.3.4.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000172
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	IPA at which the Granule will be mapped in the target Realm
addr	X3	63:0	Address	PA of the target device memory

B4.3.4.1.2 Context

The RMI_DEV_MEM_MAP command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL , RMM_RTT_TREE_PRIMARY)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
gran_state_pre	RmmGranuleState	GranuleAt (addr).state	true	Previous Granule state

B4.3.4.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.4.2 Failure conditions

ID	Condition
addr_align	pre: !AddrIsGranuleAligned (addr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
addr_bound	pre: !PaIsDelegable(addr) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsGranuleAligned(ipa) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: !AddrIsProtected(ipa, realm) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < RMM_RTT_PAGE_LEVEL post: ResultEqual(result, RMI_ERROR_RTT, walk.level)
rtte_state	pre: walk.rtte.state != UNASSIGNED post: ResultEqual(result, RMI_ERROR_RTT, walk.level)

B4.3.4.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state]
[ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.4.3 Success conditions

ID	Condition
state_private	pre: gran_state_pre == DEV_DELEGATED_PRIVATE post: GranuleAt(addr).state == DEV_PRIVATE
state_shared	pre: gran_state_pre == DEV_DELEGATED_SHARED post: GranuleAt(addr).state == DEV_SHARED
rtte_state_private	pre: gran_state_pre == DEV_DELEGATED_PRIVATE post: walk.rtte.state == ASSIGNED_DEV_PRIVATE
rtte_state_shared	pre: gran_state_pre == DEV_DELEGATED_SHARED post: walk.rtte.state == ASSIGNED_DEV_SHARED
rtte_addr	walk.rtte.addr == addr

B4.3.4.4 Footprint

ID	Value
state	GranuleAt (addr) .state
rtte	RttEntryAt (walk.rtt_addr, entry_idx)

DRAFT

B4.3.5 RMI_DEV_MEM_UNMAP command

Unmaps device memory.

Issue Consider how teardown of DRAM mappings (via RMI_DATA_DESTROY) composes with teardown of device memory mappings (via RMI_DEV_MEM_UNMAP). In each case, the command returns the IPA of the next live entry - but it doesn't tell the caller whether this is DRAM or IO. How then can the caller know which of the two commands to call next, while still avoiding a (race-prone) call to RMI_RTT_READ_ENTRY?

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.5.1 Interface

B4.3.5.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000173
rd	X1	63:0	Address	PA of the RD which owns the target device memory Granule
ipa	X2	63:0	Address	IPA at which the Granule is mapped in the target Realm

B4.3.5.1.2 Context

The RMI_DEV_MEM_UNMAP command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL, RMM_RTT_TREE_PRIMARY)	false	RTT walk result
rtte_state_pre	RmmRttEntryState	walk.rtte.state	true	RTT entry state
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
walk_top	Address	RttSkipNonLiveEntries (RttAt (walk.rtt_addr), walk.level, ipa)	false	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

B4.3.5.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
io	X1	63:0	Address	PA of the device memory Granule which was unmapped

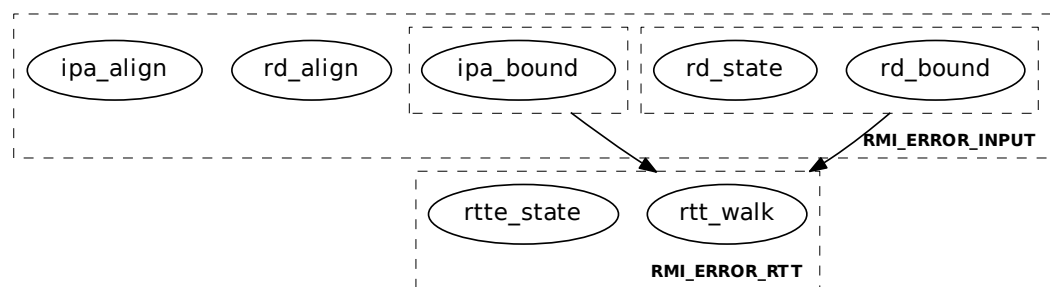
Name	Register	Bits	Type	Description
top	X2	63:0	Address	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

B4.3.5.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsGranuleAligned(ipa) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: !AddrIsProtected(ipa, realm) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < RMM_RTT_PAGE_LEVEL post: (ResultEqual(result, RMI_ERROR_RTT, walk.level) && (top == walk_top))
rtte_state	pre: (walk.rtte.state != ASSIGNED_DEV_PRIVATE && walk.rtte.state != ASSIGNED_DEV_SHARED) post: (ResultEqual(result, RMI_ERROR_RTT, walk.level) && (top == walk_top))

B4.3.5.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state]
[ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.5.3 Success conditions

ID	Condition
comm_state_private	pre: rtte_state_pre == ASSIGNED_DEV_PRIVATE post: GranuleAt(walk.rtte.addr).state == DEV_DELEGATED_PRIVATE
comm_state_shared	pre: rtte_state_pre == ASSIGNED_DEV_SHARED post: GranuleAt(walk.rtte.addr).state == DEV_DELEGATED_SHARED
rtte_state	walk.rtte.state == UNASSIGNED
ripas_ram	pre: walk.rtte.ripas == DEV post: walk.rtte.ripas == DESTROYED
io	io == walk.rtte.addr
top	top == walk_top

B4.3.5.4 Footprint

ID	Value
comm_state	GranuleAt(walk.rtte.addr).state
rtte	RttEntryAt(walk.rtt_addr, entry_idx)

B4.3.6 RMI_FEATURES command

Read feature register.

The following table indicates which feature register is returned depending on the index provided.

Index	Feature register
0	RMI feature register 0
1	RMI feature register 1

See also:

- [Chapter A3 Feature discovery and configuration](#)

B4.3.6.1 Interface

B4.3.6.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000165
index	X1	63:0	UInt64	Feature register index

B4.3.6.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
value	X1	63:0	Bits64	Feature register value

B4.3.6.2 Failure conditions

The RMI_FEATURES command does not have any failure conditions.

B4.3.6.3 Success conditions

ID	Condition
value	<code>value == RmiFeatureRegisterEncode(index)</code>

B4.3.6.4 Footprint

The RMI_FEATURES command does not have any footprint.

B4.3.7 RMI_GRANULE_DELEGATE command

Delegates a Granule.

See also:

- [A2.2 Granule](#)
- [B4.3.10 RMI_GRANULE_UNDELEGATE command](#)
- [D1.2.1 Realm creation flow](#)

B4.3.7.1 Interface

B4.3.7.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000151
addr	X1	63:0	Address	PA of the target Granule

B4.3.7.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.7.2 Failure conditions

ID	Condition
gran_align	pre: !AddrIsGranuleAligned (addr) post: ResultEqual (result, RMI_ERROR_INPUT)
gran_bound	pre: !PaIsDelegable (addr) post: ResultEqual (result, RMI_ERROR_INPUT)
gran_state	pre: GranuleAt (addr).state != UNDELEGATED post: ResultEqual (result, RMI_ERROR_INPUT)

B4.3.7.2.1 Failure condition ordering

The RMI_GRANULE_DELEGATE command does not have any failure condition orderings.

B4.3.7.3 Success conditions

ID	Condition
gran_state	GranuleAt (addr).state == DELEGATED
gran_gpt	GranuleAt (addr).gpt == GPT_REALM

B4.3.7.4 Footprint

ID	Value
gran_gpt	<code>GranuleAt(addr).gpt</code>
gran_state	<code>GranuleAt(addr).state</code>

DRAFT

B4.3.8 RMI_GRANULE_DEV_DELEGATE command

Delegate a Granule of device memory.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.8.1 Interface

B4.3.8.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000170
addr	X1	63:0	Address	PA of the target Granule
flags	X2	63:0	RmiDevDelegateFlags	Flags

B4.3.8.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.8.2 Failure conditions

ID	Condition
gran_align	pre: !AddrIsGranuleAligned (addr) post: ResultEqual (result, RMI_ERROR_INPUT)
gran_bound	pre: !PaIsDelegable (addr) post: ResultEqual (result, RMI_ERROR_INPUT)
gran_state	pre: GranuleAt (addr).state != DEV_UNDELEGATED post: ResultEqual (result, RMI_ERROR_INPUT)

B4.3.8.2.1 Failure condition ordering

The RMI_GRANULE_DEV_DELEGATE command does not have any failure condition orderings.

B4.3.8.3 Success conditions

ID	Condition
state_private	pre: flags.share == RMI_DEV_MEM_PRIVATE post: GranuleAt (addr).state == DEV_DELEGATED_PRIVATE
gpt_private	pre: flags.share == RMI_DEV_MEM_PRIVATE post: GranuleAt (addr).gpt == GPT_REALM
state_shared	pre: flags.share == RMI_DEV_MEM_SHARED post: GranuleAt (addr).state == DEV_DELEGATED_SHARED

ID	Condition
gpt_shared	pre: flags.share == RMI_DEV_MEM_SHARED post: GranuleAt(addr).gpt == GPT_AAP

B4.3.8.4 Footprint

ID	Value
gpt	GranuleAt(addr).gpt
state	GranuleAt(addr).state

DRAFT

B4.3.9 RMI_GRANULE_DEV_UNDELEGATE command

Undelegate a Granule of device memory.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.9.1 Interface

B4.3.9.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000171
addr	X1	63:0	Address	PA of the target Granule

B4.3.9.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.9.2 Failure conditions

ID	Condition
gran_align	pre: !AddrIsGranuleAligned (addr) post: ResultEqual (result, RMI_ERROR_INPUT)
gran_bound	pre: !PaIsDelegable (addr) post: ResultEqual (result, RMI_ERROR_INPUT)
gran_state	pre: (GranuleAt (addr).state != DEV_DELEGATED_PRIVATE && GranuleAt (addr).state != DEV_DELEGATED_SHARED) post: ResultEqual (result, RMI_ERROR_INPUT)

B4.3.9.2.1 Failure condition ordering

The RMI_GRANULE_DEV_UNDELEGATE command does not have any failure condition orderings.

B4.3.9.3 Success conditions

ID	Condition
gpt	GranuleAt (addr).gpt != GPT_REALM
state	GranuleAt (addr).state == DEV_UNDELEGATED

B4.3.9.4 Footprint

ID	Value
gpt	<code>GranuleAt(addr).gpt</code>
state	<code>GranuleAt(addr).state</code>

DRAFT

B4.3.10 RMI_GRANULE_UNDELEGATE command

Undelegates a Granule.

See also:

- [A2.2 Granule](#)
- [B4.3.7 RMI_GRANULE_DELEGATE command](#)
- [D1.2.5 Realm destruction flow](#)

B4.3.10.1 Interface

B4.3.10.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000152
addr	X1	63:0	Address	PA of the target Granule

B4.3.10.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.10.2 Failure conditions

ID	Condition
gran_align	pre: !AddrIsGranuleAligned (addr) post: ResultEqual (result, RMI_ERROR_INPUT)
gran_bound	pre: !PaIsDelegable (addr) post: ResultEqual (result, RMI_ERROR_INPUT)
gran_state	pre: GranuleAt (addr).state != DELEGATED post: ResultEqual (result, RMI_ERROR_INPUT)

B4.3.10.2.1 Failure condition ordering

The RMI_GRANULE_UNDELEGATE command does not have any failure condition orderings.

B4.3.10.3 Success conditions

ID	Condition
gran_gpt	GranuleAt (addr).gpt != GPT_REALM
gran_state	GranuleAt (addr).state == UNDELEGATED
gran_content	Contents of target Granule are wiped.

See also:

- [A2.2.4 Granule wiping](#)

B4.3.10.4 Footprint

ID	Value
gran_gpt	GranuleAt (addr) .gpt
gran_state	GranuleAt (addr) .state

DRAFT

B4.3.11 RMI_MEC_SET_PRIVATE command

Change state of a MEC to Private.

See also:

- [Chapter A11 Realm memory encryption](#)
- [B4.3.12 RMI_MEC_SET_SHARED command](#)

B4.3.11.1 Interface

B4.3.11.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400018D
mecid	X1	63:0	Bits64	MECID

B4.3.11.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.11.2 Failure conditions

ID	Condition
mecid_bound	pre: <code>UInt(mecid) > UInt(ImplFeatures().max_mecid)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
state	pre: <code>MecState(mecid) != MEC_STATE_SHARED</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
members	pre: <code>MecMembers(mecid) != 0</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>

B4.3.11.2.1 Failure condition ordering

The RMI_MEC_SET_PRIVATE command does not have any failure condition orderings.

B4.3.11.3 Success conditions

ID	Condition
mec_state	<code>MecState(mecid) == MEC_STATE_PRIVATE_UNASSIGNED</code>

B4.3.11.4 Footprint

The RMI_MEC_SET_PRIVATE command does not have any footprint.

B4.3.12 RMI_MEC_SET_SHARED command

Change state of a MEC to Shared.

See also:

- [Chapter A11 Realm memory encryption](#)
- [B4.3.11 RMI_MEC_SET_PRIVATE command](#)

B4.3.12.1 Interface

B4.3.12.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400018C
mecid	X1	63:0	Bits64	MECID

B4.3.12.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.12.2 Failure conditions

ID	Condition
mecid_bound	pre: <code>UInt(mecid) > UInt(ImplFeatures().max_mecid)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
state	pre: <code>MecState(mecid) != MEC_STATE_PRIVATE_UNASSIGNED</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>

B4.3.12.2.1 Failure condition ordering

The RMI_MEC_SET_SHARED command does not have any failure condition orderings.

B4.3.12.3 Success conditions

ID	Condition
mec_state	<code>MecState(mecid) == MEC_STATE_SHARED</code>

B4.3.12.4 Footprint

The RMI_MEC_SET_SHARED command does not have any footprint.

B4.3.13 RMI_PDEV_ABORT command

Abort device communication associated with a PDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.13.1 Interface

B4.3.13.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000174
pdev_ptr	X1	63:0	Address	PA of the PDEV

B4.3.13.1.2 Context

The RMI_PDEV_ABORT command operates on the following context.

Name	Type	Value	Before	Description
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV
pdev_state_pre	RmmPdevState	pdev.state	true	Previous state

B4.3.13.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.13.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_bound	pre: !PaIsDelegable (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt (pdev_ptr) .state != PDEV post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_state	pre: (pdev.state != PDEV_NEW && pdev.state != PDEV_HAS_KEY && pdev.state != PDEV_COMMUNICATING) post: ResultEqual (result, RMI_ERROR_DEVICE)

B4.3.13.2.1 Failure condition ordering

[da_supp] < [pdev_align, pdev_bound, pdev_gran_state]
[pdev_gran_state] < [pdev_state]



B4.3.13.3 Success conditions

ID	Condition
state	pre: pdev_state_pre == PDEV_COMMUNICATING post: pdev.state == PDEV_READY
comm_state	pdev.comm_state == DEV_COMM_IDLE

B4.3.13.4 Footprint

ID	Value
state	pdev.state
comm_state	pdev.comm_state

B4.3.14 RMI_PDEV_AUX_COUNT command

Get number of auxiliary Granules required for a PDEV.

B4.3.14.1 Interface

B4.3.14.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000156
flags	X1	63:0	Bits64	PDEV flags

B4.3.14.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
aux_count	X1	63:0	UInt64	Number of auxiliary Granules required for a PDEV

B4.3.14.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures().feat_da != FEATURE_TRUE post: ResultEqual(result, RMI_ERROR_NOT_SUPPORTED)

B4.3.14.3 Success conditions

ID	Condition
aux_count	aux_count == PdevAuxCount(RmiPdevFlagsDecode(flags))

B4.3.14.4 Footprint

The RMI_PDEV_AUX_COUNT command does not have any footprint.

B4.3.15 RMI_PDEV_COMMUNICATE command

Perform device communication associated with a PDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.15.1 Interface

B4.3.15.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000175
pdev_ptr	X1	63:0	Address	PA of the PDEV
data_ptr	X2	63:0	Address	PA of the communication data structure

B4.3.15.1.2 Context

The RMI_PDEV_COMMUNICATE command operates on the following context.

Name	Type	Value	Before	Description
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV
pdev_state_pre	RmmPdevState	PdevAt (pdev_ptr) .state	true	PDEV previous state
data	RmiDevCommData	RmiDevCommDataAt (data_ptr)	false	Device communication object

B4.3.15.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.15.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures ().feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_bound	pre: !PaIsDelegable (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt (pdev_ptr) .state != PDEV post: ResultEqual (result, RMI_ERROR_INPUT)
data_align	pre: !AddrIsGranuleAligned (data_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
data_bound	pre: !PaIsDelegable(data_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
data_pas	pre: !GranuleAccessPermitted(data_ptr, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
req_align	pre: !AddrIsGranuleAligned(data.enter.req_addr) post: ResultEqual(result, RMI_ERROR_INPUT)
req_bound	pre: !PaIsDelegable(data.enter.req_addr) post: ResultEqual(result, RMI_ERROR_INPUT)
req_pas	pre: !GranuleAccessPermitted(data.enter.req_addr, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
resp_align	pre: !AddrIsGranuleAligned(data.enter.resp_addr) post: ResultEqual(result, RMI_ERROR_INPUT)
resp_bound	pre: !PaIsDelegable(data.enter.resp_addr) post: ResultEqual(result, RMI_ERROR_INPUT)
resp_pas	pre: !GranuleAccessPermitted(data.enter.resp_addr, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
resp_len	pre: data.enter.resp_len > RMM_GRANULE_SIZE post: ResultEqual(result, RMI_ERROR_INPUT)
comm_state	pre: (pdev.comm_state == DEV_COMM_IDLE pdev.comm_state == DEV_COMM_ERROR) post: ResultEqual(result, RMI_ERROR_DEVICE)

B4.3.15.2.1 Failure condition ordering

```
[da_supp] < [pdev_align, pdev_bound, pdev_gran_state, data_align,
data_bound, data_pas, req_align, req_bound, req_pas, resp_align,
resp_bound, resp_pas, resp_len]
[pdev_gran_state] < [comm_state]
```



B4.3.15.3 Success conditions

ID	Condition
comm_state	pdev.comm_state == DeviceCommunicate(pdev, data)
error	pre: (DeviceCommunicate(pdev, data) == DEV_COMM_ERROR && pdev.state != PDEV_STOPPING) post: pdev.state == PDEV_ERROR
new	pre: (DeviceCommunicate(pdev, data) == DEV_COMM_IDLE && pdev_state_pre == PDEV_NEW) post: pdev.state == PDEV_NEEDS_KEY

ID	Condition
has_key	pre: (DeviceCommunicate(pdev, data) == DEV_COMM_IDLE && pdev_state_pre == PDEV_HAS_KEY) post: pdev.state == PDEV_READY
ready	pre: (DeviceCommunicate(pdev, data) == DEV_COMM_IDLE && pdev_state_pre == PDEV_READY) post: pdev.state == PDEV_READY
stopped	pre: (DeviceCommunicate(pdev, data) != DEV_COMM_ACTIVE && pdev_state_pre == PDEV_STOPPING) post: pdev.state == PDEV_STOPPED
communicating	pre: (DeviceCommunicate(pdev, data) == DEV_COMM_IDLE && pdev_state_pre == PDEV_COMMUNICATING) post: pdev.state == PDEV_READY
ide_resetting	pre: (DeviceCommunicate(pdev, data) == DEV_COMM_IDLE && pdev_state_pre == PDEV_IDE_RESETTING) post: pdev.state == PDEV_READY

B4.3.15.4 Footprint

ID	Value
state	pdev.state
comm_state	pdev.comm_state

B4.3.16 RMI_PDEV_CREATE command

Create a PDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.16.1 Interface

B4.3.16.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000176
pdev_ptr	X1	63:0	Address	PA of the PDEV
params_ptr	X2	63:0	Address	PA of PDEV parameters

B4.3.16.1.2 Context

The RMI_PDEV_CREATE command operates on the following context.

Name	Type	Value	Before	Description
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV
params	RmiPdevParams	RmiPdevParamsAt (params_ptr)	false	PDEV parameters

B4.3.16.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.16.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_bound	pre: !PaIsDelegable (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_state	pre: GranuleAt (pdev_ptr) .state != DELEGATED post: ResultEqual (result, RMI_ERROR_INPUT)
params_align	pre: !AddrIsGranuleAligned (params_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
params_bound	pre: !PaIsDelegable (params_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
params_pas	pre: !GranuleAccessPermitted(params_ptr, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
params_valid	pre: !RmiPdevParamsIsValid(params_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
num_aux	pre: params.num_aux != PdevAuxCount(params.flags) post: ResultEqual(result, RMI_ERROR_INPUT)
aux_align	pre: !AuxAligned32(params.aux, params.num_aux) post: ResultEqual(result, RMI_ERROR_INPUT)
aux_alias	pre: AuxAlias32(pdev_ptr, params.aux, params.num_aux) post: ResultEqual(result, RMI_ERROR_INPUT)
aux_state	pre: !AuxStateEqual32(params.aux, params.num_aux, DELEGATED) post: ResultEqual(result, RMI_ERROR_INPUT)

B4.3.16.2.1 Failure condition ordering

```
[da_supp] < [pdev_align, pdev_bound, pdev_state, params_align,
params_bound, params_pas, params_valid, num_aux, aux_align,
aux_alias, aux_state]
```



B4.3.16.3 Success conditions

ID	Condition
gran_state	GranuleAt(pdev_ptr).state == PDEV
pdev_id	pdev.pdev_id == params.pdev_id
prot_config	Equal(pdev.prot_config, params.flags.prot_config)
segment_id	pdev.segment_id == params.segment_id
root_id	pdev.root_id == params.root_id
cert_id	pdev.cert_id == params.cert_id
rid_base	pdev.rid_base == params.rid_base
rid_top	pdev.rid_top == params.rid_top
hash_algo	Equal(pdev.hash_algo, params.hash_algo)
ide_sid	pdev.ide_sid == params.ide_sid
iocoh_num_addr_range	pdev.iocoh_num_addr_range == params.iocoh_num_addr_range
iocoh_addr_range	RmiAddressRangesEqual16(pdev.iocoh_addr_range, params.iocoh_addr_range, params.iocoh_num_addr_range)

ID	Condition
fcoh_num_addr_range	<code>pdev.fcoh_num_addr_range == params.fcoh_num_addr_range</code>
fcoh_addr_range	<code>RmiAddressRangesEqual4(pdev.fcoh_addr_range, params.fcoh_addr_range, params.fcoh_num_addr_range)</code>
state	<code>pdev.state == PDEV_NEW</code>
comm_state	<code>pdev.comm_state == DEV_COMM_PENDING</code>
num_vdevs	<code>pdev.num_vdevs == 0</code>
aux	<code>AuxEqual32(pdev.aux, params.aux, PdevAuxCount(params.flags))</code>
num_aux	<code>pdev.num_aux == PdevAuxCount(params.flags)</code>
aux_state	<code>AuxStateEqual32(pdev.aux, PdevAuxCount(params.flags), PDEV_AUX)</code>

B4.3.16.4 Footprint

ID	Value
state	<code>GranuleAt(pdev_ptr).state</code>
aux_state	<code>AuxStates(pdev.aux, PdevAuxCount(params.flags))</code>

B4.3.17 RMI_PDEV_DESTROY command

Destroy a PDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.17.1 Interface

B4.3.17.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000177
pdev_ptr	X1	63:0	Address	PA of the PDEV

B4.3.17.1.2 Context

The RMI_PDEV_DESTROY command operates on the following context.

Name	Type	Value	Before	Description
pdev_pre	RmmPdev	PdevAt (pdev_ptr)	true	PDEV

B4.3.17.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.17.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_bound	pre: !PaIsDelegable (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt (pdev_ptr) .state != PDEV post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_state	pre: pdev_pre.state != PDEV_STOPPED post: ResultEqual (result, RMI_ERROR_DEVICE)

B4.3.17.2.1 Failure condition ordering

[pdev_gran_state] < [pdev_state]

[da_supp] < [pdev_align, pdev_bound, pdev_gran_state, pdev_state]



B4.3.17.3 Success conditions

ID	Condition
gran_state	<code>GranuleAt(pdev_ptr).state == DELEGATED</code>
aux_state	<code>AuxStateEqual32(pdev_pre.aux, pdev_pre.num_aux, DELEGATED)</code>

B4.3.17.4 Footprint

ID	Value
state	<code>GranuleAt(pdev_ptr).state</code>
aux_state	<code>AuxStates(pdev_pre.aux, prev_pre._num_aux)</code>

B4.3.18 RMI_PDEV_GET_STATE command

Get state of a PDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.18.1 Interface

B4.3.18.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000178
pdev_ptr	X1	63:0	Address	PA of the PDEV

B4.3.18.1.2 Context

The RMI_PDEV_GET_STATE command operates on the following context.

Name	Type	Value	Before	Description
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV

B4.3.18.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
state	X1	7:0	RmiPdevState	PDEV state


The following unused bits of RMI_PDEV_GET_STATE output values MBZ: X1[63:8].

B4.3.18.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_bound	pre: !PaIsDelegable (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt (pdev_ptr) .state != PDEV post: ResultEqual (result, RMI_ERROR_INPUT)

B4.3.18.2.1 Failure condition ordering

[da_supp] < [pdev_align, pdev_bound, pdev_gran_state]



B4.3.18.3 Success conditions

ID	Condition
state	<code>Equal(state, pdev.state)</code>

B4.3.18.4 Footprint

The RMI_PDEV_GET_STATE command does not have any footprint.

B4.3.19 RMI_PDEV_IDE_RESET command

Reset the IDE link of a PDEV.

B4.3.19.1 Interface

B4.3.19.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xc4000179
pdev_ptr	X1	63:0	Address	PA of the PDEV

B4.3.19.1.2 Context

The RMI_PDEV_IDE_RESET command operates on the following context.

Name	Type	Value	Before	Description
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV

B4.3.19.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.19.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures().feat_da != FEATURE_TRUE post: ResultEqual(result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned(pdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
pdev_bound	pre: !PaIsDelegable(pdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt(pdev_ptr).state != PDEV post: ResultEqual(result, RMI_ERROR_INPUT)
pdev_state	pre: pdev.state != PDEV_READY post: ResultEqual(result, RMI_ERROR_DEVICE)

B4.3.19.2.1 Failure condition ordering

```
[da_supp] < [pdev_align, pdev_bound, pdev_gran_state]
[pdev_gran_state] < [pdev_state]
```



B4.3.19.3 Success conditions

ID	Condition
pdev_state	pdev.state == PDEV_IDE_RESETTING
comm_state	pdev.comm_state == DEV_COMM_PENDING

B4.3.19.4 Footprint

ID	Value
state	pdev.state
comm_state	pdev.comm_state

B4.3.20 RMI_PDEV_NOTIFY command

Notify the RMM of an event related to a PDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.20.1 Interface

B4.3.20.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400017A
pdev_ptr	X1	63:0	Address	PA of the PDEV
ev	X2	7:0	RmiPdevEvent	Event type

The following unused bits of RMI_PDEV_NOTIFY input values SBZ: X2[63:8].

B4.3.20.1.2 Context

The RMI_PDEV_NOTIFY command operates on the following context.

Name	Type	Value	Before	Description
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV

B4.3.20.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.20.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_bound	pre: !PaIsDelegable (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt (pdev_ptr) .state != PDEV post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_state	pre: pdev.state != PDEV_READY post: ResultEqual (result, RMI_ERROR_DEVICE)
ev_valid	pre: !RmiPdevEventIsValid (ev) post: ResultEqual (result, RMI_ERROR_INPUT)

B4.3.20.2.1 Failure condition ordering

[da_supp] < [pdev_align, pdev_bound, pdev_gran_state]
[pdev_gran_state] < [pdev_state]
[pdev_state] < [ev_valid]



B4.3.20.3 Success conditions

ID	Condition
pdev_state	pdev.state == PDEV_COMMUNICATING
comm_state	pdev.comm_state == DEV_COMM_PENDING

B4.3.20.4 Footprint

ID	Value
state	pdev.state
comm_state	pdev.comm_state

B4.3.21 RMI_PDEV_SET_PUBKEY command

Provide public key associated with a PDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.21.1 Interface

B4.3.21.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400017B
pdev_ptr	X1	63:0	Address	PA of the PDEV
key	X2	63:0	Address	PA of the key
len	X3	63:0	UInt64	Length of the key in bytes
algo	X4	7:0	RmiSignatureAlgorithm	Signature algorithm

The following unused bits of RMI_PDEV_SET_PUBKEY input values SBZ: X4[63:8].

B4.3.21.1.2 Context

The RMI_PDEV_SET_PUBKEY command operates on the following context.

Name	Type	Value	Before	Description
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV

B4.3.21.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.21.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_bound	pre: !PaIsDelegable (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt (pdev_ptr) .state != PDEV post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
key_align	pre: !AddrIsGranuleAligned(key) post: ResultEqual(result, RMI_ERROR_INPUT)
key_bound	pre: !PaIsDelegable(key) post: ResultEqual(result, RMI_ERROR_INPUT)
key_pas	pre: !GranuleAccessPermitted(key, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
key_len_oflow	pre: len > RMM_GRANULE_SIZE post: ResultEqual(result, RMI_ERROR_INPUT)
key_invalid	pre: Key is invalid, for example length is invalid for specified signature algorithm. post: ResultEqual(result, RMI_ERROR_INPUT)
pdev_state	pre: pdev.state != PDEV_NEEDS_KEY post: ResultEqual(result, RMI_ERROR_DEVICE)

B4.3.21.2.1 Failure condition ordering

```
[da_supp] < [pdev_align, pdev_bound, pdev_gran_state, key_align,
             key_bound, key_pas, key_len_oflow, key_invalid]
[pdev_gran_state] < [pdev_state]
```



B4.3.21.3 Success conditions

ID	Condition
state	pdev.state == PDEV_HAS_KEY
comm_state	pdev.comm_state == DEV_COMM_PENDING

B4.3.21.4 Footprint

ID	Value
state	pdev.state
comm_state	pdev.comm_state

B4.3.22 RMI_PDEV_STOP command

Stop a PDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.22.1 Interface

B4.3.22.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400017C
pdev_ptr	X1	63:0	Address	PA of the PDEV

B4.3.22.1.2 Context

The RMI_PDEV_STOP command operates on the following context.

Name	Type	Value	Before	Description
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV

B4.3.22.1.3 Output values

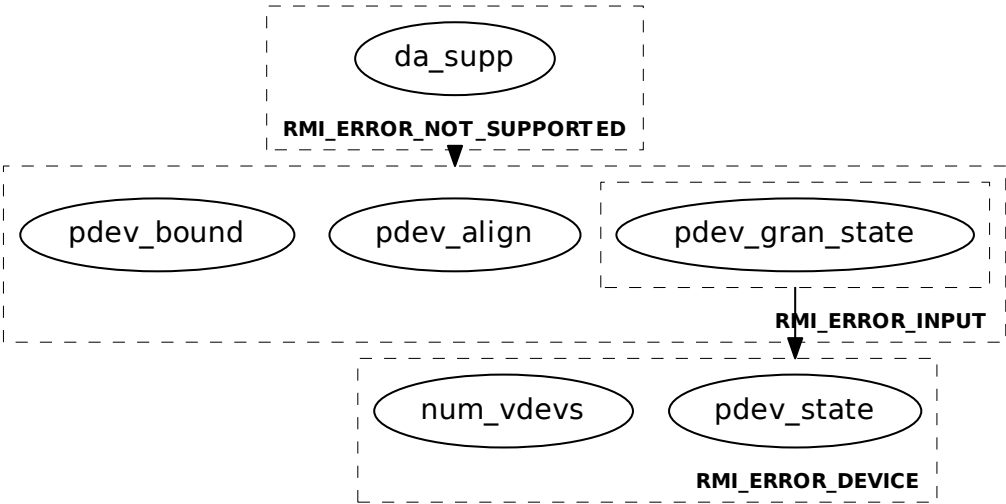
Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.22.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_bound	pre: !PaIsDelegable (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt (pdev_ptr) .state != PDEV post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_state	pre: (pdev.state == PDEV_COMMUNICATING pdev.state == PDEV_STOPPING pdev.state == PDEV_STOPPED) post: ResultEqual (result, RMI_ERROR_DEVICE)
num_vdevs	pre: pdev.num_vdevs != 0 post: ResultEqual (result, RMI_ERROR_DEVICE)

B4.3.22.2.1 Failure condition ordering

[da_supp] < [pdev_align, pdev_bound, pdev_gran_state]
[pdev_gran_state] < [pdev_state, num_vdevs]



B4.3.22.3 Success conditions

ID	Condition
pdev_state	pdev.state == PDEV_STOPPING
comm_state	pdev.comm_state == DEV_COMM_PENDING

B4.3.22.4 Footprint

ID	Value
state	pdev.state
comm_state	pdev.comm_state

B4.3.23 RMI_PSCI_COMPLETE command

Completes a pending PSCI command which was called with an MPIDR argument, by providing the corresponding REC.

See also:

- [A4.3.7 REC exit due to PSCI](#)
- [B6.3.1 PSCI_AFFINITY_INFO command](#)
- [B6.3.3 PSCI_CPU_ON command](#)
- [D1.4 PSCI flows](#)

B4.3.23.1 Interface

B4.3.23.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000164
calling_rec_ptr	X1	63:0	Address	PA of the calling REC
target_rec_ptr	X2	63:0	Address	PA of the target REC
status	X3	63:0	PsciReturnCode	Status of the PSCI request

B4.3.23.1.2 Context

The RMI_PSCI_COMPLETE command operates on the following context.

Name	Type	Value	Before	Description
calling_rec	RmmRec	RecAt (calling_rec_ptr)	false	Calling REC
target_rec	RmmRec	RecAt (target_rec_ptr)	false	Target REC

B4.3.23.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.23.2 Failure conditions

ID	Condition
alias	pre: calling_rec_ptr == target_rec_ptr post: ResultEqual (result, RMI_ERROR_INPUT)
calling_align	pre: !AddrIsGranuleAligned (calling_rec_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
calling_bound	pre: !PaIsDelegable (calling_rec_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
calling_state	pre: <code>GranuleAt(calling_rec_ptr).state != REC</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
target_align	pre: <code>!AddrIsGranuleAligned(target_rec_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
target_bound	pre: <code>!PaIsDelegable(target_rec_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
target_state	pre: <code>GranuleAt(target_rec_ptr).state != REC</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
pending	pre: <code>calling_rec.pending != REC_PENDING_PSCI</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
owner	pre: <code>target_rec.owner != calling_rec.owner</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
target	pre: <code>target_rec.mpidr != calling_rec.gprs[1]</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
status	pre: <code>!PsciReturnCodePermitted(calling_rec, target_rec, status)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>

B4.3.23.2.1 Failure condition ordering

The RMI_PSCI_COMPLETE command does not have any failure condition orderings.

B4.3.23.3 Success conditions

ID	Condition
pending	<code>calling_rec.pending == REC_PENDING_NONE</code>
on_already	pre: <code>(status == PSCI_SUCCESS</code> <code>&& calling_rec.gprs[0] == FID_PSCI_CPU_ON</code> <code>&& target_rec.flags.runnable == RUNNABLE)</code> post: <code>(calling_rec.gprs[0] ==</code> <code>PsciReturnCodeEncode(PSCI_ALREADY_ON))</code>

ID	Condition
on_success	<pre> pre: (status == PSCI_SUCCESS && calling_rec.gprs[0] == FID_PSCI_CPU_ON && target_rec.flags.runnable != RUNNABLE) post: (target_rec.gprs[0] == calling_rec.gprs[3] && target_rec.gprs[1] == Zeros(64) && target_rec.gprs[2] == Zeros(64) && target_rec.gprs[3] == Zeros(64) && target_rec.gprs[4] == Zeros(64) && target_rec.gprs[5] == Zeros(64) && target_rec.gprs[6] == Zeros(64) && target_rec.gprs[7] == Zeros(64) && target_rec.gprs[8] == Zeros(64) && target_rec.gprs[9] == Zeros(64) && target_rec.gprs[10] == Zeros(64) && target_rec.gprs[11] == Zeros(64) && target_rec.gprs[12] == Zeros(64) && target_rec.gprs[13] == Zeros(64) && target_rec.gprs[14] == Zeros(64) && target_rec.gprs[15] == Zeros(64) && target_rec.gprs[16] == Zeros(64) && target_rec.gprs[17] == Zeros(64) && target_rec.gprs[18] == Zeros(64) && target_rec.gprs[19] == Zeros(64) && target_rec.gprs[20] == Zeros(64) && target_rec.gprs[21] == Zeros(64) && target_rec.gprs[22] == Zeros(64) && target_rec.gprs[23] == Zeros(64) && target_rec.gprs[24] == Zeros(64) && target_rec.gprs[25] == Zeros(64) && target_rec.gprs[26] == Zeros(64) && target_rec.gprs[27] == Zeros(64) && target_rec.gprs[28] == Zeros(64) && target_rec.gprs[29] == Zeros(64) && target_rec.gprs[30] == Zeros(64) && target_rec.gprs[31] == Zeros(64) && target_rec.pc == calling_rec.gprs[2] && target_rec.flags.runnable == RUNNABLE && calling_rec.gprs[0] == PsciReturnCodeEncode(PSCI_SUCCESS)) </pre>
affinity_on	<pre> pre: (status == PSCI_SUCCESS && calling_rec.gprs[0] == FID_PSCI_AFFINITY_INFO && target_rec.flags.runnable == RUNNABLE) post: (calling_rec.gprs[0] == PsciReturnCodeEncode(PSCI_SUCCESS)) </pre>
affinity_off	<pre> pre: (status == PSCI_SUCCESS && calling_rec.gprs[0] == FID_PSCI_AFFINITY_INFO && target_rec.flags.runnable != RUNNABLE) post: (calling_rec.gprs[0] == PsciReturnCodeEncode(PSCI_OFF)) </pre>
status	<pre> pre: status != PSCI_SUCCESS post: (calling_rec.gprs[0] == PsciReturnCodeEncode(status)) </pre>
args	<pre> (calling_rec.gprs[1] == Zeros(64) && calling_rec.gprs[2] == Zeros(64) && calling_rec.gprs[3] == Zeros(64)) </pre>

B4.3.23.4 Footprint

ID	Value
target_flags	target_rec.flags
target_gprs	target_rec.gprs
target_pc	target_rec.pc
calling_pend	calling_rec.pending
calling_gprs	calling_rec.gprs

DRAFT

B4.3.24 RMI_REALM_ACTIVATE command

Activates a Realm.

See also:

- [A2.1 Realm](#)

B4.3.24.1 Interface

B4.3.24.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000157
rd	X1	63:0	Address	PA of the RD

B4.3.24.1.2 Context

The RMI_REALM_ACTIVATE command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm

B4.3.24.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.24.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt (rd).state != RD post: ResultEqual (result, RMI_ERROR_INPUT)
realm_state	pre: realm.state != REALM_NEW post: ResultEqual (result, RMI_ERROR_REALM)

B4.3.24.2.1 Failure condition ordering

[rd_bound, rd_state] < [realm_state]



B4.3.24.3 Success conditions

ID	Condition
realm_state	realm.state == REALM_ACTIVE

B4.3.24.4 Footprint

ID	Value
realm_state	realm.state

B4.3.25 RMI_REALM_CREATE command

Creates a Realm.

See also:

- [A2.1 Realm](#)
- [A2.1.6 Realm parameters](#)
- [B4.3.26 RMI_REALM_DESTROY command](#)
- [D1.2.1 Realm creation flow](#)

B4.3.25.1 Interface

B4.3.25.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xc4000158
rd	X1	63:0	Address	PA of the RD
params_ptr	X2	63:0	Address	PA of Realm parameters

B4.3.25.1.2 Context

The RMI_REALM_CREATE command operates on the following context.

Name	Type	Value	Before	Description
params	RmiRealmParams	RmiRealmParamsAt (params_ptr)	false	Realm parameters
realm	RmmRealm	RealmAt (rd)	false	Realm
mec_members_pre	UInt64	MecMembers (params.mecid)	true	Number of Realms which are members of the Realm's MEC
mec_state_pre	RmmMecState	MecState (params.mecid)	true	MECID state

B4.3.25.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.25.2 Failure conditions

ID	Condition
params_align	pre: !AddrIsGranuleAligned (params_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
params_bound	pre: !PaIsDelegable (params_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
params_pas	pre: <code>!GranuleAccessPermitted(params_ptr, PAS_NS)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
params_valid	pre: <code>!RmiRealmParamsIsValid(params_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
params_supp	pre: <code>!RealmParamsSupported(params)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
alias	pre: <code>AddrInRange(rd, params.rtt_base, (params.rtt_num_start - 1) * RMM_GRANULE_SIZE)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
rd_align	pre: <code>!AddrIsGranuleAligned(rd)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
rd_bound	pre: <code>!PaIsDelegable(rd)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
rd_state	pre: <code>GranuleAt(rd).state != DELEGATED</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
rtt_align	pre: <code>!AddrIsAligned(params.rtt_base, params.rtt_num_start * RMM_GRANULE_SIZE)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
rtt_num_level	pre: <code>!RttConfigIsValid(params.s2sz, params.rtt_level_start, params.rtt_num_start)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
rtt_state	pre: <code>!RttsStateEqual(params.rtt_base, params.rtt_num_start, DELEGATED)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
vmid_valid	pre: <code>(!VmidsAreValid(params.vmid, params.aux_vmid) !VmidsAreFree(params.vmid, params.aux_vmid))</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
mecid_bound	pre: <code>UInt(params.mecid) > UInt(ImplFeatures().max_mecid)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
mecid_state	pre: <code>MecState(params.mecid) == MEC_STATE_PRIVATE_ASSIGNED</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>

B4.3.25.2.1 Failure condition ordering

The RMI_REALM_CREATE command does not have any failure condition orderings.

B4.3.25.3 Success conditions

ID	Condition
rd_state	<code>GranuleAt(rd).state == RD</code>
realm_state	<code>realm.state == REALM_NEW</code>
rec_index	<code>realm.rec_index == 0</code>
rtt_base	<code>RealmRttBaseEqual(realm, params.rtt_base, params.aux_rtt_base)</code>
rtt_state	<code>RttsStateEqual(realm.rtt_base[0], realm.rtt_num_start, RTT)</code>

ID	Condition
rtte_p_states	<code>RttsAllProtectedEntriesState(realm.rtt_base[0], realm.rtt_num_start, UNASSIGNED)</code>
rtte_up_states	<code>RttsAllUnprotectedEntriesState(realm.rtt_base[0], realm.rtt_num_start, UNASSIGNED_NS)</code>
rtte_ripas	<code>RttsAllProtectedEntriesRipas(realm.rtt_base[0], realm.rtt_num_start, EMPTY)</code>
lpa2	<code>Equal(realm.feat_lpa2, params.flags0.lpa2)</code>
ipa_width	<code>realm.ipa_width == params.s2sz</code>
hash_algo	<code>Equal(realm.hash_algo, params.hash_algo)</code>
rim	<code>realm.measurements[0] == RimInit(realm.hash_algo, params)</code>
rem	<code>(realm.measurements[1] == Zeros(RMM_REALM_MEASUREMENT_WIDTH) && realm.measurements[2] == Zeros(RMM_REALM_MEASUREMENT_WIDTH) && realm.measurements[3] == Zeros(RMM_REALM_MEASUREMENT_WIDTH) && realm.measurements[4] == Zeros(RMM_REALM_MEASUREMENT_WIDTH))</code>
rtt_level	<code>realm.rtt_level_start == params.rtt_level_start</code>
rtt_num	<code>realm.rtt_num_start == params.rtt_num_start</code>
vmid	<code>RealmVmidEqual(realm, params.vmid, params.aux_vmid)</code>
rpv	<code>realm.rpv == params.rpv</code>
da	<code>Equal(realm.feat_da, params.flags0.da)</code>
rtt_tree_pp	<code>Equal(realm.rtt_tree_pp, params.flags1.rtt_tree_pp)</code>
num_aux_planes	<code>realm.num_aux_planes == params.num_aux_planes</code>
lfa_policy	<code>Equal(realm.lfa_policy, params.flags0.lfa_policy)</code>
mecid	<code>realm.mecid == params.mecid</code>
mec_policy	<code>realm.mec_policy == MecPolicy(realm.mecid)</code>
mecid_private	<code>pre: mec_state_pre == MEC_STATE_PRIVATE_UNASSIGNED post: MecState(params.mecid) == MEC_STATE_PRIVATE_ASSIGNED</code>
mec_members	<code>pre: mec_state_pre == MEC_STATE_SHARED post: MecMembers(params.mecid) == mec_members_pre + 1</code>
num_recs	<code>realm.num_recs == 0</code>
num_vdevs	<code>realm.num_vdevs == 0</code>

B4.3.25.4 RMI_REALM_CREATE initialization of RIM

On successful execution of RMI_REALM_CREATE, the initial RIM value of the target Realm is calculated by the RMM as follows:

1. Allocate a zero-filled [RmiRealmParams](#) data structure to hold the measured Realm parameters.
2. Copy the following attributes from the Host-provided [RmiRealmParams](#) data structure into the measured Realm parameters data structure:

- flags0
 - s2sz
 - sve_vl
 - num_bps
 - num_wps
 - pmu_num_ctrs
 - hash_algo
3. Using the RHA of the target Realm, compute the hash of the measured Realm parameters data structure. Set the RIM of the target Realm to this value, zero filling upper bytes if the RHA output is smaller than the size of the RIM.

See also:

- [A7.1.1 Realm Initial Measurement](#)
- [B3.71 RimInit function](#)
- [B4.4.31 RmiRealmParams type](#)

B4.3.25.5 Footprint

ID	Value
rd_state	<code>GranuleAt(rd).state</code>
rtt_state	<code>RttsGranuleState(realm.rtt_base[0], realm.rtt_num_start)</code>

B4.3.26 RMI_REALM_DESTROY command

Destroys a Realm.

See also:

- [A2.1 Realm](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [D1.2.5 Realm destruction flow](#)

B4.3.26.1 Interface

B4.3.26.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000159
rd	X1	63:0	Address	PA of the RD

B4.3.26.1.2 Context

The RMI_REALM_DESTROY command operates on the following context.

Name	Type	Value	Before	Description
realm_pre	RmmRealm	RealmAt (rd)	true	Realm
mec_members_pre	UInt64	MecMembers (realm_pre.mecid)	true	Number of Realms which are members of the Realm's MEC

B4.3.26.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.26.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt (rd).state != RD post: ResultEqual (result, RMI_ERROR_INPUT)
realm_live	pre: RealmIsLive (rd) post: ResultEqual (result, RMI_ERROR_REALM)

B4.3.26.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [realm_live]
```



B4.3.26.3 Success conditions

ID	Condition
rtt_state	<code>RttsStateEqual(realm_pre.rtt_base[0], realm_pre.rtt_num_start, DELEGATED)</code>
rd_state	<code>GranuleAt(rd).state == DELEGATED</code>
vmid	<code>VmidsAreFree(realm_pre.vmid)</code>
mecid_private	pre: <code>realm_pre.mec_policy == MEC_POLICY_PRIVATE</code> post: <code>MecState(realm_pre.mecid) == MEC_STATE_PRIVATE_UNASSIGNED</code>
mec_members	pre: <code>realm_pre.mec_policy == MEC_POLICY_SHARED</code> post: <code>MecMembers(realm_pre.mecid) == mec_members_pre - 1</code>

B4.3.26.4 Footprint

ID	Value
rd_state	<code>GranuleAt(rd).state</code>
rtt_state	<code>RttsGranuleState(realm_pre.rtt_base[0], realm_pre.rtt_num_start)</code>

B4.3.27 RMI_REC_AUX_COUNT command

Get number of auxiliary Granules required for a REC.

See also:

- [A2.3 Realm Execution Context](#)
- [B4.3.28 RMI_REC_CREATE command](#)
- [B4.4.39 RmiRecParams type](#)
- [D1.2.4 REC creation flow](#)

B4.3.27.1 Interface

B4.3.27.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000167
rd	X1	63:0	Address	PA of the RD for the target Realm

B4.3.27.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
aux_count	X1	63:0	UInt64	Number of auxiliary Granules required for a REC

B4.3.27.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt (rd).state != RD post: ResultEqual (result, RMI_ERROR_INPUT)

B4.3.27.2.1 Failure condition ordering

The RMI_REC_AUX_COUNT command does not have any failure condition orderings.

B4.3.27.3 Success conditions

ID	Condition
aux_count	aux_count == RecAuxCount (rd)

B4.3.27.4 Footprint

The RMI_REC_AUX_COUNT command does not have any footprint.

B4.3.28 RMI_REC_CREATE command

Creates a REC.

See also:

- [A2.3 Realm Execution Context](#)
- [A2.3.3 REC index and MPIDR value](#)
- [B4.3.27 RMI_REC_AUX_COUNT command](#)
- [B4.3.29 RMI_REC_DESTROY command](#)
- [D1.2.4 REC creation flow](#)

B4.3.28.1 Interface

B4.3.28.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400015A
rd	X1	63:0	Address	PA of the RD for the target Realm
rec_ptr	X2	63:0	Address	PA of the target REC
params_ptr	X3	63:0	Address	PA of REC parameters

B4.3.28.1.2 Context

The RMI_REC_CREATE command operates on the following context.

Name	Type	Value	Before	Description
realm_pre	RmmRealm	RealmAt (rd)	true	Realm
realm	RmmRealm	RealmAt (rd)	false	Realm
params	RmiRecParams	RmiRecParamsAt (params_ptr)	false	REC parameters
rec	RmmRec	RecAt (rec_ptr)	false	REC

B4.3.28.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.28.2 Failure conditions

ID	Condition
params_align	pre: !AddrIsGranuleAligned (params_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
params_bound	pre: !PaIsDelegable (params_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
params_pas	pre: !GranuleAccessPermitted(params_ptr, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
rec_align	pre: !AddrIsGranuleAligned(rec_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
rec_bound	pre: !PaIsDelegable(rec_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
rec_state	pre: GranuleAt(rec_ptr).state != DELEGATED post: ResultEqual(result, RMI_ERROR_INPUT)
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
realm_state	pre: realm_pre.state != REALM_NEW post: ResultEqual(result, RMI_ERROR_REALM)
num_recs	pre: realm_pre.num_recs == (2 ^ ImplFeatures().max_recs_order) - 1 post: ResultEqual(result, RMI_ERROR_REALM)
mpidr_index	pre: RecIndex(params.mpidr) != realm_pre.rec_index post: ResultEqual(result, RMI_ERROR_INPUT)
num_aux	pre: params.num_aux != RecAuxCount(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
aux_align	pre: !AuxAligned16(params.aux, params.num_aux) post: ResultEqual(result, RMI_ERROR_INPUT)
aux_alias	pre: AuxAlias16(rec_ptr, params.aux, params.num_aux) post: ResultEqual(result, RMI_ERROR_INPUT)
aux_state	pre: !AuxStateEqual16(params.aux, params.num_aux, DELEGATED) post: ResultEqual(result, RMI_ERROR_INPUT)

B4.3.28.2.1 Failure condition ordering

[rd_bound, rd_state] < [realm_state, num_recs]



B4.3.28.3 Success conditions

ID	Condition
rec_index	realm.rec_index == realm_pre.rec_index + 1
rec_gran_state	GranuleAt(rec_ptr).state == REC
rec_owner	rec.owner == rd

ID	Condition
rec_attest	rec.attest_state == NO_ATTEST_IN_PROGRESS
rec_mpidr	MpidrEqual(rec.mpidr, params.mpidr)
rec_state	rec.state == REC_READY
runnable	pre: params.flags.runnable == RMI_RUNNABLE post: rec.flags.runnable == RUNNABLE
not_runnable	pre: params.flags.runnable == RMI_NOT_RUNNABLE post: rec.flags.runnable == NOT_RUNNABLE
rec_gprs	(rec.gprs[0] == params.gprs[0] && rec.gprs[1] == params.gprs[1] && rec.gprs[2] == params.gprs[2] && rec.gprs[3] == params.gprs[3] && rec.gprs[4] == params.gprs[4] && rec.gprs[5] == params.gprs[5] && rec.gprs[6] == params.gprs[6] && rec.gprs[7] == params.gprs[7] && rec.gprs[8] == Zeros(64) && rec.gprs[9] == Zeros(64) && rec.gprs[10] == Zeros(64) && rec.gprs[11] == Zeros(64) && rec.gprs[12] == Zeros(64) && rec.gprs[13] == Zeros(64) && rec.gprs[14] == Zeros(64) && rec.gprs[15] == Zeros(64) && rec.gprs[16] == Zeros(64) && rec.gprs[17] == Zeros(64) && rec.gprs[18] == Zeros(64) && rec.gprs[19] == Zeros(64) && rec.gprs[20] == Zeros(64) && rec.gprs[21] == Zeros(64) && rec.gprs[22] == Zeros(64) && rec.gprs[23] == Zeros(64) && rec.gprs[24] == Zeros(64) && rec.gprs[25] == Zeros(64) && rec.gprs[26] == Zeros(64) && rec.gprs[27] == Zeros(64) && rec.gprs[28] == Zeros(64) && rec.gprs[29] == Zeros(64) && rec.gprs[30] == Zeros(64) && rec.gprs[31] == Zeros(64))
rec_pc	rec.pc == params.pc
rim	pre: params.flags.runnable == RMI_RUNNABLE post: realm.measurements[0] == RimExtendRec(realm_pre, params)
rec_aux	AuxEqual16(rec.aux, params.aux, RecAuxCount(rd))
rec_aux_state	AuxStateEqual16(rec.aux, RecAuxCount(rd), REC_AUX)
ripas_addr	rec.ripas_addr == Zeros(ADDRESS_WIDTH)
ripas_top	rec.ripas_top == Zeros(ADDRESS_WIDTH)
pending	rec.pending == REC_PENDING_NONE

ID	Condition
num_recs	<code>realm.num_recs == realm_pre.num_recs + 1</code>
gic_owner	<code>rec.gic_owner == 0</code>

B4.3.28.4 RMI_REC_CREATE extension of RIM

On successful execution of RMI_REC_CREATE, if the new REC is runnable then the new RIM value of the target Realm is calculated by the RMM as follows:

1. Allocate a zero-filled [RmiRecParams](#) data structure to hold the measured REC parameters.
2. Copy the following attributes from the Host-provided [RmiRecParams](#) data structure into the measured REC parameters data structure:
 - gprs
 - pc
 - flags
3. Using the RHA of the target Realm, compute the hash of the measured REC parameters data structure.
4. Allocate an [RmmMeasurementDescriptorRec](#) data structure.
5. Populate the measurement descriptor:
 - Set the desc_type field to the descriptor type.
 - Set the len field to the descriptor length.
 - Set the rim field to the current RIM value of the target Realm.
 - Set the content field to the hash of the measured REC parameters.
6. Using the RHA of the target Realm, compute the hash of the measurement descriptor. Set the RIM of the target Realm to this value, zero filling upper bytes if the RHA output is smaller than the size of the RIM.

See also:

- [A7.1.1 Realm Initial Measurement](#)
- [B3.68 RimExtendRec function](#)
- [B4.4.39 RmiRecParams type](#)
- [C2.16 RmmMeasurementDescriptorRec type](#)

B4.3.28.5 Footprint

ID	Value
rec_index	<code>realm.rec_index</code>
rec_state	<code>GranuleAt(rec).state</code>
rec_aux_state	<code>AuxStates(rec.aux, RecAuxCount(rd))</code>
rim	<code>realm.measurements[0]</code>
num_recs	<code>realm.num_recs</code>

B4.3.29 RMI_REC_DESTROY command

Destroys a REC.

See also:

- [A2.3 Realm Execution Context](#)
- [B4.3.28 RMI_REC_CREATE command](#)
- [D1.2.5 Realm destruction flow](#)

B4.3.29.1 Interface

B4.3.29.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400015B
rec_ptr	X1	63:0	Address	PA of the target REC

B4.3.29.1.2 Context

The RMI_REC_DESTROY command operates on the following context.

Name	Type	Value	Before	Description
rd_pre	Address	RecAt (rec_ptr).owner	true	RD address
realm_pre	RmmRealm	RealmAt (rd_pre)	true	Realm
realm	RmmRealm	RealmAt (rd_pre)	false	Realm
rec_pre	RmmRec	RecAt (rec_ptr)	true	REC
rec	RmmRec	RecAt (rec_ptr)	false	REC

B4.3.29.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.29.2 Failure conditions

ID	Condition
rec_align	pre: !AddrIsGranuleAligned (rec_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
rec_bound	pre: !PaIsDelegable (rec_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
rec_gran_state	pre: GranuleAt (rec_ptr).state != REC post: ResultEqual (result, RMI_ERROR_INPUT)
rec_state	pre: rec.state == REC_RUNNING post: ResultEqual (result, RMI_ERROR_REC)

B4.3.29.2.1 Failure condition ordering

[rec_bound, rec_gran_state] < [rec_state]



B4.3.29.3 Success conditions

ID	Condition
rec_gran_state	<code>GranuleAt(rec_ptr).state == DELEGATED</code>
rec_aux_state	<code>AuxStateEqual16(rec_pre.aux, RecAuxCount(rd_pre), DELEGATED)</code>
num_recs	<code>realm.num_recs == realm_pre.num_recs - 1</code>

B4.3.29.4 Footprint

ID	Value
rec_state	<code>GranuleAt(rec_ptr).state</code>
rec_aux_state	<code>AuxStates(rec_pre.aux, RecAuxCount(rd_pre))</code>
num_recs	<code>realm.num_recs</code>

B4.3.30 RMI_REC_ENTER command

Enter a REC.

See also:

- [A2.3 Realm Execution Context](#)
- [Chapter A4 Realm exception model](#)
- [D1.3.1 Realm entry and exit flow](#)

B4.3.30.1 Interface

B4.3.30.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400015C
rec_ptr	X1	63:0	Address	PA of the target REC
run_ptr	X2	63:0	Address	PA of RecRun object

The number of GICv3 List Register values which can be provided by the Host in RmiRecEnter, and which are returned in RmiRecExit, is reported by the RMI_FEATURES command.

See also:

- [A3.13 GICv3 virtualization](#)

B4.3.30.1.2 Context

The RMI_REC_ENTER command operates on the following context.

Name	Type	Value	Before	Description
run	RmiRecRun	RmiRecRunAt (run_ptr)	false	RecRun object
rec	RmmRec	RecAt (rec_ptr)	false	REC
realm	RmmRealm	RealmAt (rec.owner)	false	Realm

B4.3.30.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.30.2 Failure conditions

ID	Condition
run_align	pre: !AddrIsGranuleAligned (run_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
run_bound	pre: !PaIsDelegable (run_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
run_pas	pre: !GranuleAccessPermitted(run_ptr, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
rec_align	pre: !AddrIsGranuleAligned(rec_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
rec_bound	pre: !PaIsDelegable(rec_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
rec_gran_state	pre: GranuleAt(rec_ptr).state != REC post: ResultEqual(result, RMI_ERROR_INPUT)
realm_new	pre: realm.state == REALM_NEW post: ResultEqual(result, RMI_ERROR_REALM, 0)
system_off	pre: realm.state == REALM_SYSTEM_OFF post: ResultEqual(result, RMI_ERROR_REALM, 1)
rec_state	pre: rec.state == REC_RUNNING post: ResultEqual(result, RMI_ERROR_REC)
rec_runnable	pre: rec.flags.runnable == NOT_RUNNABLE post: ResultEqual(result, RMI_ERROR_REC)
rec_mmio	pre: (run.enter.flags.emul_mmio == RMI_EMULATED_MMIO && rec.emulatable_abort != EMULATABLE_ABORT) post: ResultEqual(result, RMI_ERROR_REC)
rec_gicv3	pre: !Gicv3ConfigIsValid(run.enter.gicv3_hcr, run.enter.gicv3_lrs) post: ResultEqual(result, RMI_ERROR_REC)
rec_pending	pre: rec.pending != REC_PENDING_NONE post: ResultEqual(result, RMI_ERROR_REC)

B4.3.30.2.1 Failure condition ordering

```
[rec_align, rec_bound, rec_gran_state, run_pas, run_bound, run_align]
< [rec_state, rec_runnable, rec_mmio, realm_new, system_off,
rec_gicv3, rec_pending]
```



B4.3.30.3 Success conditions

ID	Condition
rec_exit	run.exit contains Realm exit syndrome information.
rec_emul_abt	rec.emulatable_abort is updated.

B4.3.30.4 Footprint

ID	Value
emul_abt	rec.emulatable_abort

DRAFT

B4.3.31 RMI_RTT_AUX_CREATE command

Creates an auxiliary RTT.

See also:

- [A10.3.1 Auxiliary RTT](#)
- [B4.3.32 RMI_RTT_AUX_DESTROY command](#)
- [B4.3.33 RMI_RTT_AUX_FOLD command](#)

B4.3.31.1 Interface

B4.3.31.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400017D
rd	X1	63:0	Address	PA of the RD for the target Realm
rtt	X2	63:0	Address	PA of the target RTT
ipa	X3	63:0	Address	Base of the IPA range described by the RTT
level	X4	63:0	Int64	RTT level
index	X5	63:0	UInt64	RTT tree index

B4.3.31.1.2 Context

The RMI_RTT_AUX_CREATE command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, level - 1, index)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
unfold	RmmRttEntry	RttWalk (rd, ipa, level - 1, index).rtte	true	RTTE before command execution

B4.3.31.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.31.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
level_bound	pre: (!RttLevelIsValid(rd, level) RttLevelIsStarting(rd, level)) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsRttLevelAligned(ipa, level - 1) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: UInt(ipa) >= (2 ^ realm.ipa_width) post: ResultEqual(result, RMI_ERROR_INPUT)
index_bound	pre: (realm.rtt_tree_pp == FEATURE_FALSE index == RMM_RTT_TREE_PRIMARY index > realm.num_aux_planes) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_align	pre: !AddrIsGranuleAligned(rtt) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_bound	pre: !PaIsDelegable(rtt) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_state	pre: GranuleAt(rtt).state != DELEGATED post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_bound2	pre: ((realm.feat_lpa2 == FEATURE_FALSE) && (UInt(rtt) >= 2^48)) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < level - 1 post: ResultEqual(result, RMI_ERROR_RTT_AUX, walk.level)
rtte_state	pre: walk.rtte.state == TABLE post: ResultEqual(result, RMI_ERROR_RTT_AUX, walk.level)

B4.3.31.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state]
[level_bound, ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.31.3 Success conditions

ID	Condition
rtt_state	<code>GranuleAt(rtt).state == RTT</code>
rtte_state	<code>walk.rtte.state == TABLE</code>
rtte_addr	<code>walk.rtte.addr == rtt</code>
rtte_c_ripas	pre: <code>AddrIsProtected(ipa, realm)</code> post: <code>RttAllEntriesRipas(RttAt(rtt), unfold.ripas)</code>
rtte_c_state	<code>RttAllEntriesState(RttAt(rtt), unfold.state)</code>
rtte_c_addr	pre: <code>(unfold.state != UNASSIGNED</code> <code>&& unfold.state != UNASSIGNED_NS)</code> post: <code>RttAllEntriesContiguous(RttAt(rtt), unfold.addr, level)</code>

B4.3.31.4 Footprint

ID	Value
rtt_state	<code>GranuleAt(rtt).state</code>
rtte	<code>RttEntry(walk.rtt_addr, entry_idx)</code>

B4.3.32 RMI_RTT_AUX_DESTROY command

Destroys an auxiliary RTT.

See also:

- [A10.3.1 Auxiliary RTT](#)
- [B4.3.31 RMI_RTT_AUX_CREATE command](#)
- [B4.3.33 RMI_RTT_AUX_FOLD command](#)

B4.3.32.1 Interface

B4.3.32.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400017E
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	Base of the IPA range described by the RTT
level	X3	63:0	Int64	RTT level
index	X4	63:0	UInt64	RTT tree index

B4.3.32.1.2 Context

The RMI_RTT_AUX_DESTROY command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, level - 1, index)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
walk_top	Address	RttSkipNonLiveEntries (RttAt (walk.rtt_addr), walk.level, ipa)	false	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

B4.3.32.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
rtt	X1	63:0	Address	PA of the RTT which was destroyed
top	X2	63:0	Address	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

B4.3.32.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
level_bound	pre: (!RttLevelIsValid(rd, level) RttLevelIsStarting(rd, level)) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsRttLevelAligned(ipa, level - 1) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: UInt(ipa) >= (2 ^ realm.ipa_width) post: ResultEqual(result, RMI_ERROR_INPUT)
index_bound	pre: (realm.rtt_tree_pp == FEATURE_FALSE index == RMM_RTT_TREE_PRIMARY index > realm.num_aux_planes) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < level - 1 post: (ResultEqual(result, RMI_ERROR_RTT_AUX, walk.level) && (top == walk_top))
rtte_state	pre: walk.rtte.state != TABLE post: (ResultEqual(result, RMI_ERROR_RTT_AUX, walk.level) && (top == walk_top))
rtt_live	pre: RttIsLive(RttAt(walk.rtte.addr)) post: (ResultEqual(result, RMI_ERROR_RTT_AUX, level) && (top == ipa))

B4.3.32.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state, rtt_live]
[level_bound, ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.32.3 Success conditions

ID	Condition
rtte_state	walk.rtte.state == <code>AUX_DESTROYED</code>
ripas	walk.rtte.ripas == <code>DESTROYED</code>
rtt_state	<code>GranuleAt</code> (walk.rtte.addr).state == <code>DELEGATED</code>
rtt	rtt == walk.rtte.addr
top	top == walk_top

B4.3.32.4 Footprint

ID	Value
rtt_state	<code>GranuleAt</code> (walk.rtte.addr).state
rtte	<code>RttEntry</code> (walk.rtt_addr, entry_idx)

DRAFT

B4.3.33 RMI_RTT_AUX_FOLD command

Destroys a homogeneous auxiliary RTT.

See also:

- [A5.5.6 RTT folding](#)
- [A10.3.1 Auxiliary RTT](#)
- [B4.3.31 RMI_RTT_AUX_CREATE command](#)
- [B4.3.32 RMI_RTT_AUX_DESTROY command](#)

B4.3.33.1 Interface

B4.3.33.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400017F
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	Base of the IPA range described by the RTT
level	X3	63:0	Int64	RTT level
index	X4	63:0	UInt64	RTT tree index

B4.3.33.1.2 Context

The RMI_RTT_AUX_FOLD command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, level - 1, index)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
fold_pre	RmmRttEntry	RttFold (RttAt (walk.rtte.addr))	true	Result of folding RTT

B4.3.33.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
rtt	X1	63:0	Address	PA of the RTT which was destroyed

B4.3.33.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
level_bound	pre: (!RttLevelIsValid(rd, level) RttLevelIsStarting(rd, level)) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsRttLevelAligned(ipa, level - 1) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: UInt(ipa) >= (2 ^ realm.ipa_width) post: ResultEqual(result, RMI_ERROR_INPUT)
index_bound	pre: (realm.rtt_tree_pp == FEATURE_FALSE index == RMM_RTT_TREE_PRIMARY index > realm.num_aux_planes) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < level - 1 post: ResultEqual(result, RMI_ERROR_RTT_AUX, walk.level)
rtte_state	pre: walk.rtte.state != TABLE post: ResultEqual(result, RMI_ERROR_RTT_AUX, walk.level)
rtt_homo	pre: !RttIsHomogeneous(RttAt(walk.rtte.addr)) post: ResultEqual(result, RMI_ERROR_RTT_AUX, level)

B4.3.33.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state, rtt_homo]
[level_bound, ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.33.3 Success conditions

ID	Condition
rtte_state	walk.rtte.state == fold_pre.state
rtte_addr	pre: (fold_pre.state != UNASSIGNED && fold_pre.state != UNASSIGNED_NS) post: walk.rtte.addr == fold_pre.addr

ID	Condition
rtte_attr	pre: (fold_pre.state == ASSIGNED fold_pre.state == ASSIGNED_NS) post: (walk.rtte.MemAttr == fold_pre.MemAttr && walk.rtte.s2ap_base == fold_pre.s2ap_base && walk.rtte.s2ap_overlay == fold_pre.s2ap_overlay)
rtte_ripas	pre: AddrIsProtected(ipa, realm) post: walk.rtte.ripas == fold_pre.ripas
rtt_state	GranuleAt(walk.rtte.addr).state == DELEGATED
rtt	rtt == walk.rtte.addr

B4.3.33.4 Footprint

ID	Value
rtt_state	GranuleAt(walk.rtte.addr).state
rtte	RttEntry(walk.rtt_addr, entry_idx)

DRAFT

B4.3.34 RMI_RTT_AUX_MAP_PROTECTED command

Creates a mapping from an Protected IPA in an auxiliary RTT.

See also:

- [A10.3.1 Auxiliary RTT](#)
- [B4.3.36 RMI_RTT_AUX_UNMAP_PROTECTED command](#)

B4.3.34.1 Interface

B4.3.34.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000180
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	IPA in the target Realm
index	X3	63:0	UInt64	RTT tree index

B4.3.34.1.2 Context

The RMI_RTT_AUX_MAP_PROTECTED command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk_pri	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL , RMM_RTT_TREE_PRIMARY)	false	Primary RTT walk result
walk_aux	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL , index)	false	Auxiliary RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk_aux.level)	false	RTTE index

B4.3.34.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
fail_index	X1	63:0	UInt64	Index of RTT tree whose contents caused command to fail with RMI_ERROR_RTT
level_pri	X2	63:0	Int64	Level of RTTE reached by walk of primary RTT tree

Name	Register	Bits	Type	Description
state	X3	7:0	RmiRttEntryState	State of RTT entry whose contents caused command to fail with RMI_ERROR_RTT
ripas	X4	7:0	RmiRipas	RIPAS of RTT entry which caused command to fail with RMI_ERROR_RTT

The following unused bits of RMI_RTT_AUX_MAP_PROTECTED output values MBZ: X3[63:8], X4[63:8].

B4.3.34.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt (rd).state != RD post: ResultEqual (result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsGranuleAligned (ipa) post: ResultEqual (result, RMI_ERROR_INPUT)
ipa_bound	pre: !AddrIsProtected (ipa, realm) post: ResultEqual (result, RMI_ERROR_INPUT)
index_bound	pre: (realm.rtt_tree_pp == FEATURE_FALSE index == RMM_RTT_TREE_PRIMARY index > realm.num_aux_planes) post: ResultEqual (result, RMI_ERROR_INPUT)
pri_state	pre: walk_pri.rtte.state != ASSIGNED post: (ResultEqual (result, RMI_ERROR_RTT , walk_pri.level) && (fail_index == RMM_RTT_TREE_PRIMARY) && (level_pri == walk_pri.level) && (state == RttEntryStateToRmi (walk_pri.rtte.state)) && (ripas == RipasToRmi (walk_pri.rtte.ripas)))
aux_destroyed	pre: walk_aux.rtte.state == AUX_DESTROYED post: (ResultEqual (result, RMI_ERROR_RTT_AUX , walk_aux.level) && (fail_index == index) && (level_pri == walk_pri.level) && (state == RttEntryStateToRmi (walk_aux.rtte.state)) && (ripas == RipasToRmi (walk_pri.rtte.ripas)))

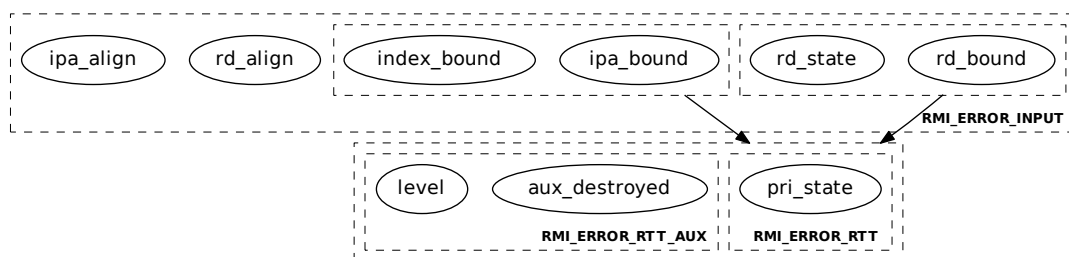
ID	Condition
level	<pre> pre: walk_aux.level < walk_pri.level post: (ResultEqual(result, RMI_ERROR_RTT_AUX, walk_aux.level) && (fail_index == index) && (level_pri == walk_pri.level) && (state == RttEntryStateToRmi(walk_aux.rtte.state)) && (ripas == RipasToRmi(walk_pri.rtte.ripas))) </pre>

B4.3.34.2.1 Failure condition ordering

```

[rd_bound, rd_state] < [pri_state, aux_destroyed, level]
[ipa_bound, index_bound] < [pri_state, aux_destroyed, level]

```



B4.3.34.3 Success conditions

ID	Condition
rtte_state	walk_aux.rtte.state == ASSIGNED
rtte_attr	walk_aux.rtte.MemAttr == walk_pri.rtte.MemAttr
rtte_addr	<pre> walk_aux.rtte.addr == walk_pri.rtte.addr + (entry_idx * RttLevelSize(walk_aux.level)) </pre>

B4.3.34.4 Footprint

ID	Value
rtte	RttEntry(walk_aux.rtt_addr, entry_idx)

B4.3.35 RMI_RTT_AUX_MAP_UNPROTECTED command

Creates a mapping from an Unprotected IPA in an auxiliary RTT.

See also:

- [A10.3.1 Auxiliary RTT](#)
- [B4.3.37 RMI_RTT_AUX_UNMAP_UNPROTECTED command](#)

B4.3.35.1 Interface

B4.3.35.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000181
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	IPA in the target Realm
index	X3	63:0	UInt64	RTT tree index

B4.3.35.1.2 Context

The RMI_RTT_AUX_MAP_UNPROTECTED command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk_pri	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL , RMM_RTT_TREE_PRIMARY)	false	Primary RTT walk result
walk_aux	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL , index)	false	Auxiliary RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk_aux.level)	false	RTTE index

B4.3.35.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
fail_index	X1	63:0	UInt64	Index of RTT tree whose contents caused command to fail with RMI_ERROR_RTT
level_pri	X2	63:0	Int64	Level of RTTE reached by walk of primary RTT tree

Name	Register	Bits	Type	Description
state	X3	7:0	RmiRttEntryState	State of RTT entry whose contents caused command to fail with RMI_ERROR_RTT

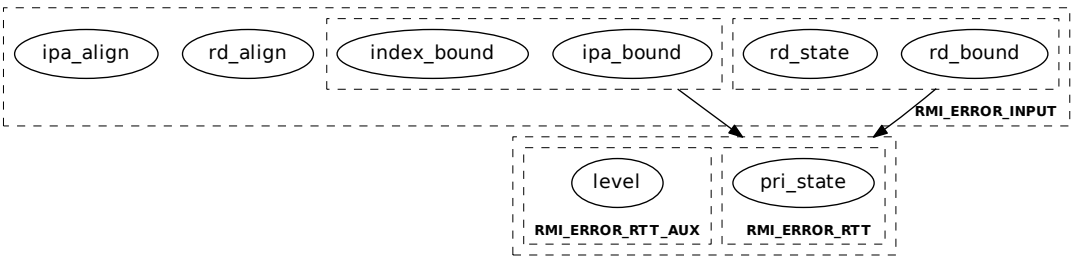
The following unused bits of RMI_RTT_AUX_MAP_UNPROTECTED output values MBZ: X3[63:8].

B4.3.35.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt (rd).state != RD post: ResultEqual (result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsGranuleAligned (ipa) post: ResultEqual (result, RMI_ERROR_INPUT)
ipa_bound	pre: (UInt(ipa) >= (2 ^ realm.ipa_width) AddrIsProtected (ipa, realm)) post: ResultEqual (result, RMI_ERROR_INPUT)
index_bound	pre: (realm.rtt_tree_pp == FEATURE_FALSE index == RMM_RTT_TREE_PRIMARY index > realm.num_aux_planes) post: ResultEqual (result, RMI_ERROR_INPUT)
pri_state	pre: walk_pri.rtte.state != ASSIGNED_NS post: (ResultEqual (result, RMI_ERROR_RTT , walk_pri.level) && (fail_index == RMM_RTT_TREE_PRIMARY) && (state == RttEntryStateToRmi (walk_pri.rtte.state)) && (level_pri == walk_pri.level))
level	pre: walk_aux.level < walk_pri.level post: (ResultEqual (result, RMI_ERROR_RTT_AUX , walk_aux.level) && (fail_index == index) && (level_pri == walk_pri.level) && (state == RttEntryStateToRmi (walk_aux.rtte.state)))

B4.3.35.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [pri_state, level]
[ipa_bound, index_bound] < [pri_state, level]
```



B4.3.35.3 Success conditions

ID	Condition
rtte_state	walk_aux.rtte.state == ASSIGNED_NS
rtte_attr	(walk_aux.rtte.MemAttr == walk_pri.rtte.MemAttr && walk_aux.rtte.s2ap_base == walk_pri.rtte.s2ap_base && walk_aux.rtte.s2ap_overlay == walk_pri.rtte.s2ap_overlay)
rtte_addr	walk_aux.rtte.addr == walk_pri.rtte.addr + (entry_idx * RttLevelSize(walk_aux.level))

B4.3.35.4 Footprint

ID	Value
rtte	RttEntry(walk_aux.rtt_addr, entry_idx)

B4.3.36 RMI_RTT_AUX_UNMAP_PROTECTED command

Removes a mapping from an Protected IPA in an auxiliary RTT.

See also:

- [A10.3.1 Auxiliary RTT](#)
- [B4.3.34 RMI_RTT_AUX_MAP_PROTECTED command](#)

B4.3.36.1 Interface

B4.3.36.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000183
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	IPA in the target Realm
index	X3	63:0	UInt64	RTT tree index

B4.3.36.1.2 Context

The RMI_RTT_AUX_UNMAP_PROTECTED command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL , index)	false	Auxiliary RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
walk_top	Address	RttSkipNonLiveEntries (RttAt (walk.rtt_addr), walk.level, ipa)	false	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

B4.3.36.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
top	X1	63:0	Address	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated
level	X2	63:0	Int64	RTT level reached by the RTT walk

B4.3.36.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsGranuleAligned(ipa) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: !AddrIsProtected(ipa, realm) post: ResultEqual(result, RMI_ERROR_INPUT)
index_bound	pre: (realm.rtt_tree_pp == FEATURE_FALSE index == RMM_RTT_TREE_PRIMARY index > realm.num_aux_planes) post: ResultEqual(result, RMI_ERROR_INPUT)
rtte_state	pre: walk.rtte.state != ASSIGNED post: (ResultEqual(result, RMI_ERROR_RTT_AUX, walk.level) && (top == walk_top))

B4.3.36.2.1 Failure condition ordering

[rd_bound, rd_state] < [rtte_state]
[ipa_bound, index_bound] < [rtte_state]



B4.3.36.3 Success conditions

ID	Condition
rtte_state	walk.rtte.state == UNASSIGNED
top	top == walk_top
level	level == walk.level

B4.3.36.4 Footprint

ID	Value
rtte	<code>RttEntry(walk.rtt_addr, entry_idx)</code>

DRAFT

B4.3.37 RMI_RTT_AUX_UNMAP_UNPROTECTED command

Removes a mapping from an Unprotected IPA in an auxiliary RTT.

See also:

- [A10.3.1 Auxiliary RTT](#)
- [B4.3.34 RMI_RTT_AUX_MAP_PROTECTED command](#)

B4.3.37.1 Interface

B4.3.37.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000184
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	IPA in the target Realm
index	X3	63:0	UInt64	RTT tree index

B4.3.37.1.2 Context

The RMI_RTT_AUX_UNMAP_UNPROTECTED command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, RMM_RTT_PAGE_LEVEL, index)	false	Auxiliary RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
walk_top	Address	RttSkipNonLiveEntries (RttAt (walk.rtt_addr), walk.level, ipa)	false	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

B4.3.37.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
top	X1	63:0	Address	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated
level	X2	63:0	Int64	RTT level reached by the RTT walk

B4.3.37.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsGranuleAligned(ipa) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: (UInt(ipa) >= (2 ^ realm.ipa_width) AddrIsProtected(ipa, realm)) post: ResultEqual(result, RMI_ERROR_INPUT)
index_bound	pre: (realm.rtt_tree_pp == FEATURE_FALSE index == RMM_RTT_TREE_PRIMARY index > realm.num_aux_planes) post: ResultEqual(result, RMI_ERROR_INPUT)
rtte_state	pre: walk.rtte.state != ASSIGNED_NS post: (ResultEqual(result, RMI_ERROR_RTT_AUX, walk.level) && (top == walk_top))

B4.3.37.2.1 Failure condition ordering

[rd_bound, rd_state] < [rtte_state]
[ipa_bound, index_bound] < [rtte_state]



B4.3.37.3 Success conditions

ID	Condition
rtte_state	walk.rtte.state == UNASSIGNED_NS
top	top == walk_top
level	level == walk.level

B4.3.37.4 Footprint

ID	Value
rtte	<code>RttEntry(walk.rtt_addr, entry_idx)</code>

DRAFT

B4.3.38 RMI_RTT_CREATE command

Creates a primary RTT.

See also:

- [A5.5 Realm Translation Table](#)
- [A5.5.7 RTT unfolding](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)
- [B4.3.40 RMI_RTT_FOLD command](#)

B4.3.38.1 Interface

B4.3.38.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400015D
rd	X1	63:0	Address	PA of the RD for the target Realm
rtt	X2	63:0	Address	PA of the target RTT
ipa	X3	63:0	Address	Base of the IPA range described by the RTT
level	X4	63:0	Int64	RTT level

B4.3.38.1.2 Context

The RMI_RTT_CREATE command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, level - 1, RMM_RTT_TREE_PRIMARY)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
walk_pre	RmmRttWalkResult	RttWalk (rd, ipa, level - 1, RMM_RTT_TREE_PRIMARY)	true	RTT walk result before command execution
rtte_pre	RmmRttEntry	walk_pre.rtte	true	RTTE before command execution

B4.3.38.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.38.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
level_bound	pre: (!RttLevelIsValid(rd, level) RttLevelIsStarting(rd, level)) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsRttLevelAligned(ipa, level - 1) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: UInt(ipa) >= (2 ^ realm.ipa_width) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_align	pre: !AddrIsGranuleAligned(rtt) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_bound	pre: !PaIsDelegable(rtt) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_state	pre: GranuleAt(rtt).state != DELEGATED post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_bound2	pre: ((realm.feat_lpa2 == FEATURE_FALSE) && (UInt(rtt) >= 2^48)) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < level - 1 post: ResultEqual(result, RMI_ERROR_RTT, walk.level)
rtte_state	pre: walk.rtte.state == TABLE post: ResultEqual(result, RMI_ERROR_RTT, walk.level)

B4.3.38.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state]
[level_bound, ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.38.3 Success conditions

ID	Condition
rtt_state	GranuleAt(rtt).state == RTT
rtte_state	walk.rtte.state == TABLE

ID	Condition
rtte_addr	walk.rtte.addr == rtt
rtte_c_ripas	pre: AddrIsProtected(ipa, realm) post: RttAllEntriesRipas(RttAt(rtt), rtte_pre.ripas)
rtte_c_state	RttAllEntriesState(RttAt(rtt), rtte_pre.state)
rtte_c_addr	pre: (rtte_pre.state != UNASSIGNED && rtte_pre.state != UNASSIGNED_NS) post: RttAllEntriesContiguous(RttAt(rtt), rtte_pre.addr, level)

B4.3.38.4 Footprint

ID	Value
rtt_state	GranuleAt(rtt).state
rtte	RttEntryAt(walk.rtt_addr, entry_idx)

DRAFT

B4.3.39 RMI_RTT_DESTROY command

Destroys a primary RTT.

See also:

- [A5.5 Realm Translation Table](#)
- [A5.5.9 RTT destruction](#)
- [B4.3.38 RMI_RTT_CREATE command](#)
- [B4.3.40 RMI_RTT_FOLD command](#)

B4.3.39.1 Interface

B4.3.39.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400015E
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	Base of the IPA range described by the RTT
level	X3	63:0	Int64	RTT level

B4.3.39.1.2 Context

The RMI_RTT_DESTROY command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, level - 1, RMM_RTT_TREE_PRIMARY)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
walk_top	Address	RttSkipNonLiveEntries (RttAt (walk.rtt_addr), walk.level, ipa)	false	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

B4.3.39.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
rtt	X1	63:0	Address	PA of the RTT which was destroyed
top	X2	63:0	Address	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

The `rtt` output value is valid only when the command result is `RMI_SUCCESS`.

The values of the `result` and `top` output values for different command outcomes are summarized in the following table.

Scenario	result	top	walk.rtte.state
Target RTT exists and is not live	<code>RMI_SUCCESS</code>	<code>> ipa</code>	Before execution: <code>TABLE</code> After execution: <code>UNASSIGNED</code> and <code>RIPAS</code> is <code>DESTROYED</code>
Missing RTT	<code>(RMI_ERROR_RTT, < level)</code>	<code>> ipa</code>	<code>UNASSIGNED</code> or <code>UNASSIGNED_NS</code>
Block mapping at lower level	<code>(RMI_ERROR_RTT, < level)</code>	<code>== ipa</code>	<code>ASSIGNED</code> or <code>ASSIGNED_NS</code>
Live RTT at target level	<code>(RMI_ERROR_RTT, level)</code>	<code>== ipa</code>	<code>TABLE</code>
RTT walk was not performed, due to any other command failure	Another error code	<code>0</code>	Unknown

See also:

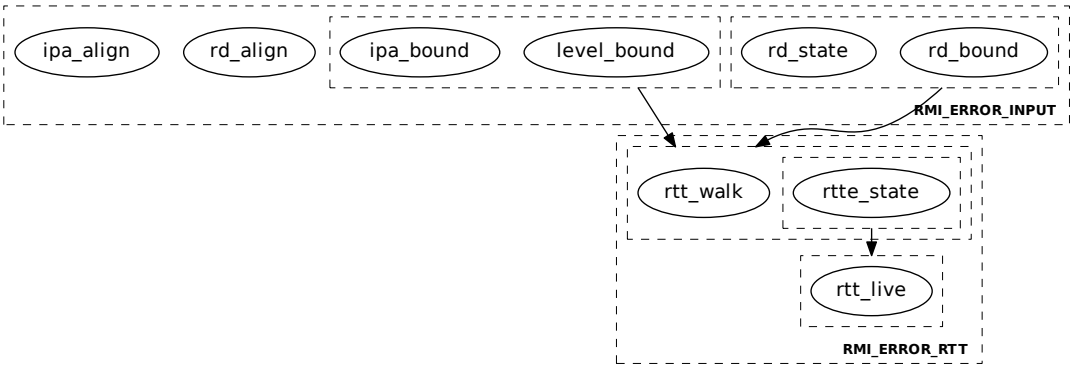
- [A5.5.8 RTTE liveness and RTT liveness](#)

B4.3.39.2 Failure conditions

ID	Condition
<code>rd_align</code>	pre: <code>!AddrIsGranuleAligned(rd)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
<code>rd_bound</code>	pre: <code>!PaIsDelegable(rd)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
<code>rd_state</code>	pre: <code>GranuleAt(rd).state != RD</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
<code>level_bound</code>	pre: <code>(!RttLevelIsValid(rd, level) RttLevelIsStarting(rd, level))</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
<code>ipa_align</code>	pre: <code>!AddrIsRttLevelAligned(ipa, level - 1)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
<code>ipa_bound</code>	pre: <code>UInt(ipa) >= (2 ^ realm.ipa_width)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
<code>rtt_walk</code>	pre: <code>walk.level < level - 1</code> post: <code>(ResultEqual(result, RMI_ERROR_RTT, walk.level) && (top == walk_top))</code>
<code>rtte_state</code>	pre: <code>walk.rtte.state != TABLE</code> post: <code>(ResultEqual(result, RMI_ERROR_RTT, walk.level) && (top == walk_top))</code>
<code>rtt_live</code>	pre: <code>RttIsLive(RttAt(walk.rtte.addr))</code> post: <code>(ResultEqual(result, RMI_ERROR_RTT, level) && (top == ipa))</code>

B4.3.39.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state]
[rtte_state] < [rtt_live]
[level_bound, ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.39.3 Success conditions

ID	Condition
rtte_state	walk.rtte.state == UNASSIGNED
ripas	walk.rtte.ripas == DESTROYED
rtt_state	GranuleAt(walk.rtte.addr).state == DELEGATED
rtt	rtt == walk.rtte.addr
top	top == walk_top

B4.3.39.4 Footprint

ID	Value
rtt_state	GranuleAt(walk.rtte.addr).state
rtte	RttEntryAt(walk.rtt_addr, entry_idx)

B4.3.40 RMI_RTT_FOLD command

Destroys a homogeneous primary RTT.

See also:

- [A5.5 Realm Translation Table](#)
- [A5.5.6 RTT folding](#)
- [B4.3.38 RMI_RTT_CREATE command](#)
- [B4.3.39 RMI_RTT_DESTROY command](#)

B4.3.40.1 Interface

B4.3.40.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000166
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	Base of the IPA range described by the RTT
level	X3	63:0	Int64	RTT level

B4.3.40.1.2 Context

The RMI_RTT_FOLD command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, level - 1, RMM_RTT_TREE_PRIMARY)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
fold_pre	RmmRttEntry	RttFold (RttAt (walk.rtte.addr))	true	Result of folding RTT

B4.3.40.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
rtt	X1	63:0	Address	PA of the RTT which was destroyed

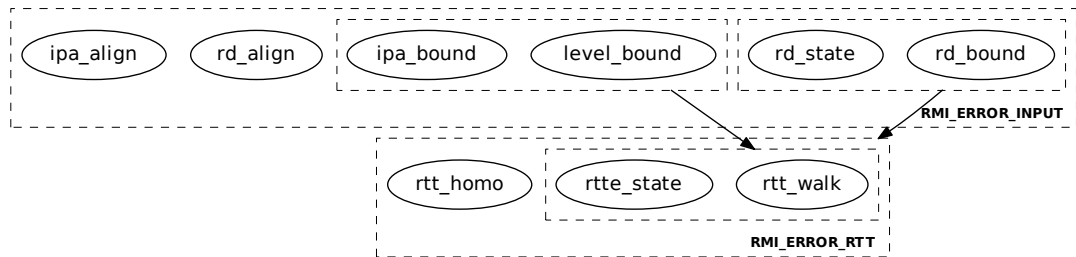
The `rtt` output value is valid only when the command result is `RMI_SUCCESS`.

B4.3.40.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
level_bound	pre: (!RttLevelIsValid(rd, level) RttLevelIsStarting(rd, level)) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsRttLevelAligned(ipa, level - 1) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: UInt(ipa) >= (2 ^ realm.ipa_width) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < level - 1 post: ResultEqual(result, RMI_ERROR_RTT, walk.level)
rtte_state	pre: walk.rtte.state != TABLE post: ResultEqual(result, RMI_ERROR_RTT, walk.level)
rtt_homo	pre: !RttIsHomogeneous(RttAt(walk.rtte.addr)) post: ResultEqual(result, RMI_ERROR_RTT, level)

B4.3.40.2.1 Failure condition ordering

[rd_bound, rd_state] < [rtt_walk, rtte_state, rtt_homo]
[level_bound, ipa_bound] < [rtt_walk, rtte_state]



B4.3.40.3 Success conditions

ID	Condition
rtte_state	walk.rtte.state == fold_pre.state
rtte_addr	pre: (fold_pre.state != UNASSIGNED && fold_pre.state != UNASSIGNED_NS) post: walk.rtte.addr == fold_pre.addr

ID	Condition
rtte_attr	pre: (fold_pre.state == ASSIGNED fold_pre.state == ASSIGNED_NS) post: (walk.rtte.MemAttr == fold_pre.MemAttr && walk.rtte.s2ap_base == fold_pre.s2ap_base && walk.rtte.s2ap_overlay == fold_pre.s2ap_overlay)
rtte_ripas	pre: AddrIsProtected(ipa, realm) post: walk.rtte.ripas == fold_pre.ripas
rtt_state	GranuleAt(walk.rtte.addr).state == DELEGATED
rtt	rtt == walk.rtte.addr

B4.3.40.4 Footprint

ID	Value
rtt_state	GranuleAt(walk.rtte.addr).state
rtte	RttEntryAt(walk.rtt_addr, entry_idx)

B4.3.41 RMI_RTT_INIT_RIPAS command

Set the RIPAS of a target IPA range to RAM, for a Realm in the REALM_NEW state.

See also:

- [A5.2.2 Realm IPA state](#)
- [D1.2.3 Initialize memory of New Realm flow](#)

B4.3.41.1 Interface

B4.3.41.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000168
rd	X1	63:0	Address	PA of the RD for the target Realm
base	X2	63:0	Address	Base of target IPA region
top	X3	63:0	Address	Top of target IPA region

B4.3.41.1.2 Context

The RMI_RTT_INIT_RIPAS command operates on the following context.

Name	Type	Value	Before	Description
realm_pre	RmmRealm	RealmAt (rd)	true	Realm
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, base, RMM_RTT_PAGE_LEVEL , RMM_RTT_TREE_PRIMARY)	false	RTT walk result
walk_top	Address	RttSkipEntriesWithRipas (RttAt (walk.rtt_addr), walk.level, base, top, FALSE)	false	Top IPA of entries which have associated RIPAS values, starting from entry at which the RTT walk terminated

B4.3.41.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
out_top	X1	63:0	Address	Top IPA of range whose RIPAS was modified

The `out_top` output value is valid only when the command result is `RMI_SUCCESS`.


When the `out_top` output value is valid, it is aligned to the size of the address range described by the RTT entry at the level where the RTT walk terminated.

B4.3.41.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
size_valid	pre: UInt(top) <= UInt(base) post: ResultEqual(result, RMI_ERROR_INPUT)
top_bound	pre: !AddrIsProtected(ToAddress(UInt(top) - RMM_GRANULE_SIZE), realm_pre) post: ResultEqual(result, RMI_ERROR_INPUT)
realm_state	pre: realm_pre.state != REALM_NEW post: ResultEqual(result, RMI_ERROR_REALM)
base_align	pre: !AddrIsRttLevelAligned(base, walk.level) post: ResultEqual(result, RMI_ERROR_RTT, walk.level)
rtte_state	pre: walk.rtte.state != UNASSIGNED post: ResultEqual(result, RMI_ERROR_RTT, walk.level)
top_gran_align	pre: !AddrIsGranuleAligned(top) post: ResultEqual(result, RMI_ERROR_INPUT)
no_progress	pre: UInt(base) == UInt(walk_top) post: ResultEqual(result, RMI_ERROR_RTT, walk.level)

B4.3.41.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [realm_state]
[rd_bound, rd_state] < [base_align, rtte_state]
[rd_bound, rd_state] < [no_progress]
[top_gran_align] < [no_progress]
```



B4.3.41.3 Success conditions

ID	Condition
rtte_ripas	<code>RttEntriesInRangeRipas (</code> <code> RttAt (walk.rtt_addr),</code> <code> walk.level,</code> <code> base, walk_top,</code> <code> RAM)</code>
rim	<code>realm.measurements[0] == RimExtendRipas (</code> <code> realm_pre, base, walk_top, walk.level)</code>
out_top	<code>out_top == walk_top</code>

B4.3.41.4 RMI_RTT_INIT_RIPAS extension of RIM

On successful execution of RMI_RTT_INIT_RIPAS, the new RIM value of the target Realm is calculated by the RMM as follows:

1. Allocate an [RmmMeasurementDescriptorRipas](#) data structure.
2. For each RTT entry in the range [base, top) described by the RMI_RTT_INIT_RIPAS input values:
 - a. Populate the measurement descriptor:
 - Set the desc_type field to the descriptor type.
 - Set the len field to the descriptor length.
 - Set the base field to the IPA of the RTT entry.
 - Set the top field to `Min(ipa + size, top)`, where
 - ipa is the IPA of the RTT entry
 - size is the size in bytes of the IPA region described by the RTT entry
 - top is the input value provided to the command
 - b. Using the RHA of the target Realm, compute the hash of the measurement descriptor. Set the RIM of the target Realm to this value, zero filling upper bytes if the RHA output is smaller than the size of the RIM.

See also:

- [A7.1.1 Realm Initial Measurement](#)
- [B3.69 RimExtendRipas function](#)
- [C2.17 RmmMeasurementDescriptorRipas type](#)

B4.3.41.5 Footprint

ID	Value
rtte	<code>RttAt (walk.rtt_addr)</code>
rim	<code>realm.measurements[0]</code>

B4.3.42 RMI_RTT_MAP_UNPROTECTED command

Creates a mapping from an Unprotected IPA to a Non-secure PA in a primary RTT.

See also:

- [A5.5 Realm Translation Table](#)
- [B4.3.46 RMI_RTT_UNMAP_UNPROTECTED command](#)

B4.3.42.1 Interface

B4.3.42.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400015F
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	IPA at which the Granule will be mapped in the target Realm
level	X3	63:0	Int64	RTT level
desc	X4	63:0	Bits64	RTTE descriptor

The layout and encoding of fields in the `desc` input value match “Attribute fields in stage 2 VMSAv8-64 Block and Page descriptors” in [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#).

If the Realm has multiple Planes and is configured to use a single RTT tree then the S2AP base index value, provided in the S2AP field of the `desc` input value, uses the encoding defined in the `RmiUnprotectedS2AP` type.

See also:

- [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#)
- [A5.5.11.3 RTT entry attributes for ASSIGNED_NS mappings](#)
- [B3.101 RttDescriptorIsValidForUnprotected function](#)
- [B4.4.48 RmiUnprotectedS2AP type](#)

B4.3.42.1.2 Context

The `RMI_RTT_MAP_UNPROTECTED` command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	<code>RealmAt(rd)</code>	false	Realm
walk	RmmRttWalkResult	<code>RttWalk(rd, ipa, level, RMM_RTT_TREE_PRIMARY)</code>	false	RTT walk result
entry_idx	UInt64	<code>RttEntryIndex(ipa, walk.level)</code>	false	RTTE index
rtte	RmmRttEntry	<code>RttDescriptorDecode(desc)</code>	false	RTT entry

B4.3.42.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.42.2 Failure conditions

ID	Condition
attr_valid	pre: !RttDescriptorIsValidForUnprotected (desc) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_align	pre: !AddrIsGranuleAligned (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt (rd).state != RD post: ResultEqual (result, RMI_ERROR_INPUT)
level_bound	pre: !RttLevelIsBlockOrPage (rd, level) post: ResultEqual (result, RMI_ERROR_INPUT)
addr_align	pre: !AddrIsRttLevelAligned (rtte.addr, level) post: ResultEqual (result, RMI_ERROR_INPUT)
addr_bound	pre: ((realm.feats_lpa2 == FEATURE_FALSE) && (UInt(rtte.addr) >= 2 ⁴⁸)) post: ResultEqual (result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsRttLevelAligned (ipa, level) post: ResultEqual (result, RMI_ERROR_INPUT)
ipa_bound	pre: (UInt(ipa) >= (2 ^ realm.ipa_width) AddrIsProtected (ipa, realm)) post: ResultEqual (result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < level post: ResultEqual (result, RMI_ERROR_RTT , walk.level)
rtte_state	pre: walk.rtte.state != UNASSIGNED_NS post: ResultEqual (result, RMI_ERROR_RTT , walk.level)

B4.3.42.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state]
[level_bound, ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.42.3 Success conditions

ID	Condition
rtte_state	walk.rtte.state == ASSIGNED_NS
rtte_contents	(walk.rtte.MemAttr == rtte.MemAttr && walk.rtte.s2ap_base == rtte.s2ap_base && walk.rtte.s2ap_overlay == 15 && walk.rtte.addr == rtte.addr)

B4.3.42.4 Footprint

ID	Value
rtte	RttEntryAt(walk.rtt_addr, entry_idx)

DRAFT

B4.3.43 RMI_RTT_READ_ENTRY command

Reads an entry from a primary RTT.

See also:

- [A5.5 Realm Translation Table](#)

B4.3.43.1 Interface

B4.3.43.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000161
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	Realm Address for which to read the RTTE
level	X3	63:0	Int64	RTT level at which to read the RTTE

B4.3.43.1.2 Context

The RMI_RTT_READ_ENTRY command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, level, RMM_RTT_TREE_PRIMARY)	false	RTT walk result
rtte	RmmRttEntry	RttDescriptorDecode (desc)	false	RTT entry value returned to Host

B4.3.43.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
walk_level	X1	63:0	UInt64	RTT level reached by the RTT walk
state	X2	7:0	RmiRttEntryState	State of RTTE reached by the walk
desc	X3	63:0	Bits64	RTTE descriptor
ripas	X4	7:0	RmiRipas	RIPAS of RTTE reached by the walk

The following unused bits of RMI_RTT_READ_ENTRY output values MBZ: X2[63:8], X4[63:8].

The layout and encoding of fields in the `rtte` output value match “Attribute fields in stage 2 VMSAv8-64 Block and Page descriptors” in [Arm Architecture Reference Manual for A-Profile architecture](#) [3].

See also:

- [Arm Architecture Reference Manual for A-Profile architecture](#) [3]
- [A5.5.11 RTT entry attributes](#)

B4.3.43.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
level_bound	pre: !RttLevelIsValid(rd, level) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsRttLevelAligned(ipa, level) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: UInt(ipa) >= (2 ^ realm.ipa_width) post: ResultEqual(result, RMI_ERROR_INPUT)

B4.3.43.2.1 Failure condition ordering

The RMI_RTT_READ_ENTRY command does not have any failure condition orderings.

B4.3.43.3 Success conditions

ID	Condition
state	state == RttEntryStateToRmi(walk.rtte.state)
state_invalid	pre: (walk.rtte.state == UNASSIGNED walk.rtte.state == UNASSIGNED_NS) post: (rtte.MemAttr == Zeros(3) && rtte.s2ap_base == 0 && rtte.s2ap_overlay == 0 && rtte.addr == Zeros(ADDRESS_WIDTH))
state_prot	pre: (walk.rtte.state == ASSIGNED walk.rtte.state == TABLE) post: (rtte.MemAttr == Zeros(3) && rtte.s2ap_base == 0 && rtte.s2ap_overlay == 0 && rtte.addr == walk.rtte.addr)
state_unprot	pre: walk.rtte.state == ASSIGNED_NS post: (rtte.MemAttr == walk.rtte.MemAttr && rtte.s2ap_base == walk.rtte.s2ap_base && rtte.s2ap_overlay == 0 && rtte.addr == walk.rtte.addr)
state_io	pre: (walk.rtte.state == ASSIGNED_DEV_PRIVATE walk.rtte.state == ASSIGNED_DEV_SHARED) post: (rtte.MemAttr == walk.rtte.MemAttr && rtte.s2ap_base == 0 && rtte.s2ap_overlay == 0 && rtte.addr == walk.rtte.addr)
ripas_prot	pre: (walk.rtte.state == UNASSIGNED walk.rtte.state == ASSIGNED) post: ripas == RipasToRmi(walk.rtte.ripas)

ID	Condition
ripas_unprot	<pre>pre: (walk.rtte.state == UNASSIGNED_NS walk.rtte.state == ASSIGNED_NS) post: ripas == RMI_EMPTY</pre>

B4.3.43.4 Footprint

The RMI_RTT_READ_ENTRY command does not have any footprint.

DRAFT

B4.3.44 RMI_RTT_SET_RIPAS command

Completes a request made by the Realm to change the RIPAS of a target IPA range.

Issue In RMI_RTT_SET_RIPAS, consider how to combine:

- Modification of a range of RTT entries in a single command, and
- Checking of output address and HIPAS values against rec.ripas_dev_pa and rec.ripas_dev_shared respectively.

See also:

- [A5.4 RIPAS change](#)
- [A9.5.3 Realm validation of device memory mappings](#)

B4.3.44.1 Interface

B4.3.44.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000169
rd	X1	63:0	Address	PA of the RD for the target Realm
rec_ptr	X2	63:0	Address	PA of the target REC
base	X3	63:0	Address	Base of target IPA region
top	X4	63:0	Address	Top of target IPA region

B4.3.44.1.2 Context

The RMI_RTT_SET_RIPAS command operates on the following context.

Name	Type	Value	Before	Description
realm_pre	RmmRealm	RealmAt (rd)	true	Realm
rec	RmmRec	RecAt (rec_ptr)	false	REC
walk	RmmRttWalkResult	RttWalk (rd, base, RMM_RTT_PAGE_LEVEL, RMM_RTT_TREE_PRIMARY)	false	RTT walk result
ripas_pre	RmmRipas	walk.rtte.ripas	true	RIPAS before the command executed
walk_top_pre	Address	RttSkipEntriesWithRipas (RttAt (walk.rtt_addr), walk.level, base, top, rec.ripas_destroyed != CHANGE_DESTROYED)	true	Top IPA of entries which have associated RIPAS values, starting from entry at which the RTT walk terminated

B4.3.44.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
out_top	X1	63:0	Address	Top IPA of range whose RIPAS was modified

The `out_top` output value is valid only when the command result is `RMI_SUCCESS`.


B4.3.44.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt (rd).state != RD post: ResultEqual (result, RMI_ERROR_INPUT)
rec_align	pre: !AddrIsGranuleAligned (rec_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
rec_bound	pre: !PaIsDelegable (rec_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
rec_gran_state	pre: GranuleAt (rec_ptr).state != REC post: ResultEqual (result, RMI_ERROR_INPUT)
rec_state	pre: rec.state == REC_RUNNING post: ResultEqual (result, RMI_ERROR_REC)
rec_owner	pre: rec.owner != rd post: ResultEqual (result, RMI_ERROR_REC)
size_valid	pre: UInt (top) <= UInt (base) post: ResultEqual (result, RMI_ERROR_INPUT)
base_bound	pre: base != rec.ripas_addr post: ResultEqual (result, RMI_ERROR_INPUT)
top_bound	pre: UInt (top) > UInt (rec.ripas_top) post: ResultEqual (result, RMI_ERROR_INPUT)
base_align	pre: (!AddrIsRttLevelAligned (base, walk.level) && ripas_pre != rec.ripas_value) post: ResultEqual (result, RMI_ERROR_RTT , walk.level)
top_gran_align	pre: !AddrIsGranuleAligned (top) post: ResultEqual (result, RMI_ERROR_INPUT)
no_progress	pre: (UInt (base) == UInt (walk_top_pre) && ripas_pre != rec.ripas_value) post: ResultEqual (result, RMI_ERROR_RTT , walk.level)
dev_pa	pre: Output address does not match value required by rec.ripas_dev_pa. post: ResultEqual (result, RMI_ERROR_RTT , walk.level)
dev_shared	pre: HIPAS value does not match value required by rec.ripas_dev_shared. post: ResultEqual (result, RMI_ERROR_RTT , walk.level)

ID	Condition
aux_live	<pre>pre: AddrRangeIsAuxLive(base, top, realm_pre) post: ResultEqual(result, RMI_ERROR_RTT, walk.level)</pre>

B4.3.44.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [base_align]
[rd_bound, rd_state] < [no_progress]
[rec_bound, rec_gran_state] < [rec_state, rec_owner]
[base_bound] < [base_align]
[top_gran_align] < [no_progress]
```



B4.3.44.3 Success conditions

ID	Condition
rtte_ripas	<code>RttEntriesInRangeRipas (RttAt (walk.rtt_addr), walk.level, base, walk_top_pre, rec.ripas_value)</code>
ripas_addr	<code>rec.ripas_addr == MinAddress (top, walk_top_pre)</code>
out_top	<code>out_top == MinAddress (top, walk_top_pre)</code>

B4.3.44.4 Footprint

ID	Value
rtte	<code>RttAt (walk.rtt_addr)</code>
ripas_addr	<code>rec.ripas_addr</code>

B4.3.45 RMI_RTT_SET_S2AP command

Completes a request made by the Realm to change the S2AP of a target IPA range.

See also:

- [A10.3.2.4 Stage 2 access permissions change](#)

B4.3.45.1 Interface

B4.3.45.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400018B
rd	X1	63:0	Address	PA of the RD for the target Realm
rec_ptr	X2	63:0	Address	PA of the target REC
base	X3	63:0	Address	Base of target IPA region
top	X4	63:0	Address	Top of target IPA region

B4.3.45.1.2 Context

The RMI_RTT_SET_S2AP command operates on the following context.

Name	Type	Value	Before	Description
realm_pre	RmmRealm	RealmAt (rd)	true	Realm
rec	RmmRec	RecAt (rec_ptr)	false	REC
not_aligned	RmmRttWalkNotAligned	RttWalkAnyNotAligned (rd, base, top, RMM_RTT_PAGE_LEVEL)	false	RTT walk result which is not aligned to page level

B4.3.45.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
out_top	X1	63:0	Address	Top IPA of range whose S2AP was modified
rtt_tree	X2	63:0	UInt64	Index of RTT tree in which base alignment check failed

If `result` is `RMI_ERROR_RTT` or `RMI_ERROR_RTT_AUX` then the following are true:

- `out_top` is the IPA of the RTTE at which the base alignment check failed.
- `rtt_tree` is the index of the RTT in which the base alignment check failed.

B4.3.45.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
rec_align	pre: !AddrIsGranuleAligned(rec_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
rec_bound	pre: !PaIsDelegable(rec_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
rec_gran_state	pre: GranuleAt(rec_ptr).state != REC post: ResultEqual(result, RMI_ERROR_INPUT)
rec_state	pre: rec.state == REC_RUNNING post: ResultEqual(result, RMI_ERROR_REC)
rec_owner	pre: rec.owner != rd post: ResultEqual(result, RMI_ERROR_REC)
size_valid	pre: UInt(top) <= UInt(base) post: ResultEqual(result, RMI_ERROR_INPUT)
base_bound	pre: base != rec.s2ap_addr post: ResultEqual(result, RMI_ERROR_INPUT)
top_bound	pre: UInt(top) > UInt(rec.s2ap_top) post: ResultEqual(result, RMI_ERROR_INPUT)
top_gran_align	pre: !AddrIsGranuleAligned(top) post: ResultEqual(result, RMI_ERROR_INPUT)
base_align_pri	pre: (not_aligned.valid == RMM_TRUE && !AddrRangeIsWithin(base, top, AlignDownToRttLevel(not_aligned.addr, not_aligned.walk.level), AlignUpToRttLevel(not_aligned.addr, not_aligned.walk.level)) && not_aligned.index == RMM_RTT_TREE_PRIMARY && not_aligned.walk.rtte.s2ap_overlay != rec.s2ap_overlay) post: ResultEqual(result, RMI_ERROR_RTT, not_aligned.walk.level)

ID	Condition
base_align_aux	<pre> pre: (not_aligned.valid == RMM_TRUE && !AddrRangeIsWithin(base, top, AlignDownToRttLevel(not_aligned.addr, not_aligned.walk.level), AlignUpToRttLevel(not_aligned.addr, not_aligned.walk.level)) && not_aligned.index != RMM_RTT_TREE_PRIMARY && not_aligned.walk.rtte.s2ap_overlay != rec.s2ap_overlay) post: ResultEqual(result, RMI_ERROR_RTT_AUX, not_aligned.walk.level) </pre>

B4.3.45.2.1 Failure condition ordering

The RMI_RTT_SET_S2AP command does not have any failure condition orderings.

B4.3.45.3 Success conditions

The RMI_RTT_SET_S2AP command does not have any success conditions.

B4.3.45.4 Footprint

The RMI_RTT_SET_S2AP command does not have any footprint.

B4.3.46 RMI_RTT_UNMAP_UNPROTECTED command

Removes a mapping at an Unprotected IPA.

See also:

- [A5.5 Realm Translation Table](#)
- [B4.3.42 RMI_RTT_MAP_UNPROTECTED command](#)

B4.3.46.1 Interface

B4.3.46.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000162
rd	X1	63:0	Address	PA of the RD for the target Realm
ipa	X2	63:0	Address	IPA at which the Granule is mapped in the target Realm
level	X3	63:0	Int64	RTT level

B4.3.46.1.2 Context

The RMI_RTT_UNMAP_UNPROTECTED command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	RealmAt (rd)	false	Realm
walk	RmmRttWalkResult	RttWalk (rd, ipa, level, RMM_RTT_TREE_PRIMARY)	false	RTT walk result
entry_idx	UInt64	RttEntryIndex (ipa, walk.level)	false	RTTE index
walk_top	Address	RttSkipNonLiveEntries (RttAt (walk.rtt_addr), walk.level, ipa)	false	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

B4.3.46.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
top	X1	63:0	Address	Top IPA of non-live RTT entries, from entry at which the RTT walk terminated

The values of the `result` and `top` output values for different command outcomes are summarized in the following table.

Scenario	result	top	walk.rtte.state
ipa is mapped at the target level	RMI_SUCCESS > ipa Before execution: ASSIGNED_NS After execution: UNASSIGNED_NS		
ipa is not mapped	(RMI_ERROR_RTT, <= level) > ipa UNASSIGNED_NS		
ipa is mapped at a lower level	(RMI_ERROR_RTT, < level) == ipa ASSIGNED_NS		
RTT walk was not performed, due to any other command failure	Another error code	0	Unknown

See also:


- [A5.5.8 RTTE liveness and RTT liveness](#)

B4.3.46.2 Failure conditions

ID	Condition
rd_align	pre: !AddrIsGranuleAligned(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable(rd) post: ResultEqual(result, RMI_ERROR_INPUT)
rd_state	pre: GranuleAt(rd).state != RD post: ResultEqual(result, RMI_ERROR_INPUT)
level_bound	pre: !RttLevelIsBlockOrPage(rd, level) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_align	pre: !AddrIsRttLevelAligned(ipa, level) post: ResultEqual(result, RMI_ERROR_INPUT)
ipa_bound	pre: (UInt(ipa) >= (2 ^ realm.ipa_width) AddrIsProtected(ipa, realm)) post: ResultEqual(result, RMI_ERROR_INPUT)
rtt_walk	pre: walk.level < level post: (ResultEqual(result, RMI_ERROR_RTT, walk.level) && (top == walk_top))
rtte_state	pre: walk.rtte.state != ASSIGNED_NS post: (ResultEqual(result, RMI_ERROR_RTT, walk.level) && (top == walk_top))

B4.3.46.2.1 Failure condition ordering

```
[rd_bound, rd_state] < [rtt_walk, rtte_state]
[level_bound, ipa_bound] < [rtt_walk, rtte_state]
```



B4.3.46.3 Success conditions

ID	Condition
rtte_state	walk.rtte.state == UNASSIGNED_NS
top	top == walk_top

B4.3.46.4 Footprint

ID	Value
rtte	RttEntryAt(walk.rtt_addr, entry_idx)

B4.3.47 RMI_VDEV_ABORT command

Abort device communication associated with a VDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.47.1 Interface

B4.3.47.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000185
vdev_ptr	X1	63:0	Address	PA of the VDEV

B4.3.47.1.2 Context

The RMI_VDEV_ABORT command operates on the following context.

Name	Type	Value	Before	Description
vdev	RmmVdev	VdevAt (vdev_ptr)	false	VDEV

B4.3.47.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.47.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
vdev_align	pre: !AddrIsGranuleAligned (vdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
vdev_bound	pre: !PaIsDelegable (vdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
vdev_gran_state	pre: GranuleAt (vdev_ptr) .state != VDEV post: ResultEqual (result, RMI_ERROR_INPUT)
vdev_state	pre: vdev.state != VDEV_COMMUNICATING post: ResultEqual (result, RMI_ERROR_DEVICE)

B4.3.47.2.1 Failure condition ordering

[da_supp] < [vdev_align, vdev_bound, vdev_gran_state]
[vdev_gran_state] < [vdev_state]



B4.3.47.3 Success conditions

ID	Condition
state	vdev.state == VDEV_READY
comm_state	vdev.comm_state == DEV_COMM_IDLE

B4.3.47.4 Footprint

ID	Value
state	vdev.state
comm_state	vdev.comm_state

B4.3.48 RMI_VDEV_AUX_COUNT command

Get number of auxiliary Granules required for a VDEV.

B4.3.48.1 Interface

B4.3.48.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000160
pdev_flags	X1	63:0	Bits64	PDEV flags
vdev_flags	X2	63:0	Bits64	VDEV flags

B4.3.48.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
aux_count	X1	63:0	UInt64	Number of auxiliary Granules required for a VDEV

B4.3.48.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures().feat_da != FEATURE_TRUE post: ResultEqual(result, RMI_ERROR_NOT_SUPPORTED)

B4.3.48.3 Success conditions

ID	Condition
aux_count	aux_count == VdevAuxCount(RmiPdevFlagsDecode(pdev_flags), RmiVdevFlagsDecode(vdev_flags))

B4.3.48.4 Footprint

The RMI_VDEV_AUX_COUNT command does not have any footprint.

B4.3.49 RMI_VDEV_COMMUNICATE command

Perform device communication associated with a VDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.49.1 Interface

B4.3.49.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000186
pdev_ptr	X1	63:0	Address	PA of the PDEV
vdev_ptr	X2	63:0	Address	PA of the VDEV
data_ptr	X3	63:0	Address	PA of the communication data structure

B4.3.49.1.2 Context

The RMI_VDEV_COMMUNICATE command operates on the following context.

Name	Type	Value	Before	Description
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV
vdev	RmmVdev	VdevAt (vdev_ptr)	false	VDEV
data	RmiDevCommData	RmiDevCommDataAt (data_ptr)	false	Device communication object
num_vdevs_pre	UInt64	pdev.num_vdevs	true	Number of VDEVs associated with the PDEV

B4.3.49.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.49.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures ().feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
pdev_bound	pre: !PaIsDelegable(pdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt(pdev_ptr).state != PDEV post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_align	pre: !AddrIsGranuleAligned(vdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_bound	pre: !PaIsDelegable(vdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_gran_state	pre: GranuleAt(vdev_ptr).state != VDEV post: ResultEqual(result, RMI_ERROR_INPUT)
data_align	pre: !AddrIsGranuleAligned(data_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
data_bound	pre: !PaIsDelegable(data_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
data_pas	pre: !GranuleAccessPermitted(data_ptr, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
req_align	pre: !AddrIsGranuleAligned(data.enter.req_addr) post: ResultEqual(result, RMI_ERROR_INPUT)
req_bound	pre: !PaIsDelegable(data.enter.req_addr) post: ResultEqual(result, RMI_ERROR_INPUT)
req_pas	pre: !GranuleAccessPermitted(data.enter.req_addr, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
resp_align	pre: !AddrIsGranuleAligned(data.enter.resp_addr) post: ResultEqual(result, RMI_ERROR_INPUT)
resp_bound	pre: !PaIsDelegable(data.enter.resp_addr) post: ResultEqual(result, RMI_ERROR_INPUT)
resp_pas	pre: !GranuleAccessPermitted(data.enter.resp_addr, PAS_NS) post: ResultEqual(result, RMI_ERROR_INPUT)
resp_len	pre: data.enter.resp_len > RMM_GRANULE_SIZE post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_pdev	pre: vdev.pdev != pdev_ptr post: ResultEqual(result, RMI_ERROR_DEVICE)
vdev_state	pre: (vdev.state != VDEV_COMMUNICATING && vdev.state != VDEV_STOPPING) post: ResultEqual(result, RMI_ERROR_DEVICE)

B4.3.49.2.1 Failure condition ordering

```
[da_supp] < [pdev_align, pdev_bound, pdev_gran_state, vdev_align,
vdev_bound, vdev_gran_state, data_align, data_bound, data_pas,
req_align, req_bound, req_pas, resp_align, resp_bound, resp_pas,
resp_len]
[pdev_gran_state, vdev_gran_state] < [vdev_pdev, vdev_state]
```



B4.3.49.3 Success conditions

ID	Condition
comm_state	<code>vdev.comm_state == DeviceCommunicate(vdev, data)</code>
error	<code>pre: (DeviceCommunicate(vdev, data) == DEV_COMM_ERROR && vdev.state == VDEV_COMMUNICATING) post: vdev.state == VDEV_ERROR</code>
ready	<code>pre: (DeviceCommunicate(vdev, data) == DEV_COMM_IDLE && vdev.state == VDEV_COMMUNICATING) post: vdev.state == VDEV_READY</code>
stopped	<code>pre: (DeviceCommunicate(vdev, data) != DEV_COMM_ACTIVE && vdev.state == VDEV_STOPPING) post: (vdev.state == VDEV_STOPPED && pdev.num_vdevs == num_vdevs_pre - 1)</code>

B4.3.49.4 Footprint

ID	Value
state	<code>vdev.state</code>
comm_state	<code>vdev.comm_state</code>

B4.3.50 RMI_VDEV_COMPLETE command

Completes a pending VDEV request.

See also:

- [A4.3.14 REC exit due to VDEV request](#)
- [A9.2.2 Mapping from virtual device ID to VDEV object](#)

B4.3.50.1 Interface

B4.3.50.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400018E
rec_ptr	X1	63:0	Address	PA of the REC
vdev_ptr	X2	63:0	Address	PA of the VDEV

B4.3.50.1.2 Context

The RMI_VDEV_COMPLETE command operates on the following context.

Name	Type	Value	Before	Description
rec	RmmRec	RecAt (rec_ptr)	false	REC
vdev	RmmVdev	VdevAt (vdev_ptr)	false	VDEV

B4.3.50.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.50.2 Failure conditions

ID	Condition
rec_align	pre: !AddrIsGranuleAligned (rec_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
rec_bound	pre: !PaIsDelegable (rec_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
recv_state	pre: GranuleAt (rec_ptr).state != REC post: ResultEqual (result, RMI_ERROR_INPUT)
vdev_align	pre: !AddrIsGranuleAligned (vdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
vdev_bound	pre: !PaIsDelegable (vdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
vdev_state	pre: <code>GranuleAt(vdev_ptr).state != VDEV</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
pending	pre: <code>rec.pending != REC_PENDING_VDEV_REQUEST</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
owner	pre: <code>rec.owner != vdev.realm</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
vdev_id	pre: <code>rec.vdev_id != vdev.vdev_id</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
inst_id	pre: <code>(rec.inst_id_valid == RMM_TRUE && rec.inst_id != vdev.inst_id)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>

B4.3.50.2.1 Failure condition ordering

The RMI_VDEV_COMPLETE command does not have any failure condition orderings.

B4.3.50.3 Success conditions

ID	Condition
pending	<code>rec.pending == REC_PENDING_NONE</code>

B4.3.50.4 Footprint

ID	Value
pend	<code>rec.pending</code>

B4.3.51 RMI_VDEV_CREATE command

Create a VDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.51.1 Interface

B4.3.51.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000187
rd	X1	63:0	Address	PA of the RD
pdev_ptr	X2	63:0	Address	PA of the PDEV
vdev_ptr	X3	63:0	Address	PA of the VDEV
params_ptr	X4	63:0	Address	PA of VDEV parameters

B4.3.51.1.2 Context

The RMI_VDEV_CREATE command operates on the following context.

Name	Type	Value	Before	Description
realm_pre	RmmRealm	RealmAt (rd)	true	Realm
realm	RmmRealm	RealmAt (rd)	false	Realm
pdev	RmmPdev	PdevAt (pdev_ptr)	false	PDEV
num_vdevs_pre	UInt64	pdev.num_vdevs	true	Number of VDEVs associated with the PDEV
vdev	RmmVdev	VdevAt (vdev_ptr)	false	VDEV
params	RmiVdevParams	RmiVdevParamsAt (params_ptr)	false	VDEV parameters
num_aux	UInt64	VdevAuxCount (PdevFlags (pdev) , params.flags)	false	Number of auxiliary Granules
rdev	RmmRdev	RdevFromIds (realm, params.vdev_id, realm_pre.num_vdevs)	false	RDEV

B4.3.51.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.51.2 Failure conditions

ID	Condition
da_supp	pre: <code>ImplFeatures().feat_da != FEATURE_TRUE</code> post: <code>ResultEqual(result, RMI_ERROR_NOT_SUPPORTED)</code>
rd_align	pre: <code>!AddrIsGranuleAligned(rd)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
rd_bound	pre: <code>!PaIsDelegable(rd)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
rd_state	pre: <code>GranuleAt(rd).state != RD</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
pdev_align	pre: <code>!AddrIsGranuleAligned(pdev_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
pdev_bound	pre: <code>!PaIsDelegable(pdev_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
pdev_gran_state	pre: <code>GranuleAt(pdev_ptr).state != PDEV</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
pdev_state	pre: <code>pdev.state != PDEV_READY</code> post: <code>ResultEqual(result, RMI_ERROR_DEVICE)</code>
vdev_align	pre: <code>!AddrIsGranuleAligned(vdev_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
vdev_bound	pre: <code>!PaIsDelegable(vdev_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
vdev_gran_state	pre: <code>GranuleAt(vdev_ptr).state != DELEGATED</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
params_align	pre: <code>!AddrIsGranuleAligned(params_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
params_bound	pre: <code>!PaIsDelegable(params_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
params_pas	pre: <code>!GranuleAccessPermitted(params_ptr, PAS_NS)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
params_valid	pre: <code>!RmiVdevParamsIsValid(params_ptr)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
da_en	pre: <code>realm.feat_da != FEATURE_TRUE</code> post: <code>ResultEqual(result, RMI_ERROR_REALM)</code>
num_aux	pre: <code>params.num_aux != num_aux</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
aux_align	pre: <code>!AuxAligned32(params.aux, params.num_aux)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
aux_alias	pre: <code>AuxAlias32(vdev_ptr, params.aux, params.num_aux)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
aux_state	pre: <code>!AuxStateEqual32(params.aux, params.num_aux, DELEGATED)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
tdi_id_free	pre: <code>!TdiIdIsFree(params.tdi_id, pdev.segment_id)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>
tdi_id_bound	pre: <code>(UInt(params.tdi_id) < pdev.rid_base UInt(params.tdi_id) >= pdev.rid_top)</code> post: <code>ResultEqual(result, RMI_ERROR_INPUT)</code>

B4.3.51.2.1 Failure condition ordering

```
[da_supp] < [rd_align, rd_bound, pdev_bound, vdev_align, vdev_bound,
             vdev_gran_state, params_align, params_bound, params_pas,
             params_valid, num_aux, aux_align, aux_alias, aux_state]
[da_supp] < [pdev_gran_state]
[da_supp] < [rd_state]
[pdev_gran_state] < [pdev_state]
[rd_state] < [da_en]
```



B4.3.51.3 Success conditions

ID	Condition
pdev_num_vdevs	<code>pdev.num_vdevs == num_vdevs_pre + 1</code>
gran_state	<code>GranuleAt(vdev_ptr).state == VDEV</code>
vdev_id	<code>vdev.vdev_id == params.vdev_id</code>
tdi_id	<code>vdev.tdi_id == params.tdi_id</code>
pdev	<code>vdev.pdev == pdev_ptr</code>
realm	<code>vdev.realm == rd</code>
state	<code>vdev.state == VDEV_READY</code>
comm_state	<code>vdev.comm_state == DEV_COMM_IDLE</code>
rdev_state	<code>rdev.state == RDEV_NEW</code>
rdev_op	<code>rdev.operation == RDEV_OP_NONE</code>
rdev_vdev_ptr	<code>rdev.vdev_ptr == vdev_ptr</code>
aux	<code>AuxEqual32(vdev.aux, params.aux, num_aux)</code>
num_aux	<code>vdev.num_aux == num_aux</code>
aux_state	<code>AuxStateEqual32(vdev.aux, num_aux, VDEV_AUX)</code>
tdi_id_used	<code>!TdiIdIsFree(params.tdi_id, pdev.segment_id)</code>
inst_id	<code>vdev.inst_id == realm_pre.num_vdevs</code>
realm_num_vdevs	<code>realm.num_vdevs == realm_pre.num_vdevs + 1</code>

B4.3.51.4 Footprint

ID	Value
state	<code>GranuleAt(vdev_ptr).state</code>

ID	Value
aux_state	<code>AuxStates(vdev.aux, num_aux)</code>
pdev_num_vdevs	<code>pdev.num_vdevs</code>
realm_num_vdevs	<code>realm.num_vdevs</code>

DRAFT

B4.3.52 RMI_VDEV_DESTROY command

Destroy a VDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.52.1 Interface

B4.3.52.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000188
rd	X1	63:0	Address	PA of the RD
pdev_ptr	X2	63:0	Address	PA of the PDEV
vdev_ptr	X3	63:0	Address	PA of the VDEV

B4.3.52.1.2 Context

The RMI_VDEV_DESTROY command operates on the following context.

Name	Type	Value	Before	Description
vdev_pre	RmmVdev	VdevAt (vdev_ptr)	true	VDEV
pdev_pre	RmmPdev	PdevAt (pdev_ptr)	true	PDEV

B4.3.52.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status


B4.3.52.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
rd_align	pre: !AddrIsGranuleAligned (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_bound	pre: !PaIsDelegable (rd) post: ResultEqual (result, RMI_ERROR_INPUT)
rd_gran_state	pre: GranuleAt (rd) .state != RD post: ResultEqual (result, RMI_ERROR_INPUT)
pdev_align	pre: !AddrIsGranuleAligned (pdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)

ID	Condition
pdev_bound	pre: !PaIsDelegable(pdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
pdev_gran_state	pre: GranuleAt(pdev_ptr).state != PDEV post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_align	pre: !AddrIsGranuleAligned(vdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_bound	pre: !PaIsDelegable(vdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_gran_state	pre: GranuleAt(vdev_ptr).state != VDEV post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_realm	pre: vdev_pre.realm != rd post: ResultEqual(result, RMI_ERROR_DEVICE)
vdev_pdev	pre: vdev_pre.pdev != pdev_ptr post: ResultEqual(result, RMI_ERROR_DEVICE)
vdev_state	pre: vdev_pre.state != VDEV_STOPPED post: ResultEqual(result, RMI_ERROR_DEVICE)

B4.3.52.2.1 Failure condition ordering

```
[da_supp] < [rd_align, rd_bound, rd_gran_state, pdev_align,
pdev_bound, pdev_gran_state, vdev_align, vdev_bound,
vdev_gran_state]
[rd_gran_state, pdev_gran_state, vdev_gran_state] < [vdev_realm,
vdev_pdev, vdev_state]
```



B4.3.52.3 Success conditions

ID	Condition
gran_state	GranuleAt(vdev_ptr).state == DELEGATED
aux_state	AuxStateEqual32(vdev_pre.aux, vdev_pre.num_aux, DELEGATED)
tdi_id_free	TdiIdIsFree(vdev_pre.tdi_id, pdev_pre.segment_id)

B4.3.52.4 Footprint

ID	Value
state	<code>GranuleAt(vdev_ptr).state</code>
aux_state	<code>AuxStates(vdev_pre.aux, vdev_pre.num_aux)</code>
num_vdevs	<code>pdev.num_vdevs</code>

DRAFT

B4.3.53 RMI_VDEV_GET_STATE command

Get state of a VDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.53.1 Interface

B4.3.53.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000189
vdev_ptr	X1	63:0	Address	PA of the VDEV

B4.3.53.1.2 Context

The RMI_VDEV_GET_STATE command operates on the following context.

Name	Type	Value	Before	Description
vdev	RmmVdev	VdevAt (vdev_ptr)	false	VDEV

B4.3.53.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
state	X1	7:0	RmiVdevState	VDEV state


The following unused bits of RMI_VDEV_GET_STATE output values MBZ: X1[63:8].

B4.3.53.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures().feat_da != FEATURE_TRUE post: ResultEqual(result, RMI_ERROR_NOT_SUPPORTED)
vdev_align	pre: !AddrIsGranuleAligned(vdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_bound	pre: !PaIsDelegable(vdev_ptr) post: ResultEqual(result, RMI_ERROR_INPUT)
vdev_gran_state	pre: GranuleAt(vdev_ptr).state != VDEV post: ResultEqual(result, RMI_ERROR_INPUT)

B4.3.53.2.1 Failure condition ordering

[da_supp] < [vdev_align, vdev_bound, vdev_gran_state]



B4.3.53.3 Success conditions

ID	Condition
state	<code>Equal(state, vdev.state)</code>

B4.3.53.4 Footprint

The RMI_VDEV_GET_STATE command does not have any footprint.

B4.3.54 RMI_VDEV_STOP command

Stop a VDEV.

See also:

- [Chapter A9 Realm device assignment](#)

B4.3.54.1 Interface

B4.3.54.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC400018A
vdev_ptr	X1	63:0	Address	PA of the VDEV

B4.3.54.1.2 Context

The RMI_VDEV_STOP command operates on the following context.

Name	Type	Value	Before	Description
vdev	RmmVdev	VdevAt (vdev_ptr)	false	VDEV

B4.3.54.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status

B4.3.54.2 Failure conditions

ID	Condition
da_supp	pre: ImplFeatures () .feat_da != FEATURE_TRUE post: ResultEqual (result, RMI_ERROR_NOT_SUPPORTED)
vdev_align	pre: !AddrIsGranuleAligned (vdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
vdev_bound	pre: !PaIsDelegable (vdev_ptr) post: ResultEqual (result, RMI_ERROR_INPUT)
vdev_gran_state	pre: GranuleAt (vdev_ptr) .state != VDEV post: ResultEqual (result, RMI_ERROR_INPUT)
vdev_state	pre: (vdev.state != VDEV_READY && vdev.state != VDEV_ERROR) post: ResultEqual (result, RMI_ERROR_DEVICE)

B4.3.54.2.1 Failure condition ordering

[da_supp] < [vdev_align, vdev_bound, vdev_gran_state]
[vdev_gran_state] < [vdev_state]



B4.3.54.3 Success conditions

ID	Condition
state	vdev.state == VDEV_STOPPING
comm_state	vdev.comm_state == DEV_COMM_PENDING

B4.3.54.4 Footprint

ID	Value
state	vdev.state
comm_state	vdev.comm_state

B4.3.55 RMI_VERSION command

Allows the Host and the RMM to determine whether there exists a mutually acceptable revision of the RMM via which the two components can communicate.

On calling this command, the Host provides a requested RMI version.

The output values include a status code and two revisions which are supported by the RMM: a *lower revision* and a *higher revision*.

- The *higher revision* value is the highest interface revision which is supported by the RMM.
- The *lower revision* is less than or equal to the *higher revision*.

The status code and *lower revision* output values indicate which of the following is true, in order of precedence:

- a) The RMM supports an interface revision which is compatible with the requested revision.
 - The status code is RMI_SUCCESS.
 - The *lower revision* is equal to the requested revision.
- b) The RMM does not support an interface revision which is compatible with the requested revision The RMM supports an interface revision which is incompatible with and less than the requested revision.
 - The status code is RMI_ERROR_INPUT.
 - The *lower revision* is the highest interface revision which is both less than the requested revision and supported by the RMM.
- c) The RMM does not support an interface revision which is compatible with the requested revision The RMM supports an interface revision which is incompatible with and greater than the requested revision.
 - The status code is RMI_ERROR_INPUT.
 - The *lower revision* is equal to the *higher revision*.

See also:

- [Chapter B2 Interface versioning](#)
- [B4.1 RMI version](#)

B4.3.55.1 Interface

B4.3.55.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000150
req	X1	63:0	RmiInterfaceVersion	Requested interface revision

B4.3.55.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RmiCommandReturnCode	Command return status
lower	X1	63:0	RmiInterfaceVersion	Lower implemented interface revision
higher	X2	63:0	RmiInterfaceVersion	Higher implemented interface revision

B4.3.55.2 Failure conditions

The RMI_VERSION command does not have any failure conditions.

B4.3.55.3 Success conditions

The RMI_VERSION command does not have any success conditions.

B4.3.55.4 Footprint

The RMI_VERSION command does not have any footprint.

DRAFT

B4.4 RMI types

This section defines types which are used in the RMI interface.

B4.4.1 RmiAddressRange type

The RmiAddressRange structure contains address range.

The RmiAddressRange structure is a [concrete type](#).

The width of the RmiAddressRange structure is 16 (0x10) bytes.

The members of the RmiAddressRange structure are shown in the following table.

Name	Byte offset	Type	Description
base	0x0	Address	Base of address range (inclusive)
top	0x8	Address	Top of address range (exclusive)

The RmiAddressRange structure is used in the following types:

- [RmiPdevParams](#)

B4.4.2 RmiBoolean type

The RmiBoolean enumeration represents a boolean value.

The RmiBoolean enumeration is a [concrete type](#).

The width of the RmiBoolean enumeration is 1 bits.

The values of the RmiBoolean enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_FALSE	False
1	RMI_TRUE	True

The RmiBoolean enumeration is used in the following types:

- [RmiDevCommExitFlags](#)

B4.4.3 RmiCommandReturnCode type

The RmiCommandReturnCode fieldset contains a return code from an RMI command.


The RmiCommandReturnCode fieldset is a [concrete type](#).

The width of the RmiCommandReturnCode fieldset is 64 bits.

See also:

- [Chapter B1 Commands](#)

The fields of the RmiCommandReturnCode fieldset are shown in the following diagram.



The fields of the RmiCommandReturnCode fieldset are shown in the following table.

Name	Bits	Description	Value
status	7:0	Status of the command	RmiStatusCode
index	15:8	Index which identifies the reason for a command failure	UInt8
	63:16	Reserved	MBZ


B4.4.4 RmiDataFlags type

The RmiDataFlags fieldset contains flags provided by the Host during DATA Granule creation.

The RmiDataFlags fieldset is a [concrete type](#).

The width of the RmiDataFlags fieldset is 64 bits.

The fields of the RmiDataFlags fieldset are shown in the following diagram.



The fields of the RmiDataFlags fieldset are shown in the following table.

Name	Bits	Description	Value
measure	0	Whether to measure DATA Granule contents	RmiDataMeasureContent
	63:1	Reserved	SBZ

B4.4.5 RmiDataMeasureContent type

The RmiDataMeasureContent enumeration represents whether to measure DATA Granule contents.

The RmiDataMeasureContent enumeration is a [concrete type](#).

The width of the RmiDataMeasureContent enumeration is 1 bits.

The values of the RmiDataMeasureContent enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_NO_MEASURE_CONTENT	Do not measure DATA Granule contents.
1	RMI_MEASURE_CONTENT	Measure DATA Granule contents.

The RmiDataMeasureContent enumeration is used in the following types:

- [RmiDataFlags](#)

B4.4.6 RmiDevCommData type

The RmiDevCommData structure contains data structure shared between Host and RMM for device communication.

The RmiDevCommData structure is a [concrete type](#).

The width of the RmiDevCommData structure is 4096 (0x1000) bytes.

The members of the RmiDevCommData structure are shown in the following table.

Name	Byte offset	Type	Description
enter	0x0	RmiDevCommEnter	Entry information
exit	0x800	RmiDevCommExit	Exit information

Unused bits of the RmiDevCommData structure SBZ.

B4.4.7 RmiDevCommEnter type

The RmiDevCommEnter structure contains data passed from the Host to the RMM during device communication.

The RmiDevCommEnter structure is a [concrete type](#).

The width of the RmiDevCommEnter structure is 256 (0x100) bytes.

See also:

- [A9.2.5.2 Device communication enter data structure](#)

The members of the RmiDevCommEnter structure are shown in the following table.

Name	Byte offset	Type	Description
status	0x0	RmiDevCommStatus	Status of device transaction
req_addr	0x8	Address	Address of request buffer
resp_addr	0x10	Address	Address of response buffer
resp_len	0x18	UInt64	Amount of valid data in response buffer in bytes

Unused bits of the RmiDevCommEnter structure SBZ.

The RmiDevCommEnter structure is used in the following types:

- [RmiDevCommData](#)

B4.4.8 RmiDevCommExit type

The RmiDevCommExit structure contains data passed from the RMM to the Host during device communication.

The RmiDevCommExit structure is a [concrete type](#).

The width of the RmiDevCommExit structure is 256 (0x100) bytes.

See also:

- [A9.2.5.1 Device communication exit data structure](#)

The members of the RmiDevCommExit structure are shown in the following table.

Name	Byte offset	Type	Description
flags	0x0	RmiDevCommExitFlags	Flags indicating action(s) which the Host is requested to perform
cache_offset	0x8	UInt64	If flags.cache is true, offset in the device response buffer to the start of data to be cached, in bytes
cache_len	0x10	UInt64	If flags.cache is true, amount of data to be cached, in bytes
protocol	0x18	RmiDevCommProtocol	If flags.send is true, protocol to use
req_len	0x20	UInt64	If flags.send is true, amount of valid data in request buffer in bytes
timeout	0x28	UInt64	If flags.wait is true, amount of time to wait for device response in milliseconds

Unused bits of the RmiDevCommExit structure MBZ.

The RmiDevCommExit structure is used in the following types:

- [RmiDevCommData](#)


B4.4.9 RmiDevCommExitFlags type

The RmiDevCommExitFlags fieldset contains flags provided by the RMM during a device transaction.

The RmiDevCommExitFlags fieldset is a [concrete type](#).

The width of the RmiDevCommExitFlags fieldset is 64 bits.

The fields of the RmiDevCommExitFlags fieldset are shown in the following diagram.



The fields of the RmiDevCommExitFlags fieldset are shown in the following table.

Name	Bits	Description	Value
cache	0	Whether the Host is requested to cache data from the device response buffer	RmiBoolean
send	1	Whether the Host is requested to send data from the device request buffer to the device	RmiBoolean
wait	2	Whether the RMM is waiting for a response from the device	RmiBoolean
multi	3	Whether the device transaction contains more than one (device request, device response) tuple	RmiBoolean
	63:4	Reserved	MBZ

The RmiDevCommExitFlags fieldset is used in the following types:

- [RmiDevCommExit](#)

B4.4.10 RmiDevCommProtocol type

The RmiDevCommProtocol enumeration represents protocol used for device communication.

The RmiDevCommProtocol enumeration is a [concrete type](#).

The width of the RmiDevCommProtocol enumeration is 8 bits.

The values of the RmiDevCommProtocol enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_PROTOCOL_SPDM	SPDM See Security Protocol and Data Model (SPDM) [18]
1	RMI_PROTOCOL_SECURE_SPDM	Secure SPDM See Security Protocol and Data Model (SPDM) [18]

Unused encodings for the RmiDevCommProtocol enumeration are reserved for use by future versions of this specification.

The RmiDevCommProtocol enumeration is used in the following types:

- [RmiDevCommExit](#)

B4.4.11 RmiDevCommStatus type

The RmiDevCommStatus enumeration represents status passed from the Host to the RMM during device communication.

The RmiDevCommStatus enumeration is a [concrete type](#).

The width of the RmiDevCommStatus enumeration is 8 bits.

The values of the RmiDevCommStatus enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_DEV_COMM_SUCCESS	Either: <ul style="list-style-type: none"> • The device transaction is PENDING, or • The device transaction is ACTIVE and a device response has been received from the device.
1	RMI_DEV_COMM_ERROR	Either: <ul style="list-style-type: none"> • The device did not provide a device response within the expected time period, or • The device indicated an error.
2	RMI_DEV_COMM_NONE	There are no pending actions for the device.

Unused encodings for the RmiDevCommStatus enumeration are reserved for use by future versions of this specification.

The RmiDevCommStatus enumeration is used in the following types:

- [RmiDevCommEnter](#)


B4.4.12 RmiDevDelegateFlags type

The RmiDevDelegateFlags fieldset contains flags provided by the Host during device memory Granule delegation.

The RmiDevDelegateFlags fieldset is a [concrete type](#).

The width of the RmiDevDelegateFlags fieldset is 64 bits.

The fields of the RmiDevDelegateFlags fieldset are shown in the following diagram.



The fields of the RmiDevDelegateFlags fieldset are shown in the following table.

Name	Bits	Description	Value
share	0	Whether device memory Granule should be shared	RmiDevMemShared
	63:1	Reserved	SBZ

B4.4.13 RmiDevMemShared type

The RmiDevMemShared enumeration represents whether device memory is shared.

The RmiDevMemShared enumeration is a [concrete type](#).

The width of the RmiDevMemShared enumeration is 1 bits.

The values of the RmiDevMemShared enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_DEV_MEM_PRIVATE	Device memory is private
1	RMI_DEV_MEM_SHARED	Device memory is shared

The RmiDevMemShared enumeration is used in the following types:

- [RmiDevDelegateFlags](#)
- [RmiRecExitFlags](#)

B4.4.14 RmiEmulatedMmio type

The RmiEmulatedMmio enumeration represents whether the host has completed emulation for an Emulatable Abort.

The RmiEmulatedMmio enumeration is a [concrete type](#).

The width of the RmiEmulatedMmio enumeration is 1 bits.

The values of the RmiEmulatedMmio enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_NOT_EMULATED_MMIO	Host has not completed emulation for an Emulatable Abort.
1	RMI_EMULATED_MMIO	Host has completed emulation for an Emulatable Abort.

The RmiEmulatedMmio enumeration is used in the following types:

- [RmiRecEnterFlags](#)

B4.4.15 RmiFeature type

The RmiFeature enumeration represents whether a feature is supported or enabled.

The RmiFeature enumeration is a [concrete type](#).

The width of the RmiFeature enumeration is 1 bits.

The values of the RmiFeature enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_FEATURE_FALSE	<ul style="list-style-type: none"> During discovery: Feature is not supported. During selection: Feature is not enabled.
1	RMI_FEATURE_TRUE	<ul style="list-style-type: none"> During discovery: Feature is supported. During selection: Feature is enabled.

The RmiFeature enumeration is used in the following types:

- [RmiRealmFlags1](#)
- [RmiRealmFlags0](#)
- [RmiFeatureRegister0](#)

B4.4.16 RmiFeatureRegister0 type

The RmiFeatureRegister0 fieldset contains RMI feature register 0.


The RmiFeatureRegister0 fieldset is a [concrete type](#).

The width of the RmiFeatureRegister0 fieldset is 64 bits.

See also:

- [Chapter A3 Feature discovery and configuration](#)
- [B4.3.6 RMI_FEATURES command](#)

The fields of the RmiFeatureRegister0 fieldset are shown in the following diagram.



The fields of the RmiFeatureRegister0 fieldset are shown in the following table.

Name	Bits	Description	Value
S2SZ	7:0	Maximum Realm IPA width supported by the RMM. Specifies the input address size for stage 2 translation to be 2^{S2SZ} . Note this format expresses the IPA width directly and is therefore different from the <code>VTCR_EL2.T0SZ</code> encoding.	UInt8
LPA2	8	Whether LPA2 is supported.	RmiFeature

Name	Bits	Description	Value
SVE	9	Whether SVE is supported.	RmiFeature
SVE_VL	13:10	Maximum SVE vector length supported by the RMM. The effective vector length supported by the RMM is $(SVE_VL + 1) * 128$, similar to the value of <code>ZCR_ELx.LEN</code> .	UInt4
NUM_BPS	19:14	Number of breakpoints available, minus one. The value 0 is reserved.	UInt6
NUM_WPS	25:20	Number of watchpoints available, minus one. The value 0 is reserved.	UInt6
PMU	26	Whether PMU is supported	RmiFeature
PMU_NUM_CTRS	31:27	Number of PMU counters available	UInt5
HASH_SHA_256	32	Whether SHA-256 is supported	RmiFeature
HASH_SHA_512	33	Whether SHA-512 is supported	RmiFeature
GICV3_NUM_LRS	37:34	Number of GICv3 List Registers which are available, minus one.	UInt4
MAX_RECS_ORDER	41:38	Order of the maximum number of RECs which can be created per Realm. The maximum number of RECs is computed as follows: $MAX_RECS = (2^{\wedge} MAX_RECS_ORDER) - 1$	UInt4
DA	42	Whether Realm device assignment is supported	RmiFeature
PLANE_RTT	44:43	RTT usage models supported for multi-Plane Realms. If only a single Plane is supported (that is, <code>MAX_NUM_AUX_PLANES</code> is 0), this field can be ignored.	RmiPlaneRttFeature
MAX_NUM_AUX_PLANES	48:45	Maximum number of auxiliary Planes	UInt4
	63:49	Reserved	MBZ

B4.4.17 RmiFeatureRegister1 type

The `RmiFeatureRegister1` fieldset contains RMI feature register 1.


The `RmiFeatureRegister1` fieldset is a [concrete type](#).

The width of the `RmiFeatureRegister1` fieldset is 64 bits.

See also:

- [Chapter A3 Feature discovery and configuration](#)
- [Chapter A11 Realm memory encryption](#)
- [B4.3.6 RMI_FEATURES command](#)

The fields of the `RmiFeatureRegister1` fieldset are shown in the following diagram.



The fields of the RmiFeatureRegister1 fieldset are shown in the following table.

Name	Bits	Description	Value
MAX_MECID	63:0	Maximum MECID.	Bits64

B4.4.18 RmiHashAlgorithm type

The RmiHashAlgorithm enumeration represents hash algorithm.

The RmiHashAlgorithm enumeration is a [concrete type](#).

The width of the RmiHashAlgorithm enumeration is 8 bits.

The values of the RmiHashAlgorithm enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_HASH_SHA_256	SHA-256 (Secure Hash Standard (SHS) [19])
1	RMI_HASH_SHA_512	SHA-512 (Secure Hash Standard (SHS) [19])

Unused encodings for the RmiHashAlgorithm enumeration are reserved for use by future versions of this specification.

The RmiHashAlgorithm enumeration is used in the following types:

- [RmiPdevParams](#)
- [RmiRealmParams](#)

B4.4.19 RmiInjectSea type

The RmiInjectSea enumeration represents whether to inject a Synchronous External Abort into the Realm.

The RmiInjectSea enumeration is a [concrete type](#).

The width of the RmiInjectSea enumeration is 1 bits.

The values of the RmiInjectSea enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_NO_INJECT_SEA	Do not inject an SEA into the Realm.
1	RMI_INJECT_SEA	Inject an SEA into the Realm.

The RmiInjectSea enumeration is used in the following types:

- [RmiRecEnterFlags](#)

B4.4.20 RmiInterfaceVersion type

The RmiInterfaceVersion fieldset contains an RMI interface version.


The RmiInterfaceVersion fieldset is a [concrete type](#).

The width of the RmiInterfaceVersion fieldset is 64 bits.

See also:

- [B4.1 RMI version](#)
- [B4.3.55 RMI_VERSION command](#)

The fields of the RmiInterfaceVersion fieldset are shown in the following diagram.



The fields of the RmiInterfaceVersion fieldset are shown in the following table.

Name	Bits	Description	Value
minor	15:0	Interface minor version number (the value y in interface version $x.y$)	UInt16
major	30:16	Interface major version number (the value x in interface version $x.y$)	UInt15
	63:31	Reserved	MBZ

B4.4.21 RmiLfaPolicy type

The RmiLfaPolicy enumeration represents a Live Firmware Activation policy.

The RmiLfaPolicy enumeration is a [concrete type](#).

The width of the RmiLfaPolicy enumeration is 2 bits.

See also:

- [A3.12 Live Firmware Activation](#)

The values of the RmiLfaPolicy enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_LFA_DISALLOW	LFA is not permitted.
1	RMI_LFA_ALLOW	LFA is permitted.

Unused encodings for the RmiLfaPolicy enumeration are reserved for use by future versions of this specification.

The RmiLfaPolicy enumeration is used in the following types:

- [RmiRealmFlags0](#)

B4.4.22 RmiPdevEvent type

The RmiPdevEvent enumeration represents physical device event.

The RmiPdevEvent enumeration is a [concrete type](#).

The width of the RmiPdevEvent enumeration is 8 bits.

The values of the RmiPdevEvent enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_IDE_KEY_REFRESH	IDE key refresh.

Unused encodings for the RmiPdevEvent enumeration are reserved for use by future versions of this specification.


B4.4.23 RmiPdevFlags type

The RmiPdevFlags fieldset contains flags provided by the Host during PDEV creation.

The RmiPdevFlags fieldset is a [concrete type](#).

The width of the RmiPdevFlags fieldset is 64 bits.

The fields of the RmiPdevFlags fieldset are shown in the following diagram.



The fields of the RmiPdevFlags fieldset are shown in the following table.

Name	Bits	Description	Value
prot_config	1:0	Configuration of protection between system and device	RmiPdevProtConfig
	63:2	Reserved	SBZ

The RmiPdevFlags fieldset is used in the following types:

- [RmiPdevParams](#)

B4.4.24 RmiPdevParams type

The RmiPdevParams structure contains parameters provided by the Host during PDEV creation.

The RmiPdevParams structure is a [concrete type](#).

The width of the RmiPdevParams structure is 4096 (0x1000) bytes.

The members of the RmiPdevParams structure are shown in the following table.

Name	Byte offset	Type	Description
flags	0x0	RmiPdevFlags	Flags
pdev_id	0x8	Bits64	Physical device identifier For a PCIe device this is the PCIe routing identifier of the endpoint. For function 0 of a PCIe device, pdev_id[11:0] MBZ.
segment_id	0x10	Bits16	Segment identifier PCIe Segment identifier of the Root Port and endpoint.
root_id	0x18	Bits16	Root Port identifier Physical PCIe routing identifier of the Root Port to which the endpoint is connected.
cert_id	0x20	UInt64	Certificate identifier
rid_base	0x28	UInt64	Base of requester ID range (inclusive)
rid_top	0x30	UInt64	Top of requester ID range (exclusive)
hash_algo	0x38	RmiHashAlgorithm	Algorithm used to generate device digests
num_aux	0x40	UInt64	Number of auxiliary Granules
ide_sid	0x48	UInt64	IDE stream ID
iocoh_num_addr_range	0x50	UInt64	Number of IO-coherent address ranges
fcoh_num_addr_range	0x58	UInt64	Number of fully-coherent address ranges.
aux[32]	0x100	Address	Addresses of auxiliary Granules
iocoh_addr_range[16]	0x200	RmiAddressRange	IO-coherent address range
fcoh_addr_range[4]	0x300	RmiAddressRange	Fully-coherent address range

Unused bits of the RmiPdevParams structure SBZ.

B4.4.25 RmiPdevProtConfig type

The RmiPdevProtConfig enumeration represents configuration of protection between system and device.

The RmiPdevProtConfig enumeration is a [concrete type](#).

The width of the RmiPdevProtConfig enumeration is 2 bits.

The values of the RmiPdevProtConfig enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_PDEV_IOCOH_E2E_IDE	IO-coherent device with end-to-end protection provided by IDE.
1	RMI_PDEV_IOCOH_E2E_SYS	IO-coherent device with end-to-end protection provided by system construction.
2	RMI_PDEV_FCOH_E2E_IDE	Fully-coherent device with end-to-end protection provided by IDE.
3	RMI_PDEV_FCOH_E2E_SYS	Fully-coherent device with end-to-end protection provided by system construction.

The RmiPdevProtConfig enumeration is used in the following types:

- [RmiPdevFlags](#)

B4.4.26 RmiPdevState type

The RmiPdevState enumeration represents the state of a PDEV.

The RmiPdevState enumeration is a [concrete type](#).

The width of the RmiPdevState enumeration is 8 bits.

The values of the RmiPdevState enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_PDEV_NEW	Initial state of the device.
1	RMI_PDEV_NEEDS_KEY	RMM needs device public key.
2	RMI_PDEV_HAS_KEY	RMM has device public key.
3	RMI_PDEV_READY	Secure connection between the RMM and the device has been established. Physical link between the device and memory is secured. Ready for creation of VDEV instances.
4	RMI_PDEV_COMMUNICATING	The RMM is communicating with the device.
5	RMI_PDEV_STOPPING	The RMM is communicating with the device to terminate the secure connection between the RMM and the device.
6	RMI_PDEV_STOPPED	Secure connection between the RMM and the device has been terminated.
7	RMI_PDEV_ERROR	Device has reported a fatal error.

Unused encodings for the RmiPdevState enumeration are reserved for use by future versions of this specification.

B4.4.27 RmiPlaneRttFeature type

The RmiPlaneRttFeature enumeration represents RTT usage models supported for multi-Plane Realms.

The RmiPlaneRttFeature enumeration is a [concrete type](#).

The width of the RmiPlaneRttFeature enumeration is 2 bits.

See also:

- [A3.11 Support for auxiliary Planes](#)

The values of the RmiPlaneRttFeature enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_PLANE_RTT_AUX	A multi-Plane Realm uses auxiliary RTTs
1	RMI_PLANE_RTT_AUX_SINGLE	A multi-Plane Realm can be configured to either use auxiliary RTTs, or a single RTT
2	RMI_PLANE_RTT_SINGLE	A multi-Plane Realm uses a single RTT

Unused encodings for the RmiPlaneRttFeature enumeration are reserved for use by future versions of this specification.

The RmiPlaneRttFeature enumeration is used in the following types:

- [RmiFeatureRegister0](#)

B4.4.28 RmiPmuOverflowStatus type

The RmiPmuOverflowStatus enumeration represents PMU overflow status.

The RmiPmuOverflowStatus enumeration is a [concrete type](#).

The width of the RmiPmuOverflowStatus enumeration is 8 bits.

The values of the RmiPmuOverflowStatus enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_PMU_OVERFLOW_NOT_ACTIVE	PMU overflow is not active.
1	RMI_PMU_OVERFLOW_ACTIVE	PMU overflow is active.

Unused encodings for the RmiPmuOverflowStatus enumeration are reserved for use by future versions of this specification.

The RmiPmuOverflowStatus enumeration is used in the following types:

- [RmiRecExit](#)


B4.4.29 RmiRealmFlags0 type

The RmiRealmFlags0 fieldset contains flags provided by the Host during Realm creation, which are reflected in Realm Initial Measurement.

The RmiRealmFlags0 fieldset is a [concrete type](#).

The width of the RmiRealmFlags0 fieldset is 64 bits.

The fields of the RmiRealmFlags0 fieldset are shown in the following diagram.



The fields of the RmiRealmFlags0 fieldset are shown in the following table.

Name	Bits	Description	Value
lpa2	0	Whether LPA2 is enabled	RmiFeature
sve	1	Whether SVE is enabled	RmiFeature
pmu	2	Whether PMU is enabled	RmiFeature
da	3	Whether Realm device assignment is enabled	RmiFeature
	4	Reserved	SBZ
lfa_policy	6:5	Live Firmware Activation policy for components within the Realm's TCB	RmiLfaPolicy
	63:7	Reserved	SBZ

The RmiRealmFlags0 fieldset is used in the following types:

- [RmiRealmParams](#)


B4.4.30 RmiRealmFlags1 type

The RmiRealmFlags1 fieldset contains flags provided by the Host during Realm creation, which are not reflected in Realm Initial Measurement.

The RmiRealmFlags1 fieldset is a [concrete type](#).

The width of the RmiRealmFlags1 fieldset is 64 bits.

The fields of the RmiRealmFlags1 fieldset are shown in the following diagram.



The fields of the RmiRealmFlags1 fieldset are shown in the following table.

Name	Bits	Description	Value
rtt_tree_pp	0	RMI_FEATURE_FALSE: all Planes share a single RTT tree RMI_FEATURE_TRUE: each Plane has a separate RTT tree	RmiFeature
	63:1	Reserved	SBZ

The RmiRealmFlags1 fieldset is used in the following types:

- [RmiRealmParams](#)

B4.4.31 RmiRealmParams type

The RmiRealmParams structure contains parameters provided by the Host during Realm creation.

The RmiRealmParams structure is a [concrete type](#).

The width of the RmiRealmParams structure is 4096 (0x1000) bytes.

See also:

- [A2.1.6 Realm parameters](#)
- [B4.3.25 RMI_REALM_CREATE command](#)

The members of the RmiRealmParams structure are shown in the following table.

Name	Byte offset	Type	Description
flags0	0x0	RmiRealmFlags0	Flags
s2sz	0x8	UInt8	IPA width. Specifies the input address size for stage 2 translation to be 2^{s2sz} . Note this format expresses the IPA width directly and is therefore different from the VTCR_EL2.T0SZ encoding.

Name	Byte offset	Type	Description
sve_vl	0x10	UInt8	SVE vector length. The effective vector length requested is $(sve_vl + 1) * 128$, similar to the value of <code>ZCR_ELx.LEN</code> .
num_bps	0x18	UInt8	Number of breakpoints, minus one. The value 0 is reserved.
num_wps	0x20	UInt8	Number of watchpoints, minus one. The value 0 is reserved.
pmu_num_ctrs	0x28	UInt8	Number of PMU counters
hash_algo	0x30	RmiHashAlgorithm	Algorithm used to measure the initial state of the Realm
num_aux_planes	0x38	UInt64	Number of auxiliary Planes
rpv	0x400	Bits512	Realm Personalization Value
vmid	0x800	Bits16	Primary Virtual Machine Identifier
rtt_base	0x808	Address	Base address of primary RTT
rtt_level_start	0x810	Int64	RTT starting level
rtt_num_start	0x818	UInt32	Number of starting level RTTs
flags1	0x820	RmiRealmFlags1	Flags
mecid	0x828	Bits64	MECID
aux_vmid[3]	0xf00	Bits16	Auxiliary Virtual Machine Identifiers
aux_rtt_base[3]	0xf80	Address	Base address of auxiliary RTTs

Unused bits of the RmiRealmParams structure SBZ.


B4.4.32 RmiRecCreateFlags type

The RmiRecCreateFlags fieldset contains flags provided by the Host during REC creation.

The RmiRecCreateFlags fieldset is a [concrete type](#).

The width of the RmiRecCreateFlags fieldset is 64 bits.

The fields of the RmiRecCreateFlags fieldset are shown in the following diagram.



The fields of the RmiRecCreateFlags fieldset are shown in the following table.

Name	Bits	Description	Value
runnable	0	Whether REC is eligible for execution	RmiRecRunnable
	63:1	Reserved	SBZ

The RmiRecCreateFlags fieldset is used in the following types:

- [RmiRecParams](#)

B4.4.33 RmiRecEnter type

The RmiRecEnter structure contains data passed from the Host to the RMM on REC entry.

The RmiRecEnter structure is a [concrete type](#).

The width of the RmiRecEnter structure is 2048 (0x800) bytes.

See also:

- [A4.2.1 RmiRecEnter object](#)
- [B4.3.30 RMI_REC_ENTER command](#)
- [B4.4.35 RmiRecExit type](#)

The members of the RmiRecEnter structure are shown in the following table.

Name	Byte offset	Type	Description
flags	0x0	RmiRecEnterFlags	Flags
gprs[31]	0x200	Bits64	Registers
gicv3_hcr	0x300	Bits64	GICv3 Hypervisor Control Register value
gicv3_lrs[16]	0x308	Bits64	GICv3 List Register values

Unused bits of the RmiRecEnter structure SBZ.

The RmiRecEnter structure is used in the following types:

- [RmiRecRun](#)


B4.4.34 RmiRecEnterFlags type

The RmiRecEnterFlags fieldset contains flags provided by the Host during REC entry.

The RmiRecEnterFlags fieldset is a [concrete type](#).

The width of the RmiRecEnterFlags fieldset is 64 bits.

The fields of the RmiRecEnterFlags fieldset are shown in the following diagram.



The fields of the RmiRecEnterFlags fieldset are shown in the following table.

Name	Bits	Description	Value
emul_mmio	0	Whether the host has completed emulation for an Emulatable Data Abort	RmiEmulatedMmio
inject_sea	1	Whether to inject a Synchronous External Abort into the Realm.	RmiInjectSea
trap_wfi	2	Whether to trap WFI execution by the Realm.	RmiTrap
trap_wfe	3	Whether to trap WFE execution by the Realm.	RmiTrap
ripas_response	4	Host response to RIPAS change request.	RmiResponse
s2ap_response	5	Host response to S2AP change request.	RmiResponse
	63:6	Reserved	SBZ

The RmiRecEnterFlags fieldset is used in the following types:

- [RmiRecEnter](#)

B4.4.35 RmiRecExit type

The RmiRecExit structure contains data passed from the RMM to the Host on REC exit.

The RmiRecExit structure is a [concrete type](#).

The width of the RmiRecExit structure is 2048 (0x800) bytes.

See also:

- [A4.3.1 RmiRecExit object](#)
- [B4.3.30 RMI_REC_ENTER command](#)
- [B4.4.33 RmiRecEnter type](#)

The members of the RmiRecExit structure are shown in the following table.

Name	Byte offset	Type	Description
exit_reason	0x0	RmiRecExitReason	Exit reason
flags	0x8	RmiRecExitFlags	Flags
esr	0x100	Bits64	Exception Syndrome Register
far	0x108	Bits64	Fault Address Register
hpfar	0x110	Bits64	Hypervisor IPA Fault Address register
rtt_tree	0x118	UInt64	Index of RTT tree active at time of the exit
rtt_level	0x120	Int64	Level of requested RTT
gprs[31]	0x200	Bits64	Registers
gicv3_hcr	0x300	Bits64	GICv3 Hypervisor Control Register value
gicv3_lrs[16]	0x308	Bits64	GICv3 List Register values
gicv3_misr	0x388	Bits64	GICv3 Maintenance Interrupt State Register value
gicv3_vmcr	0x390	Bits64	GICv3 Virtual Machine Control Register value
cntp_ctl	0x400	Bits64	Counter-timer Physical Timer Control Register value
cntp_cval	0x408	Bits64	Counter-timer Physical Timer CompareValue Register value
cntv_ctl	0x410	Bits64	Counter-timer Virtual Timer Control Register value
cntv_cval	0x418	Bits64	Counter-timer Virtual Timer CompareValue Register value
ripas_base	0x500	Bits64	Base address of target region for pending RIPAS change
ripas_top	0x508	Bits64	Top address of target region for pending RIPAS change
ripas_value	0x510	RmiRipas	RIPAS value of pending RIPAS change
ripas_dev_pa	0x518	Address	Base PA of device memory region, if RIPAS change is pending due to execution of RSI_RDEV_VALIDATE_MAPPING
s2ap_base	0x520	Bits64	Base address of target region for pending S2AP change
s2ap_top	0x528	Bits64	Top address of target region for pending S2AP change
vdev_id	0x530	Bits64	Virtual device ID
imm	0x600	Bits16	Host call immediate value
plane	0x608	UInt64	Plane index

Name	Byte offset	Type	Description
vdev	0x610	Address	VDEV which triggered REC exit due to device communication
vdev_action	0x618	RmiVdevAction	Action which triggered REC exit due to device communication
pmu_ovf_status	0x700	RmiPmuOverflowStatus	PMU overflow status

Unused bits of the RmiRecExit structure MBZ.

The RmiRecExit structure is used in the following types:

- [RmiRecRun](#)


B4.4.36 RmiRecExitFlags type

The RmiRecExitFlags fieldset contains flags provided by the RMM during REC exit.

The RmiRecExitFlags fieldset is a [concrete type](#).

The width of the RmiRecExitFlags fieldset is 64 bits.

The fields of the RmiRecExitFlags fieldset are shown in the following diagram.



The fields of the RmiRecExitFlags fieldset are shown in the following table.

Name	Bits	Description	Value
ripas_dev_shared	0	Value of shared bit, if RIPAS change is pending due to execution of RSI_RDEV_VALIDATE_MAPPING	RmiDevMemShared
	63:1	Reserved	MBZ

The RmiRecExitFlags fieldset is used in the following types:

- [RmiRecExit](#)

B4.4.37 RmiRecExitReason type

The RmiRecExitReason enumeration represents the reason for a REC exit.

The RmiRecExitReason enumeration is a [concrete type](#).

The width of the RmiRecExitReason enumeration is 8 bits.

The values of the RmiRecExitReason enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_EXIT_SYNC	REC exit due to synchronous exception
1	RMI_EXIT_IRQ	REC exit due to IRQ
2	RMI_EXIT_FIQ	REC exit due to FIQ
3	RMI_EXIT_PSCI	REC exit due to PSCI
4	RMI_EXIT_RIPAS_CHANGE	REC exit due to RIPAS change pending
5	RMI_EXIT_HOST_CALL	REC exit due to Host call
6	RMI_EXIT_SERROR	REC exit due to SError
7	RMI_EXIT_DEV_COMM	REC exit due to device communication
8	RMI_EXIT_RTT_REQUEST	REC exit due to RTT request
9	RMI_EXIT_S2AP_CHANGE	REC exit due to S2AP change pending
10	RMI_EXIT_VDEV_REQUEST	REC exit due to VDEV request

Unused encodings for the RmiRecExitReason enumeration are reserved for use by future versions of this specification.

The RmiRecExitReason enumeration is used in the following types:

- [RmiRecExit](#)

B4.4.38 RmiRecMpidr type

The RmiRecMpidr fieldset contains MPIDR value which identifies a REC.


The RmiRecMpidr fieldset is a [concrete type](#).

The width of the RmiRecMpidr fieldset is 64 bits.

See also:

- [A2.3.3 REC index and MPIDR value](#)
- [B4.3.28 RMI_REC_CREATE command](#)

The fields of the RmiRecMpidr fieldset are shown in the following diagram.



The fields of the RmiRecMpidr fieldset are shown in the following table.

Name	Bits	Description	Value
aff0	3:0	Affinity level 0	Bits4
	7:4	Reserved	SBZ
aff1	15:8	Affinity level 1	Bits8
aff2	23:16	Affinity level 2	Bits8
aff3	31:24	Affinity level 3	Bits8
	63:32	Reserved	SBZ

The RmiRecMpidr fieldset is used in the following types:

- [RmiRecParams](#)

B4.4.39 RmiRecParams type

The RmiRecParams structure contains parameters provided by the Host during REC creation.

The RmiRecParams structure is a [concrete type](#).

The width of the RmiRecParams structure is 4096 (0x1000) bytes.

The number of valid entries in the `aux` array is determined by the return value from the `RMI_REC_AUX_COUNT` command.

See also:

- [B4.3.27 RMI_REC_AUX_COUNT command](#)

The members of the RmiRecParams structure are shown in the following table.

Name	Byte offset	Type	Description
flags	0x0	RmiRecCreateFlags	Flags
mpidr	0x100	RmiRecMpidr	MPIDR of the REC
pc	0x200	Bits64	Program counter
gprs[8]	0x300	Bits64	General-purpose registers
num_aux	0x800	UInt64	Number of auxiliary Granules
aux[16]	0x808	Address	Addresses of auxiliary Granules

Unused bits of the RmiRecParams structure SBZ.

B4.4.40 RmiRecRun type

The RmiRecRun structure contains fields used to share information between RMM and Host during REC entry and REC exit.

The RmiRecRun structure is a [concrete type](#).

The width of the RmiRecRun structure is 4096 (0x1000) bytes.

See also:

- [A4.2.1 RmiRecEnter object](#)

- [A4.3.1 RmiRecExit object](#)
- [B4.3.30 RMI_REC_ENTER command](#)

The members of the RmiRecRun structure are shown in the following table.

Name	Byte offset	Type	Description
enter	0x0	RmiRecEnter	Entry information
exit	0x800	RmiRecExit	Exit information

B4.4.41 RmiRecRunnable type

The RmiRecRunnable enumeration represents whether a REC is eligible for execution.

The RmiRecRunnable enumeration is a [concrete type](#).

The width of the RmiRecRunnable enumeration is 1 bits.

The values of the RmiRecRunnable enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_NOT_RUNNABLE	Not eligible for execution.
1	RMI_RUNNABLE	Eligible for execution.

The RmiRecRunnable enumeration is used in the following types:

- [RmiRecCreateFlags](#)

B4.4.42 RmiResponse type

The RmiResponse enumeration represents whether the Host accepted or rejected a Realm request.

The RmiResponse enumeration is a [concrete type](#).

The width of the RmiResponse enumeration is 1 bits.

The values of the RmiResponse enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_ACCEPT	Host accepted the Realm request.
1	RMI_REJECT	Host rejected the Realm request.

The RmiResponse enumeration is used in the following types:

- [RmiRecEnterFlags](#)

B4.4.43 RmiRipas type

The RmiRipas enumeration represents realm IPA state.

The RmiRipas enumeration is a [concrete type](#).

The width of the RmiRipas enumeration is 8 bits.

The values of the RmiRipas enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_EMPTY	Address where no Realm resources are mapped.
1	RMI_RAM	Address where private code or data owned by the Realm is mapped.
2	RMI_DESTROYED	Address which is inaccessible to the Realm due to an action taken by the Host.
3	RMI_DEV	Address where memory of an assigned Realm device is mapped.

Unused encodings for the RmiRipas enumeration are reserved for use by future versions of this specification.

The RmiRipas enumeration is used in the following types:

- [RmiRecExit](#)

B4.4.44 RmiRttEntryState type

The RmiRttEntryState enumeration represents the state of an RTTE.

The RmiRttEntryState enumeration is a [concrete type](#).

The width of the RmiRttEntryState enumeration is 8 bits.

The values of the RmiRttEntryState enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_UNASSIGNED	This RTTE is not associated with any Granule.
1	RMI_ASSIGNED	The output address of this RTTE points to: <ul style="list-style-type: none">• a DATA Granule, if the input address is a Protected IPA, or• an NS Granule, if the input address is an Unprotected IPA.
2	RMI_TABLE	The output address of this RTTE points to the next-level RTT.
3	RMI_ASSIGNED_DEV_PRIVATE	The output address of this RTTE points to an DEV_PRIVATE Granule.
4	RMI_ASSIGNED_DEV_SHARED	The output address of this RTTE points to an DEV_SHARED Granule.
5	RMI_AUX_DESTROYED	An auxiliary RTT was destroyed.

Unused encodings for the RmiRttEntryState enumeration are reserved for use by future versions of this specification.

B4.4.45 RmiSignatureAlgorithm type

The RmiSignatureAlgorithm enumeration represents signature algorithm.

The RmiSignatureAlgorithm enumeration is a [concrete type](#).

The width of the RmiSignatureAlgorithm enumeration is 8 bits.

The values of the RmiSignatureAlgorithm enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_SIG_RSASSA_3072	SSA-3072 (<i>RSA Cryptography Specifications Version 2.2</i> [20])
1	RMI_SIG_ECDSA_P256	ECDSA-P256 (<i>Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA)</i> [21])
2	RMI_SIG_ECDSA_P384	ECDSA-P384 (<i>Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA)</i> [21])

Unused encodings for the RmiSignatureAlgorithm enumeration are reserved for use by future versions of this specification.

B4.4.46 RmiStatusCode type

The RmiStatusCode enumeration represents the status of an RMI operation.

The RmiStatusCode enumeration is a [concrete type](#).

The width of the RmiStatusCode enumeration is 8 bits.

See also:

- [B1.3 Command registers](#)
- [B1.5 Command context values](#)

The values of the RmiStatusCode enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_SUCCESS	Command completed successfully
1	RMI_ERROR_INPUT	The value of a command input value caused the command to fail
2	RMI_ERROR_REALM	An attribute of a Realm does not match the expected value
3	RMI_ERROR_REC	An attribute of a REC does not match the expected value
4	RMI_ERROR_RTT	An RTT walk terminated before reaching the target RTT level, or reached an RTTE with an unexpected value
5	RMI_ERROR_NOT_SUPPORTED	The command is not supported
6	RMI_ERROR_DEVICE	An attribute of a device does not match the expected value
7	RMI_ERROR_RTT_AUX	RTTE in an auxiliary RTT contained an unexpected value

Unused encodings for the RmiStatusCode enumeration are reserved for use by future versions of this specification.

The RmiStatusCode enumeration is used in the following types:

- [RmiCommandReturnCode](#)

B4.4.47 RmiTrap type

The RmiTrap enumeration represents whether a trap is enabled.

The RmiTrap enumeration is a [concrete type](#).

The width of the RmiTrap enumeration is 1 bits.

The values of the RmiTrap enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_NO_TRAP	Trap is disabled.
1	RMI_TRAP	Trap is enabled.

The RmiTrap enumeration is used in the following types:

- [RmiRecEnterFlags](#)

B4.4.48 RmiUnprotectedS2AP type

The RmiUnprotectedS2AP enumeration represents mapping from Stage 2 base permission index to Stage 2 base permission value for an Unprotected IPA.

The RmiUnprotectedS2AP enumeration is a [concrete type](#).

The width of the RmiUnprotectedS2AP enumeration is 4 bits.

The values of the RmiUnprotectedS2AP enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_UNPROTECTED_S2AP_NO_ACCESS	No access
1	RMI_UNPROTECTED_S2AP_RO	Read only
2	RMI_UNPROTECTED_S2AP_WO	Write only
3	RMI_UNPROTECTED_S2AP_RW	Read write

Unused encodings for the RmiUnprotectedS2AP enumeration are reserved for use by future versions of this specification.

B4.4.49 RmiVdevAction type

The RmiVdevAction enumeration represents realm action which triggered REC exit due to device communication.

The RmiVdevAction enumeration is a [concrete type](#).

The width of the RmiVdevAction enumeration is 8 bits.

The values of the RmiVdevAction enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_VDEV_ACTION_GET_INTERFACE_REPORT	Exit triggered by RSI_RDEV_GET_INTERFACE_REPORT

Encoding	Name	Description
1	RMI_VDEV_ACTION_GET_MEASUREMENTS	Exit triggered by RSI_RDEV_GET_MEASUREMENTS
2	RMI_VDEV_ACTION_LOCK	Exit triggered by RSI_RDEV_LOCK
3	RMI_VDEV_ACTION_START	Exit triggered by RSI_RDEV_START
4	RMI_VDEV_ACTION_STOP	Exit triggered by RSI_RDEV_STOP

Unused encodings for the RmiVdevAction enumeration are reserved for use by future versions of this specification.
The RmiVdevAction enumeration is used in the following types:

- [RmiRecExit](#)


B4.4.50 RmiVdevFlags type

The RmiVdevFlags fieldset contains flags provided by the Host during VDEV creation.

The RmiVdevFlags fieldset is a [concrete type](#).

The width of the RmiVdevFlags fieldset is 64 bits.

The fields of the RmiVdevFlags fieldset are shown in the following diagram.



The fields of the RmiVdevFlags fieldset are shown in the following table.

Name	Bits	Description	Value
	63:0	Reserved	SBZ

The RmiVdevFlags fieldset is used in the following types:

- [RmiVdevParams](#)

B4.4.51 RmiVdevParams type

The RmiVdevParams structure contains parameters provided by the Host during VDEV creation.

The RmiVdevParams structure is a [concrete type](#).

The width of the RmiVdevParams structure is 4096 (0x1000) bytes.

The members of the RmiVdevParams structure are shown in the following table.

Name	Byte offset	Type	Description
flags	0x0	RmiVdevFlags	Flags

Name	Byte offset	Type	Description
vdev_id	0x8	Bits64	Virtual device identifier For a PCIe device this is the PCIe routing identifier of the virtual endpoint.
tdi_id	0x10	Bits64	TDI identifier
num_aux	0x18	UInt64	Number of auxiliary Granules
aux[32]	0x100	Address	Addresses of auxiliary Granules

Unused bits of the RmiVdevParams structure SBZ.

B4.4.52 RmiVdevState type

The RmiVdevState enumeration represents the state of a VDEV.

The RmiVdevState enumeration is a [concrete type](#).

The width of the RmiVdevState enumeration is 8 bits.

The values of the RmiVdevState enumeration are shown in the following table.

Encoding	Name	Description
0	RMI_VDEV_READY	No device transaction is associated with the VDEV.
1	RMI_VDEV_COMMUNICATING	The RMM is communicating with the VDEV.
2	RMI_VDEV_STOPPING	The RMM is communicating with the VDEV to stop the device interface.
3	RMI_VDEV_STOPPED	Device interface is stopped.
4	RMI_VDEV_ERROR	Device interface has reported a fatal error.

Unused encodings for the RmiVdevState enumeration are reserved for use by future versions of this specification.

Chapter B5

Realm Services Interface

This chapter defines the interface used by Realm software to request services from the RMM.

B5.1 RSI version

R_{QKLGZ} This specification defines version 1.1 of the Realm Services Interface.

See also:

- [Chapter B2 Interface versioning](#)
- [B5.3.25 RSI_VERSION command](#)

B5.2 RSI command return codes

I_{CYQDJ} An RSI command return code indicates whether the command

- succeeded, or
- failed, and the reason for the failure.

I_{DQJSP} If an RSI command succeeds then it returns RSI_SUCCESS.

I_{YMHKC} Multiple failure conditions in an RSI command may return the same return code.

R_{MLBDM} If an input to an RSI command uses an invalid encoding then the command fails and returns RSI_ERROR_INPUT.

Command inputs include registers and in-memory data structures.

Invalid encodings include:

- using a reserved encoding in an enumeration

See also:

- [B5.4.2 RsiCommandReturnCode type](#)

B5.3 RSI commands

The following table summarizes the FIDs of commands in the RSI interface.

FID	Command
0xC4000190	RSI_VERSION
0xC4000191	RSI_FEATURES
0xC4000192	RSI_MEASUREMENT_READ
0xC4000193	RSI_MEASUREMENT_EXTEND
0xC4000194	RSI_ATTESTATION_TOKEN_INIT
0xC4000195	RSI_ATTESTATION_TOKEN_CONTINUE
...	
0xC4000197	RSI_IPA_STATE_SET
0xC4000198	RSI_IPA_STATE_GET
0xC4000199	RSI_HOST_CALL
...	
0xC40001A0	RSI_MEM_GET_PERM_VALUE
0xC40001A1	RSI_MEM_SET_PERM_INDEX
0xC40001A2	RSI_MEM_SET_PERM_VALUE
0xC40001A3	RSI_PLANE_ENTER
0xC40001A4	RSI_RDEV_CONTINUE
0xC40001A5	RSI_RDEV_GET_INFO
0xC40001A6	RSI_RDEV_GET_INTERFACE_REPORT
0xC40001A7	RSI_RDEV_GET_MEASUREMENTS
0xC40001A8	RSI_RDEV_GET_STATE
0xC40001A9	RSI_RDEV_LOCK
0xC40001AA	RSI_RDEV_START
0xC40001AB	RSI_RDEV_STOP
0xC40001AC	RSI_RDEV_VALIDATE_MAPPING
0xC40001AD	RSI_REALM_CONFIG
0xC40001AE	RSI_PLANE_REG_READ
0xC40001AF	RSI_PLANE_REG_WRITE

B5.3.1 RSI_ATTESTATION_TOKEN_CONTINUE command

Continue the operation to retrieve an attestation token.

See also:

- [A7.2 Realm attestation](#)
- [B5.3.2 RSI_ATTESTATION_TOKEN_INIT command](#)

B5.3.1.1 Interface

B5.3.1.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000195
addr	X1	63:0	Address	IPA of the Granule to which the token will be written
offset	X2	63:0	UInt64	Offset within Granule to start of buffer in bytes
size	X3	63:0	UInt64	Size of buffer in bytes

B5.3.1.1.2 Context

The RSI_ATTESTATION_TOKEN_CONTINUE command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rec	RmmRec	CurrentRec()	false	Current REC

B5.3.1.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
len	X1	63:0	UInt64	Number of bytes written to buffer

B5.3.1.2 Failure conditions

ID	Condition
addr_align	pre: !AddrIsGranuleAligned(addr) post: result == RSI_ERROR_INPUT
addr_bound	pre: !AddrIsProtected(addr, realm) post: result == RSI_ERROR_INPUT
offset_bound	pre: offset >= RMM_GRANULE_SIZE post: result == RSI_ERROR_INPUT

ID	Condition
size_overflow	pre: offset + size < offset post: result == RSI_ERROR_INPUT
size_bound	pre: offset + size > RMM_GRANULE_SIZE post: result == RSI_ERROR_INPUT
state	pre: rec.attest_state != ATTEST_IN_PROGRESS post: result == RSI_ERROR_STATE
unknown	pre: Token generation failed for an unknown or IMPDEF reason. post: result == RSI_ERROR_UNKNOWN

B5.3.1.2.1 Failure condition ordering

The RSI_ATTESTATION_TOKEN_CONTINUE command does not have any failure condition orderings.

B5.3.1.3 Success conditions

ID	Condition
incomplete	pre: Token generation is not complete. post: result == RSI_INCOMPLETE
complete	pre: Token generation is complete. post: rec.attest_state == NO_ATTEST_IN_PROGRESS

B5.3.1.4 Footprint

ID	Value
state	rec.attest_state

B5.3.2 RSI_ATTESTATION_TOKEN_INIT command

Initialize the operation to retrieve an attestation token.

See also:

- [A7.2 Realm attestation](#)
- [B5.3.1 RSI_ATTESTATION_TOKEN_CONTINUE command](#)

B5.3.2.1 Interface

B5.3.2.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000194
challenge_0	X1	63:0	Bits64	Doubleword 0 of the challenge value
challenge_1	X2	63:0	Bits64	Doubleword 1 of the challenge value
challenge_2	X3	63:0	Bits64	Doubleword 2 of the challenge value
challenge_3	X4	63:0	Bits64	Doubleword 3 of the challenge value
challenge_4	X5	63:0	Bits64	Doubleword 4 of the challenge value
challenge_5	X6	63:0	Bits64	Doubleword 5 of the challenge value
challenge_6	X7	63:0	Bits64	Doubleword 6 of the challenge value
challenge_7	X8	63:0	Bits64	Doubleword 7 of the challenge value

B5.3.2.1.2 Context

The RSI_ATTESTATION_TOKEN_INIT command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rec	RmmRec	CurrentRec()	false	Current REC

B5.3.2.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
size	X1	63:0	UInt64	Upper bound on attestation token size in bytes

B5.3.2.2 Failure conditions

The RSI_ATTESTATION_TOKEN_INIT command does not have any failure conditions.

B5.3.2.3 Success conditions

ID	Condition
state	<code>rec.attest_state == ATTEST_IN_PROGRESS</code>
challenge	<code>rec.attest_challenge == [challenge_0, challenge_1, challenge_2, challenge_3, challenge_4, challenge_5, challenge_6, challenge_7]</code>

B5.3.2.4 Footprint

ID	Value
state	<code>rec.attest_state</code>
challenge	<code>rec.attest_challenge</code>

B5.3.3 RSI_FEATURES command

Read feature register.

The following table indicates which feature register is returned depending on the index provided.

Index	Feature register
0	RSI feature register 0

See also:

- [Chapter A3 Feature discovery and configuration](#)

B5.3.3.1 Interface

B5.3.3.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000191
index	X1	63:0	UInt64	Feature register index

B5.3.3.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
value	X1	63:0	Bits64	Feature register value

B5.3.3.2 Failure conditions

The RSI_FEATURES command does not have any failure conditions.

B5.3.3.3 Success conditions

ID	Condition
value	<code>value == RsiFeatureRegisterEncode(index)</code>

B5.3.3.4 Footprint

The RSI_FEATURES command does not have any footprint.

B5.3.4 RSI_HOST_CALL command

Make a Host call.

See also:

- [A4.5 Host call](#)

B5.3.4.1 Interface

B5.3.4.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000199
addr	X1	63:0	Address	IPA of the Host call data structure

B5.3.4.1.2 Context

The RSI_HOST_CALL command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rec	RmmRec	CurrentRec()	false	Current REC
data	RsiHostCall	RsiHostCallAt(addr)	false	Host call data structure

B5.3.4.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

B5.3.4.2 Failure conditions

ID	Condition
addr_align	pre: !AddrIsAligned(addr, 256) post: result == RSI_ERROR_INPUT
addr_bound	pre: !AddrIsProtected(addr, realm) post: result == RSI_ERROR_INPUT

B5.3.4.2.1 Failure condition ordering

The RSI_HOST_CALL command does not have any failure condition orderings.

B5.3.4.3 Success conditions

The RSI_HOST_CALL command does not have any success conditions.

B5.3.4.4 Footprint

ID	Value
pending	rec.pending

DRAFT

B5.3.5 RSI_IPA_STATE_GET command

Get RIPAS of a target IPA range.

See also:

- [A5.2 Realm view of memory management](#)
- [B5.3.6 RSI_IPA_STATE_SET command](#)

B5.3.5.1 Interface

B5.3.5.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000198
base	X1	63:0	Address	Base of target IPA region
top	X2	63:0	Address	End of target IPA region

B5.3.5.1.2 Context

The RSI_IPA_STATE_GET command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm

B5.3.5.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
out_top	X1	63:0	Address	Top of IPA region which has the reported RIPAS value
ripas	X2	7:0	RsiRipas	RIPAS value

The following unused bits of RSI_IPA_STATE_GET output values MBZ: X2[63:8].

If `result == RSI_SUCCESS` then all of the following are true:

- `out_top > base`
- `out_top <= top`
- All addresses within the range `[base, out_top)` have the RIPAS value `ripas`.

Note that the RIPAS of a Protected IPA can change at any time to DESTROYED without the Realm taking any action.

See also:

- [A5.2.5 Changes to RIPAS while Realm state is REALM_ACTIVE](#)

B5.3.5.2 Failure conditions

ID	Condition
base_align	pre: !AddrIsGranuleAligned(base) post: result == RSI_ERROR_INPUT
end_align	pre: !AddrIsGranuleAligned(top) post: result == RSI_ERROR_INPUT
size_valid	pre: UInt(top) <= UInt(base) post: result == RSI_ERROR_INPUT
rgn_bound	pre: !AddrRangeIsProtected(base, top, realm) post: result == RSI_ERROR_INPUT

B5.3.5.2.1 Failure condition ordering

The RSI_IPA_STATE_GET command does not have any failure condition orderings.

B5.3.5.3 Success conditions

The RSI_IPA_STATE_GET command does not have any success conditions.

B5.3.5.4 Footprint

The RSI_IPA_STATE_GET command does not have any footprint.

DRAFT

B5.3.6 RSI_IPA_STATE_SET command

Request RIPAS of a target IPA range to be changed to a specified value.

See also:

- [A5.2 Realm view of memory management](#)
- [A5.4 RIPAS change](#)
- [B5.3.5 RSI_IPA_STATE_GET command](#)

B5.3.6.1 Interface

B5.3.6.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000197
base	X1	63:0	Address	Base of target IPA region
top	X2	63:0	Address	Top of target IPA region
ripas	X3	7:0	RsiRipas	RIPAS value
flags	X4	63:0	RsiRipasChangeFlags	Flags

The following unused bits of RSI_IPA_STATE_SET input values SBZ: X3[63:8].

B5.3.6.1.2 Context

The RSI_IPA_STATE_SET command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rec	RmmRec	CurrentRec()	false	Current REC

B5.3.6.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
new_base	X1	63:0	Address	Base of IPA region which was not modified by the command
response	X2	0:0	RsiResponse	Whether the Host accepted or rejected the request

The following unused bits of RSI_IPA_STATE_SET output values MBZ: X2[63:1].

If the Host rejects the request then:

- `result == RSI_SUCCESS`
- `new_base == base`
- `response == RSI_REJECT`

B5.3.6.2 Failure conditions

ID	Condition
base_align	pre: !AddrIsGranuleAligned(base) post: result == RSI_ERROR_INPUT
top_align	pre: !AddrIsGranuleAligned(top) post: result == RSI_ERROR_INPUT
size_valid	pre: UInt(top) <= UInt(base) post: result == RSI_ERROR_INPUT
rgn_bound	pre: !AddrRangeIsProtected(base, top, realm) post: result == RSI_ERROR_INPUT
ripas_valid	pre: (ripas != RSI_EMPTY) && (ripas != RSI_RAM) post: result == RSI_ERROR_INPUT

B5.3.6.2.1 Failure condition ordering

The RSI_IPA_STATE_SET command does not have any failure condition orderings.

B5.3.6.3 Success conditions

ID	Condition
new_base	new_base == rec.ripas_addr
response	response == RecRipasResponseToRsi(rec)

B5.3.6.4 Footprint

The RSI_IPA_STATE_SET command does not have any footprint.

B5.3.7 RSI_MEASUREMENT_EXTEND command

Extend Realm Extensible Measurement (REM) value.

B5.3.7.1 Interface

B5.3.7.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000193
index	X1	63:0	UInt64	Measurement index
size	X2	63:0	UInt64	Measurement size in bytes
value_0	X3	63:0	Bits64	Doubleword 0 of the measurement value
value_1	X4	63:0	Bits64	Doubleword 1 of the measurement value
value_2	X5	63:0	Bits64	Doubleword 2 of the measurement value
value_3	X6	63:0	Bits64	Doubleword 3 of the measurement value
value_4	X7	63:0	Bits64	Doubleword 4 of the measurement value
value_5	X8	63:0	Bits64	Doubleword 5 of the measurement value
value_6	X9	63:0	Bits64	Doubleword 6 of the measurement value
value_7	X10	63:0	Bits64	Doubleword 7 of the measurement value

B5.3.7.1.2 Context

The RSI_MEASUREMENT_EXTEND command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
realm_pre	RmmRealm	CurrentRealm()	true	Current Realm
meas_pre	RmmRealmMeasurement	realm_pre.measurements[index]	true	Previous measurement value

B5.3.7.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

B5.3.7.2 Failure conditions

ID	Condition
index_bound	pre: index < 1 index > 4 post: result == RSI_ERROR_INPUT
size_bound	pre: size > 64 post: result == RSI_ERROR_INPUT

B5.3.7.2.1 Failure condition ordering

The RSI_MEASUREMENT_EXTEND command does not have any failure condition orderings.

B5.3.7.3 Success conditions

ID	Condition
realm_meas	realm.measurements[index] == RemExtend(realm.hash_algo, meas_pre, [value_0, value_1, value_2, value_3, value_4, value_5, value_6, value_7][(RMM_REALM_MEASUREMENT_WIDTH-1):0], size)

B5.3.7.4 Footprint

ID	Value
realm_meas	realm.measurements[index]

B5.3.8 RSI_MEASUREMENT_READ command

Read measurement for the current Realm.

See also:

- [A7.1 Realm measurements](#)
- [D1.2.1 Realm creation flow](#)

B5.3.8.1 Interface

B5.3.8.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000192
index	X1	63:0	UInt64	Measurement index

`index` 0 selects the RIM. An `index` of 1 or greater selects the corresponding REM.

B5.3.8.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
value_0	X1	63:0	Bits64	Doubleword 0 of the Realm measurement identified by “index”
value_1	X2	63:0	Bits64	Doubleword 1 of the Realm measurement identified by “index”
value_2	X3	63:0	Bits64	Doubleword 2 of the Realm measurement identified by “index”
value_3	X4	63:0	Bits64	Doubleword 3 of the Realm measurement identified by “index”
value_4	X5	63:0	Bits64	Doubleword 4 of the Realm measurement identified by “index”
value_5	X6	63:0	Bits64	Doubleword 5 of the Realm measurement identified by “index”
value_6	X7	63:0	Bits64	Doubleword 6 of the Realm measurement identified by “index”
value_7	X8	63:0	Bits64	Doubleword 7 of the Realm measurement identified by “index”

If the size of the measurement value is smaller than 512 bits, the output values are padded with zeroes.

B5.3.8.2 Failure conditions

ID	Condition
index_bound	pre: index > 4 post: result == RSI_ERROR_INPUT

B5.3.8.3 Success conditions

The RSI_MEASUREMENT_READ command does not have any success conditions.

B5.3.8.4 Footprint

The RSI_MEASUREMENT_READ command does not have any footprint.

DRAFT

B5.3.9 RSI_MEM_GET_PERM_VALUE command

Get overlay permission value for a specified (plane index, overlay permission index) tuple.

See also:

- [A10.3.2 Stage 2 access permissions](#)

B5.3.9.1 Interface

B5.3.9.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A0
plane_index	X1	63:0	UInt64	Plane index
perm_index	X2	63:0	UInt64	Permission index

B5.3.9.1.2 Context

The RSI_MEM_GET_PERM_VALUE command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm

B5.3.9.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
value	X1	63:0	Bits64	Memory permission value

B5.3.9.2 Failure conditions

ID	Condition
plane_bound	pre: plane_index > realm.num_aux_planes post: result == RSI_ERROR_INPUT
perm_bound	pre: perm_index >= RMM_NUM_PERM_OVERLAY_INDICES post: result == RSI_ERROR_INPUT

B5.3.9.2.1 Failure condition ordering

The RSI_MEM_GET_PERM_VALUE command does not have any failure condition orderings.

B5.3.9.3 Success conditions

ID	Condition
label	<code>value == realm.overlay_perms[plane_index].values[perm_index]</code>

B5.3.9.4 Footprint

The RSI_MEM_GET_PERM_VALUE command does not have any footprint.

DRAFT

B5.3.10 RSI_MEM_SET_PERM_INDEX command

Set overlay permission index for a specified IPA range.

See also:

- [A10.3.2 Stage 2 access permissions](#)

B5.3.10.1 Interface

B5.3.10.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A1
base	X1	63:0	Address	Base of target IPA region
top	X2	63:0	Address	Top of target IPA region
perm_index	X3	63:0	UInt64	Permission index
cookie	X4	63:0	Bits64	Cookie value

B5.3.10.1.2 Context

The RSI_MEM_SET_PERM_INDEX command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rec	RmmRec	CurrentRec()	false	Current REC

B5.3.10.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
new_base	X1	63:0	Address	Base of IPA region which was not modified by the command
response	X2	0:0	RsiResponse	Whether the Host accepted or rejected the request
new_cookie	X3	63:0	Bits64	New cookie value

The following unused bits of RSI_MEM_SET_PERM_INDEX output values MBZ: X2[63:1].

B5.3.10.2 Failure conditions

ID	Condition
base_align	pre: <code>!AddrIsGranuleAligned(base)</code> post: <code>result == RSI_ERROR_INPUT</code>

ID	Condition
top_align	pre: !AddrIsGranuleAligned(top) post: result == RSI_ERROR_INPUT
size_valid	pre: UInt(top) <= UInt(base) post: result == RSI_ERROR_INPUT
rgn_bound	pre: !AddrRangeIsProtected(base, top, realm) post: result == RSI_ERROR_INPUT
perm_bound	pre: perm_index >= RMM_NUM_PERM_OVERLAY_INDICES post: result == RSI_ERROR_INPUT
cookie	pre: Cookie is invalid post: result == RSI_ERROR_INPUT

B5.3.10.2.1 Failure condition ordering

The RSI_MEM_SET_PERM_INDEX command does not have any failure condition orderings.

B5.3.10.3 Success conditions

ID	Condition
locked	realm.overlay_locked[perm_index] == MEM_PERM_LOCKED
new_base	new_base == rec.s2ap_addr
response	response == RecS2APResponseToRsi(rec)

B5.3.10.4 Footprint

The RSI_MEM_SET_PERM_INDEX command does not have any footprint.

B5.3.11 RSI_MEM_SET_PERM_VALUE command

Set overlay permission value for a specified (plane index, overlay permission index) tuple.

See also:

- [A10.3.2 Stage 2 access permissions](#)

B5.3.11.1 Interface

B5.3.11.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A2
plane_index	X1	63:0	UInt64	Plane index
perm_index	X2	63:0	UInt64	Permission index
value	X3	63:0	Bits64	Memory permission value

B5.3.11.1.2 Context

The RSI_MEM_SET_PERM_VALUE command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm

B5.3.11.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

B5.3.11.2 Failure conditions

ID	Condition
plane_bound	pre: (plane_index == 0 plane_index > realm.num_aux_planes) post: result == RSI_ERROR_INPUT
perm_bound	pre: perm_index >= RMM_NUM_PERM_OVERLAY_INDICES post: result == RSI_ERROR_INPUT
locked	pre: realm.overlay_locked[perm_index] == MEM_PERM_LOCKED post: result == RSI_ERROR_INPUT
supported	pre: ! MemPermLabelSupported (value) post: result == RSI_ERROR_INPUT

B5.3.11.2.1 Failure condition ordering

The RSI_MEM_SET_PERM_VALUE command does not have any failure condition orderings.

B5.3.11.3 Success conditions

ID	Condition
label	<code>realm.overlay_perms[plane_index].values[perm_index] == value</code>

B5.3.11.4 Footprint

The RSI_MEM_SET_PERM_VALUE command does not have any footprint.

DRAFT

B5.3.12 RSI_PLANE_ENTER command

Enter a Plane.

See also:

- [A10.2 Planes exception model](#)

B5.3.12.1 Interface

B5.3.12.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A3
plane_idx	X1	63:0	UInt64	Index of target Plane
run_ptr	X2	63:0	Address	IPA of PlaneRun object

B5.3.12.1.2 Context

The RSI_PLANE_ENTER command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
run	RsiPlaneRun	RsiPlaneRunAt (realm, run_ptr)	false	PlaneRun object

B5.3.12.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

B5.3.12.2 Failure conditions

ID	Condition
idx_bound	pre: (plane_idx == 0 plane_idx > realm.num_aux_planes) post: result == RSI_ERROR_INPUT
run_align	pre: !AddrIsGranuleAligned (run_ptr) post: result == RSI_ERROR_INPUT
run_bound	pre: !AddrIsProtected (run_ptr, realm) post: result == RSI_ERROR_INPUT

B5.3.12.2.1 Failure condition ordering

The RSI_PLANE_ENTER command does not have any failure condition orderings.

B5.3.12.3 Success conditions

ID	Condition
plane_exit	run.exit contains Plane exit syndrome information.

B5.3.12.4 Footprint

The RSI_PLANE_ENTER command does not have any footprint.

DRAFT

B5.3.13 RSI_PLANE_REG_READ command

Read a Plane register.

See also:

- [A10.2.6 Pn system registers](#)

B5.3.13.1 Interface

B5.3.13.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001AE
plane_idx	X1	63:0	UInt64	Index of target Plane
encoding	X2	63:0	Bits64	Encoding of target register

The encoding value is an architecturally-defined system register encoding.

See also:

- [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#)

B5.3.13.1.2 Context

The RSI_PLANE_REG_READ command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm

B5.3.13.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
value	X1	63:0	Bits64	Value of target register

B5.3.13.2 Failure conditions

ID	Condition
idx_bound	pre: plane_idx > realm.num_aux_planes post: result == RSI_ERROR_INPUT
reg_valid	pre: ! PlaneRegIsValid (realm, encoding) post: result == RSI_ERROR_INPUT

B5.3.13.2.1 Failure condition ordering

The RSI_PLANE_REG_READ command does not have any failure condition orderings.

B5.3.13.3 Success conditions

ID	Condition
value	value == <code>PlaneRegValue</code> (realm, plane_idx, encoding)

B5.3.13.4 Footprint

The RSI_PLANE_REG_READ command does not have any footprint.

DRAFT

B5.3.14 RSI_PLANE_REG_WRITE command

Write a Plane register.

See also:

- [A10.2.6 Pn system registers](#)

B5.3.14.1 Interface

B5.3.14.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001AF
plane_idx	X1	63:0	UInt64	Index of target Plane
encoding	X2	63:0	Bits64	Encoding of target register
value	X3	63:0	Bits64	Value to write to target register

The encoding value is an architecturally-defined system register encoding.

See also:

- [Arm Architecture Reference Manual for A-Profile architecture \[3\]](#)

B5.3.14.1.2 Context

The RSI_PLANE_REG_WRITE command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm

B5.3.14.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

B5.3.14.2 Failure conditions

ID	Condition
idx_bound	pre: plane_idx > realm.num_aux_planes post: result == RSI_ERROR_INPUT
reg_valid	pre: !PlaneRegIsValid (realm, encoding) post: result == RSI_ERROR_INPUT

B5.3.14.2.1 Failure condition ordering

The RSI_PLANE_REG_WRITE command does not have any failure condition orderings.

B5.3.14.3 Success conditions

ID	Condition
value	<code>PlaneRegValue(realm, plane_idx, encoding) == value</code>

B5.3.14.4 Footprint

The RSI_PLANE_REG_WRITE command does not have any footprint.

DRAFT

B5.3.15 RSI_RDEV_CONTINUE command

Continue an interruptible Realm device operation.

See also:

- [A9.5.1 Interruptible Realm device operations](#)

B5.3.15.1 Interface

B5.3.15.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A4
vdev_id	X1	63:0	Bits64	Realm device identifier
inst_id	X2	63:0	UInt64	Device instance identifier

B5.3.15.1.2 Context

The RSI_RDEV_CONTINUE command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rdev	RmmRdev	RdevFromIds(realm, vdev_id, inst_id)	false	Realm device

B5.3.15.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status


B5.3.15.2 Failure conditions

ID	Condition
da_en	pre: realm.feat_da != FEATURE_TRUE post: result == RSI_ERROR_STATE
vdev_id	pre: !RdevIdsAreValid(realm, vdev_id, inst_id) post: result == RSI_ERROR_INPUT
state	pre: (rdev.state != RDEV_NEW_BUSY && rdev.state != RDEV_LOCKED_BUSY && rdev.state != RDEV_STARTED_BUSY && rdev.state != RDEV_STOPPING) post: result == RSI_ERROR_INPUT

B5.3.15.2.1 Failure condition ordering

[da_en] < [vdev_id]
[vdev_id] < [state]

DRAFT



B5.3.15.3 Success conditions

ID	Condition
error	pre: (DeviceCommunicate(rdev) == DEV_COMM_ERROR && rdev.state != RDEV_STOPPING) post: rdev.state == RDEV_ERROR
new	pre: (DeviceCommunicate(rdev) == DEV_COMM_IDLE && rdev.state == RDEV_NEW_BUSY && rdev.operation != RDEV_OP_LOCK) post: rdev.state == RDEV_NEW
to_locked	pre: (DeviceCommunicate(rdev) == DEV_COMM_IDLE && rdev.state == RDEV_NEW_BUSY && rdev.operation == RDEV_OP_LOCK) post: rdev.state == RDEV_LOCKED
locked	pre: (DeviceCommunicate(rdev) == DEV_COMM_IDLE && rdev.state == RDEV_LOCKED_BUSY && rdev.operation != RDEV_OP_LOCK) post: rdev.state == RDEV_LOCKED

ID	Condition
to_started	<pre>pre: (DeviceCommunicate(rdev) == DEV_COMM_IDLE && rdev.state == RDEV_LOCKED_BUSY && rdev.operation == RDEV_OP_START) post: rdev.state == RDEV_STARTED</pre>
started	<pre>pre: (DeviceCommunicate(rdev) == DEV_COMM_IDLE && rdev.state == RDEV_STARTED_BUSY) post: rdev.state == RDEV_STARTED</pre>
stopped	<pre>pre: (DeviceCommunicate(rdev) != DEV_COMM_ACTIVE && rdev.state == RDEV_STOPPING) post: rdev.state == RDEV_STOPPED</pre>

B5.3.15.4 Footprint

ID	Value
state	rdev.state

DRAFT

B5.3.16 RSI_RDEV_GET_INFO command

Get information for a device.

Device configuration information, including digests of attestation evidence for the device are written to an `RsiDeviceInfo` structure, at an address specified by the caller. Digests are calculated using the PDEV Hash Algorithm.

See also:

- [A9.5.2 Realm retrieval of device attestation evidence](#)
- [B5.3.24 RSI_REALM_CONFIG command](#)
- [B5.4.3 RsiDeviceInfo type](#)

B5.3.16.1 Interface

B5.3.16.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A5
vdev_id	X1	63:0	Bits64	Realm device identifier
addr	X2	63:0	Address	IPA of the Granule to which the configuration data will be written

B5.3.16.1.2 Context

The `RSI_RDEV_GET_INFO` command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	<code>CurrentRealm()</code>	false	Current Realm
rdev	RmmRdev	<code>RdevFromId(realm, vdev_id)</code>	false	Realm device
vdev	RmmVdev	<code>VdevAt(rdev.vdev_ptr)</code>	false	Virtual device
pdev	RmmPdev	<code>PdevAt(vdev.pdev)</code>	false	Physical device
cfg	RsiDeviceInfo	<code>RsiDeviceInfoAt(addr)</code>	false	Device configuration

B5.3.16.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status


B5.3.16.2 Failure conditions

ID	Condition
da_en	pre: <code>realm.feats_da != FEATURE_TRUE</code> post: <code>result == RSI_ERROR_STATE</code>

ID	Condition
vdev_id	pre: !RdevIdIsValid(realm, vdev_id) post: result == RSI_ERROR_INPUT
addr_align	pre: !AddrIsGranuleAligned(addr) post: result == RSI_ERROR_INPUT
addr_bound	pre: !AddrIsProtected(addr, realm) post: result == RSI_ERROR_INPUT

B5.3.16.2.1 Failure condition ordering

[da_en] < [vdev_id, addr_align, addr_bound]



B5.3.16.3 Success conditions

ID	Condition
hash_algo	Equal(cfg.hash_algo, pdev.hash_algo)

B5.3.16.4 Footprint

ID	Value
state	rdev.state
operation	rdev.operation

B5.3.17 RSI_RDEV_GET_INTERFACE_REPORT command

Get Realm device interface report.

See also:

- [A9.5.2 Realm retrieval of device attestation evidence](#)

B5.3.17.1 Interface

B5.3.17.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A6
vdev_id	X1	63:0	Bits64	Realm device identifier
inst_id	X2	63:0	UInt64	Device instance identifier
version_max	X3	63:0	UInt64	Maximum TDISP version accepted by caller

B5.3.17.1.2 Context

The RSI_RDEV_GET_INTERFACE_REPORT command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rdev	RmmRdev	RdevFromIds (realm, vdev_id, inst_id)	false	Realm device

B5.3.17.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
version	X1	63:0	UInt64	TDISP version


B5.3.17.2 Failure conditions

ID	Condition
da_en	pre: realm.feats_da != FEATURE_TRUE post: result == RSI_ERROR_STATE
vdev_id	pre: !RdevIdsAreValid(realm, vdev_id, inst_id) post: result == RSI_ERROR_INPUT
version	pre: TDISP version is not supported. post: result == RSI_ERROR_INPUT

ID	Condition
state	<pre>pre: (rdev.state != RDEV_LOCKED && rdev.state != RDEV_STARTED) post: result == RSI_ERROR_INPUT</pre>

B5.3.17.2.1 Failure condition ordering

<pre>[da_en] < [vdev_id, version] [vdev_id] < [state]</pre>



B5.3.17.3 Success conditions

ID	Condition
locked	<pre>pre: rdev.state == RDEV_LOCKED post: rdev.state == RDEV_LOCKED_BUSY</pre>
started	<pre>pre: rdev.state == RDEV_STARTED post: rdev.state == RDEV_STARTED_BUSY</pre>
operation	<pre>rdev.operation == RDEV_OP_GET_INTERFACE_REPORT</pre>

B5.3.17.4 Footprint

ID	Value
state	rdev.state
operation	rdev.operation

DRAFT

B5.3.18 RSI_RDEV_GET_MEASUREMENTS command

Get Realm device measurements.

See also:

- [A9.5.2 Realm retrieval of device attestation evidence](#)

B5.3.18.1 Interface

B5.3.18.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A7
vdev_id	X1	63:0	Bits64	Virtual device identifier
inst_id	X2	63:0	UInt64	Device instance identifier
params_ptr	X3	63:0	Address	IPA of measurement parameters

B5.3.18.1.2 Context

The RSI_RDEV_GET_MEASUREMENTS command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rdev	RmmRdev	RdevFromIds (realm, vdev_id, inst_id)	false	Realm device
params	RsiDeviceMeasurements	RsiDeviceMeasParamsAt (params_ptr)	false	Measurement parameters

B5.3.18.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

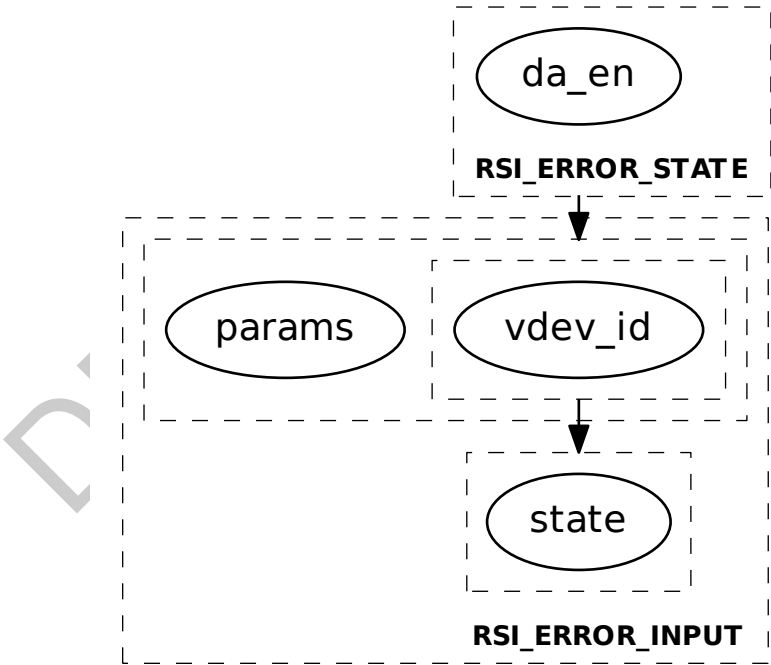
B5.3.18.2 Failure conditions

ID	Condition
da_en	pre: realm.featur_da != FEATURE_TRUE post: result == RSI_ERROR_STATE
vdev_id	pre: !RdevIdsAreValid (realm, vdev_id, inst_id) post: result == RSI_ERROR_INPUT
params	pre: !RdevMeasurementParamsValid (params) post: result == RSI_ERROR_INPUT

ID	Condition
state	<pre>pre: (rdev.state != RDEV_NEW && rdev.state != RDEV_LOCKED && rdev.state != RDEV_STARTED) post: result == RSI_ERROR_INPUT</pre>

B5.3.18.2.1 Failure condition ordering

[da_en] < [vdev_id, params]
[vdev_id] < [state]



B5.3.18.3 Success conditions

ID	Condition
new	<pre>pre: rdev.state == RDEV_NEW post: rdev.state == RDEV_NEW_BUSY</pre>
locked	<pre>pre: rdev.state == RDEV_LOCKED post: rdev.state == RDEV_LOCKED_BUSY</pre>
started	<pre>pre: rdev.state == RDEV_STARTED post: rdev.state == RDEV_STARTED_BUSY</pre>
operation	<pre>rdev.operation == RDEV_OP_GET_MEASUREMENTS</pre>

B5.3.18.4 Footprint

ID	Value
state	rdev.state
operation	rdev.operation

DRAFT

B5.3.19 RSI_RDEV_GET_STATE command

Get state of a Realm device.

See also:

- [A9.5.5 Realm device lifecycle](#)

B5.3.19.1 Interface

B5.3.19.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A8
vdev_id	X1	63:0	Bits64	Realm device identifier
inst_id	X2	63:0	UInt64	Device instance identifier

B5.3.19.1.2 Context

The RSI_RDEV_GET_STATE command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rdev	RmmRdev	RdevFromIds (realm, vdev_id, inst_id)	false	Realm device

B5.3.19.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
state	X1	63:0	RsiDeviceState	Realm device state

B5.3.19.2 Failure conditions

ID	Condition
da_en	pre: realm.feats_da != FEATURE_TRUE post: result == RSI_ERROR_STATE
vdev_id	pre: !RdevIdsAreValid (realm, vdev_id, inst_id) post: result == RSI_ERROR_INPUT

B5.3.19.2.1 Failure condition ordering

[da_en] < [vdev_id]



B5.3.19.3 Success conditions

ID	Condition
state	<code>Equal(state, rdev.state)</code>

B5.3.19.4 Footprint

The RSI_RDEV_GET_STATE command does not have any footprint.

B5.3.20 RSI_RDEV_LOCK command

Lock a Realm device.

See also:

- [A9.5.5 Realm device lifecycle](#)

B5.3.20.1 Interface

B5.3.20.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001A9
vdev_id	X1	63:0	Bits64	Realm device identifier
inst_id	X2	63:0	UInt64	Device instance identifier

B5.3.20.1.2 Context

The RSI_RDEV_LOCK command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rdev	RmmRdev	RdevFromIds (realm, vdev_id, inst_id)	false	Realm device

B5.3.20.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

B5.3.20.2 Failure conditions

ID	Condition
da_en	pre: realm.feats_da != FEATURE_TRUE post: result == RSI_ERROR_STATE
vdev_id	pre: !RdevIdsAreValid (realm, vdev_id, inst_id) post: result == RSI_ERROR_INPUT
state	pre: rdev.state != RDEV_NEW post: result == RSI_ERROR_INPUT

B5.3.20.2.1 Failure condition ordering

[da_en] < [vdev_id]
[vdev_id] < [state]



B5.3.20.3 Success conditions

ID	Condition
state	rdev.state == RDEV_NEW_BUSY
operation	rdev.operation == RDEV_OP_LOCK

B5.3.20.4 Footprint

ID	Value
state	rdev.state
operation	rdev.operation

B5.3.21 RSI_RDEV_START command

Start a Realm device.

See also:

- [A9.5.5 Realm device lifecycle](#)

B5.3.21.1 Interface

B5.3.21.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001AA
vdev_id	X1	63:0	Bits64	Realm device identifier
inst_id	X2	63:0	UInt64	Device instance identifier

B5.3.21.1.2 Context

The RSI_RDEV_START command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rdev	RmmRdev	RdevFromIds (realm, vdev_id, inst_id)	false	Realm device

B5.3.21.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

B5.3.21.2 Failure conditions

ID	Condition
da_en	pre: realm.feat_da != FEATURE_TRUE post: result == RSI_ERROR_STATE
vdev_id	pre: !RdevIdsAreValid (realm, vdev_id, inst_id) post: result == RSI_ERROR_INPUT
state	pre: rdev.state != RDEV_LOCKED post: result == RSI_ERROR_INPUT

B5.3.21.2.1 Failure condition ordering

[da_en] < [vdev_id]
[vdev_id] < [state]



B5.3.21.3 Success conditions

ID	Condition
state	rdev.state == RDEV_LOCKED_BUSY
operation	rdev.operation == RDEV_OP_START

B5.3.21.4 Footprint

ID	Value
state	rdev.state
operation	rdev.operation

B5.3.22 RSI_RDEV_STOP command

Stop a Realm device.

See also:

- [A9.5.5 Realm device lifecycle](#)

B5.3.22.1 Interface

B5.3.22.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001AB
vdev_id	X1	63:0	Bits64	Realm device identifier
inst_id	X2	63:0	UInt64	Device instance identifier

B5.3.22.1.2 Context

The RSI_RDEV_STOP command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rdev	RmmRdev	RdevFromIds (realm, vdev_id, inst_id)	false	Realm device

B5.3.22.1.3 Output values


Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

B5.3.22.2 Failure conditions

ID	Condition
da_en	pre: realm.feat_da != FEATURE_TRUE post: result == RSI_ERROR_STATE
vdev_id	pre: !RdevIdsAreValid (realm, vdev_id, inst_id) post: result == RSI_ERROR_INPUT
state	pre: (rdev.state != RDEV_NEW && rdev.state != RDEV_LOCKED && rdev.state != RDEV_STARTED && rdev.state != RDEV_ERROR) post: result == RSI_ERROR_INPUT

B5.3.22.2.1 Failure condition ordering

[da_en] < [vdev_id]
[vdev_id] < [state]



B5.3.22.3 Success conditions

ID	Condition
state	rdev.state == RDEV_STOPPING

B5.3.22.4 Footprint

ID	Value
state	rdev.state

B5.3.23 RSI_RDEV_VALIDATE_MAPPING command

Validate Realm device memory mappings.

See also:

- [A9.5.3 Realm validation of device memory mappings](#)

B5.3.23.1 Interface

B5.3.23.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001AC
vdev_id	X1	63:0	Bits64	Realm device identifier
inst_id	X2	63:0	UInt64	Device instance identifier
ipa_base	X3	63:0	Address	Base of target IPA region
ipa_top	X4	63:0	Address	Top of target IPA region
pa_base	X5	63:0	Address	Base of target PA region
flags	X6	63:0	RsiRdevValidateIoFlags	Flags

B5.3.23.1.2 Context

The RSI_RDEV_VALIDATE_MAPPING command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
rec	RmmRec	CurrentRec()	false	Current REC
rdev	RmmRdev	RdevFromIds() realm, vdev_id, inst_id)	false	Realm device

B5.3.23.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
new_ipa_base	X1	63:0	Address	Base of IPA region which was not modified by the command
response	X2	0:0	RsiResponse	Whether the Host accepted or rejected the request

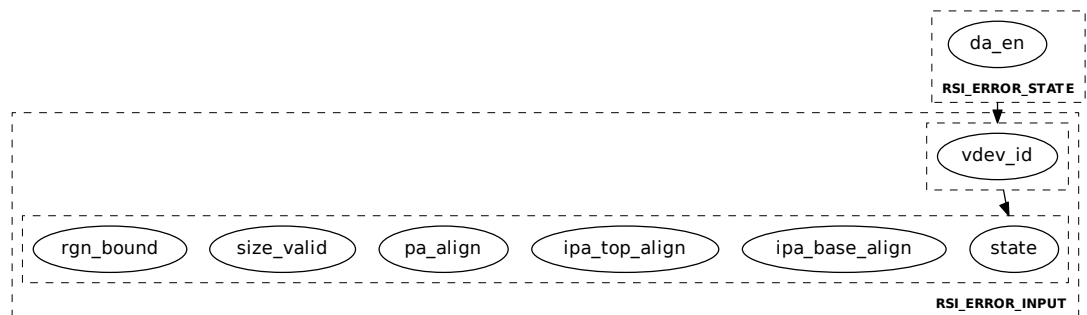
The following unused bits of RSI_RDEV_VALIDATE_MAPPING output values MBZ: X2[63:1].

B5.3.23.2 Failure conditions

ID	Condition
da_en	pre: realm.feats_da != FEATURE_TRUE post: result == RSI_ERROR_STATE
vdev_id	pre: !RdevIdsAreValid(realm, vdev_id, inst_id) post: result == RSI_ERROR_INPUT
state	pre: (rdev.state != RDEV_LOCKED && rdev.state != RDEV_STARTED) post: result == RSI_ERROR_INPUT
ipa_base_align	pre: !AddrIsGranuleAligned(ipa_base) post: result == RSI_ERROR_INPUT
ipa_top_align	pre: !AddrIsGranuleAligned(ipa_top) post: result == RSI_ERROR_INPUT
pa_align	pre: !AddrIsGranuleAligned(pa_base) post: result == RSI_ERROR_INPUT
size_valid	pre: UInt(ipa_top) <= UInt(ipa_base) post: result == RSI_ERROR_INPUT
rgn_bound	pre: !AddrRangeIsProtected(ipa_base, ipa_top, realm) post: result == RSI_ERROR_INPUT

B5.3.23.2.1 Failure condition ordering

```
[da_en] < [vdev_id]
[vdev_id] < [state, ipa_base_align, ipa_top_align, pa_align,
             size_valid, rgn_bound]
```



B5.3.23.3 Success conditions

ID	Condition
new_ipa_base	new_ipa_base == rec.ripas_addr
response	response == RecRipasResponseToRsi(rec)

B5.3.23.4 Footprint

The RSI_RDEV_VALIDATE_MAPPING command does not have any footprint.

B5.3.24 RSI_REALM_CONFIG command

Read configuration for the current Realm.

B5.3.24.1 Interface

B5.3.24.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC40001AD
addr	X1	63:0	Address	IPA of the Granule to which the configuration data will be written

B5.3.24.1.2 Context

The RSI_REALM_CONFIG command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm
cfg	RsiRealmConfig	RsiRealmConfigAt(addr)	false	Realm configuration

B5.3.24.1.3 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status

B5.3.24.2 Failure conditions

ID	Condition
addr_align	pre: !AddrIsGranuleAligned(addr) post: result == RSI_ERROR_INPUT
addr_bound	pre: !AddrIsProtected(addr, realm) post: result == RSI_ERROR_INPUT

B5.3.24.2.1 Failure condition ordering

The RSI_REALM_CONFIG command does not have any failure condition orderings.

ID	Condition
----	-----------

B5.3.24.3 Success conditions

ID	Condition
ipa_width	<code>cfg.ipa_width == realm.ipa_width</code>
hash_algo	<code>Equal(cfg.hash_algo, realm.hash_algo)</code>
num_aux_planes	<code>cfg.num_aux_planes == realm.num_aux_planes</code>

B5.3.24.4 Footprint

The RSI_REALM_CONFIG command does not have any footprint.

DRAFT

B5.3.25 RSI_VERSION command

Returns RSI version.

On calling this command, the Realm provides a requested RSI version.

The output values include a status code and two revisions which are supported by the RMM: a *lower revision* and a *higher revision*.

- The *higher revision* value is the highest interface revision which is supported by the RMM.
- The *lower revision* is less than or equal to the *higher revision*.

The status code and *lower revision* output values indicate which of the following is true, in order of precedence:

- a) The RMM supports an interface revision which is compatible with the requested revision.
 - The status code is RSI_SUCCESS.
 - The *lower revision* is equal to the requested revision.
- b) The RMM does not support an interface revision which is compatible with the requested revision. The RMM supports an interface revision which is incompatible with and less than the requested revision.
 - The status code is RSI_ERROR_INPUT.
 - The *lower revision* is the highest interface revision which is both less than the requested revision and supported by the RMM.
- c) The RMM does not support an interface revision which is compatible with the requested revision. The RMM supports an interface revision which is incompatible with and greater than the requested revision.
 - The status code is RSI_ERROR_INPUT.
 - The *lower revision* is equal to the *higher revision*.

See also:

- [Chapter B2 Interface versioning](#)
- [B5.1 RSI version](#)

B5.3.25.1 Interface

B5.3.25.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xc4000190
req	X1	63:0	RsiInterfaceVersion	Requested interface revision

B5.3.25.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	RsiCommandReturnCode	Command return status
lower	X1	63:0	RsiInterfaceVersion	Lower implemented interface revision
higher	X2	63:0	RsiInterfaceVersion	Higher implemented interface revision

B5.3.25.2 Failure conditions

The RSI_VERSION command does not have any failure conditions.

B5.3.25.3 Success conditions

The RSI_VERSION command does not have any success conditions.

B5.3.25.4 Footprint

The RSI_VERSION command does not have any footprint.

DRAFT

B5.4 RSI types

This section defines types which are used in the RSI interface.

B5.4.1 RsiBoolean type

The RsiBoolean enumeration represents a boolean value.

The RsiBoolean enumeration is a [concrete type](#).

The width of the RsiBoolean enumeration is 1 bits.

The values of the RsiBoolean enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_FALSE	False
1	RSI_TRUE	True

The RsiBoolean enumeration is used in the following types:

- [RsiDeviceMeasurementsParams](#)

B5.4.2 RsiCommandReturnCode type

The RsiCommandReturnCode enumeration represents a return code from an RSI command.

The RsiCommandReturnCode enumeration is a [concrete type](#).

The width of the RsiCommandReturnCode enumeration is 64 bits.

See also:

- [Chapter B1 Commands](#)

The values of the RsiCommandReturnCode enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_SUCCESS	Command completed successfully
1	RSI_ERROR_INPUT	The value of a command input value caused the command to fail
2	RSI_ERROR_STATE	The state of the current Realm or current REC does not match the state expected by the command
3	RSI_INCOMPLETE	The operation requested by the command is not complete
4	RSI_ERROR_UNKNOWN	The operation requested by the command failed for an unknown reason
5	RSI_ERROR_DEVICE	The state of a Realm device does not match the state expected by the command

Unused encodings for the RsiCommandReturnCode enumeration are reserved for use by future versions of this specification.

B5.4.3 RsiDeviceInfo type

The RsiDeviceInfo structure contains device configuration information.

The RsiDeviceInfo structure is a [concrete type](#).

The width of the RsiDeviceInfo structure is 512 (0x200) bytes.

See also:

- [A9.5 Realm management of an assigned device interface](#)
- [B5.3.16 RSI_RDEV_GET_INFO command](#)

The members of the RsiDeviceInfo structure are shown in the following table.

Name	Byte offset	Type	Description
inst_id	0x0	UInt64	Instance identifier
cert_id	0x8	UInt64	Certificate identifier
hash_algo	0x10	RsiHashAlgorithm	Algorithm used to generate device digests
cert_digest	0x40	Bits512	Certificate digest
key_digest	0x80	Bits512	Device public key digest
meas_digest	0xc0	Bits512	Measurement block digest
report_digest	0x100	Bits512	Interface report digest

Unused bits of the RsiDeviceInfo structure MBZ.

B5.4.4 RsiDeviceMeasurementsParams type

The RsiDeviceMeasurementsParams structure contains parameters for retrieval of Realm device measurements.

The RsiDeviceMeasurementsParams structure is a [concrete type](#).

The width of the RsiDeviceMeasurementsParams structure is 64 (0x40) bytes.

The members of the RsiDeviceMeasurementsParams structure are shown in the following table.

Name	Byte offset	Type	Description
meas_ids[256]	0x0	RsiBoolean	Measurement indices For each index, if the array element is RSI_TRUE then the RMM is requested to retrieve the corresponding device measurement.
meas_params[256]	0x20	RsiBoolean	Measurement parameters

B5.4.5 RsiDeviceState type

The RsiDeviceState enumeration represents state of an assigned Realm device.

The RsiDeviceState enumeration is a [concrete type](#).

The width of the RsiDeviceState enumeration is 64 bits.

See also:

- [A9.5 Realm management of an assigned device interface](#)

The values of the RsiDeviceState enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_RDEV_NEW	Device interface is unlocked.
1	RSI_RDEV_NEW_BUSY	Device interface is unlocked and is handling an interruptible Realm device operation.
2	RSI_RDEV_LOCKED	Device interface is locked.
3	RSI_RDEV_LOCKED_BUSY	Device interface is locked and is handling an interruptible Realm device operation.
4	RSI_RDEV_STARTED	Device interface is started.
5	RSI_RDEV_STARTED_BUSY	Device interface is started and is handling an interruptible Realm device operation.
6	RSI_RDEV_STOPPING	Device interface is stopping.
7	RSI_RDEV_STOPPED	Device interface is stopped.
8	RSI_RDEV_ERROR	Device interface has reported a fatal error.

Unused encodings for the RsiDeviceState enumeration are reserved for use by future versions of this specification.

B5.4.6 RsiDevMemCoherent type

The RsiDevMemCoherent enumeration represents whether a device memory location is within the system coherent memory space.

The RsiDevMemCoherent enumeration is a [concrete type](#).

The width of the RsiDevMemCoherent enumeration is 1 bits.

The values of the RsiDevMemCoherent enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_DEV_MEM_NON_COHERENT	A device memory location is not within the system coherent memory space
1	RSI_DEV_MEM_COHERENT	A device memory location is within the system coherent memory space

The RsiDevMemCoherent enumeration is used in the following types:

- [RsiRdevValidateIoFlags](#)

B5.4.7 RsiDevMemShared type

The RsiDevMemShared enumeration represents whether a device memory mapping is shared.

The RsiDevMemShared enumeration is a [concrete type](#).

The width of the RsiDevMemShared enumeration is 1 bits.

The values of the RsiDevMemShared enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_DEV_MEM_PRIVATE	Device memory mapping is private
1	RSI_DEV_MEM_SHARED	Device memory mapping is shared

The RsiDevMemShared enumeration is used in the following types:

- [RsiRdevValidateIoFlags](#)

B5.4.8 RsiFeature type

The RsiFeature enumeration represents whether a feature is enabled.

The RsiFeature enumeration is a [concrete type](#).

The width of the RsiFeature enumeration is 1 bits.

The values of the RsiFeature enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_FEATURE_FALSE	Feature is not enabled.
1	RSI_FEATURE_TRUE	Feature is enabled.

The RsiFeature enumeration is used in the following types:

- [RsiFeatureRegister0](#)


B5.4.9 RsiFeatureRegister0 type

The RsiFeatureRegister0 fieldset contains RSI feature register 0.

The RsiFeatureRegister0 fieldset is a [concrete type](#).

The width of the RsiFeatureRegister0 fieldset is 64 bits.

The fields of the RsiFeatureRegister0 fieldset are shown in the following diagram.



The fields of the RsiFeatureRegister0 fieldset are shown in the following table.

Name	Bits	Description	Value
DA	0	Whether Realm device assignment is supported	RsiFeature
MRO	1	Whether “mostly read-only” permissions are supported	RsiFeature
	63:2	Reserved	MBZ

B5.4.10 RsiGicOwner type

The RsiGicOwner enumeration represents which Plane is GIC owner.

The RsiGicOwner enumeration is a [concrete type](#).

The width of the RsiGicOwner enumeration is 1 bits.

The values of the RsiGicOwner enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_GIC_OWNER_0	Plane 0 is GIC owner.
1	RSI_GIC_OWNER_N	Plane N is GIC owner.

The RsiGicOwner enumeration is used in the following types:

- [RsiPlaneEnterFlags](#)

B5.4.11 RsiHashAlgorithm type

The RsiHashAlgorithm enumeration represents hash algorithm.

The RsiHashAlgorithm enumeration is a [concrete type](#).

The width of the RsiHashAlgorithm enumeration is 8 bits.

See also:

- [B5.3.24 RSI_REALM_CONFIG command](#)

The values of the RsiHashAlgorithm enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_HASH_SHA_256	SHA-256 (Secure Hash Standard (SHS) [19])
1	RSI_HASH_SHA_512	SHA-512 (Secure Hash Standard (SHS) [19])

Unused encodings for the RsiHashAlgorithm enumeration are reserved for use by future versions of this specification.

The RsiHashAlgorithm enumeration is used in the following types:

- [RsiDeviceInfo](#)
- [RsiRealmConfig](#)

B5.4.12 RsiHostCall type

The RsiHostCall structure contains data structure used to pass Host call arguments and return values.

The RsiHostCall structure is a [concrete type](#).

The width of the RsiHostCall structure is 256 (0x100) bytes.

See also:

- [A4.5 Host call](#)
- [B5.3.4 RSI_HOST_CALL command](#)

The members of the RsiHostCall structure are shown in the following table.

Name	Byte offset	Type	Description
imm	0x0	UInt16	Immediate value
gprs[31]	0x8	Bits64	Registers

Unused bits of the RsiHostCall structure SBZ.

B5.4.13 RsiInterfaceVersion type

The RsiInterfaceVersion fieldset contains an RSI interface version.


The RsiInterfaceVersion fieldset is a [concrete type](#).

The width of the RsiInterfaceVersion fieldset is 64 bits.

See also:

- [B5.1 RSI version](#)
- [B5.3.25 RSI_VERSION command](#)

The fields of the RsiInterfaceVersion fieldset are shown in the following diagram.



The fields of the RsiInterfaceVersion fieldset are shown in the following table.

Name	Bits	Description	Value
minor	15:0	Interface minor version number (the value <i>y</i> in interface version <i>x.y</i>)	UInt16
major	30:16	Interface major version number (the value <i>x</i> in interface version <i>x.y</i>)	UInt15
	63:31	Reserved	SBZ

B5.4.14 RsiPlaneEnter type

The RsiPlaneEnter structure contains data passed from P0 to the RMM on Plane entry.

The RsiPlaneEnter structure is a [concrete type](#).

The width of the RsiPlaneEnter structure is 2048 (0x800) bytes.

The members of the RsiPlaneEnter structure are shown in the following table.

Name	Byte offset	Type	Description
flags	0x0	RsiPlaneEnterFlags	Flags
pc	0x8	Bits64	Program counter
gprs[31]	0x100	Bits64	Registers
gicv3_hcr	0x200	Bits64	GICv3 Hypervisor Control Register value
gicv3_lrs[16]	0x208	Bits64	GICv3 List Register values

Unused bits of the RsiPlaneEnter structure SBZ.

The RsiPlaneEnter structure is used in the following types:

- [RsiPlaneRun](#)


B5.4.15 RsiPlaneEnterFlags type

The RsiPlaneEnterFlags fieldset contains flags provided by P0 during Plane entry.

The RsiPlaneEnterFlags fieldset is a [concrete type](#).

The width of the RsiPlaneEnterFlags fieldset is 64 bits.

The fields of the RsiPlaneEnterFlags fieldset are shown in the following diagram.



The fields of the RsiPlaneEnterFlags fieldset are shown in the following table.

Name	Bits	Description	Value
trap_wfi	0	Whether to trap WFI execution by the Plane.	RsiTrap
trap_wfe	1	Whether to trap WFE execution by the Plane.	RsiTrap

Name	Bits	Description	Value
trap_hc	2	Whether to trap RSI_HOST_CALL execution by the Plane. RSI_TRAP: execution of RSI_HOST_CALL causes Plane exit RSI_NO_TRAP: execution of RSI_HOST_CALL causes REC exit to Host	RsiTrap
gic_owner	3	Whether to transfer GIC ownership to the target Plane.	RsiGicOwner
	63:4	Reserved	SBZ

The RsiPlaneEnterFlags fieldset is used in the following types:

- [RsiPlaneEnter](#)

B5.4.16 RsiPlaneExit type

The RsiPlaneExit structure contains data passed from the RMM to P0 on Plane exit.

The RsiPlaneExit structure is a [concrete type](#).

The width of the RsiPlaneExit structure is 2048 (0x800) bytes.

The members of the RsiPlaneExit structure are shown in the following table.

Name	Byte offset	Type	Description
reason	0x0	RsiPlaneExitReason	Exit reason
elr_el2	0x100	Bits64	Exception Link Register
esr_el2	0x108	Bits64	Exception Syndrome Register
far_el2	0x110	Bits64	Fault Address Register
hpfar_el2	0x118	Bits64	Hypervisor IPA Fault Address register
gprs[31]	0x200	Bits64	Registers
gicv3_hcr	0x300	Bits64	GICv3 Hypervisor Control Register value
gicv3_lrs[16]	0x308	Bits64	GICv3 List Register values
gicv3_misr	0x388	Bits64	GICv3 Maintenance Interrupt State Register value
gicv3_vmcr	0x390	Bits64	GICv3 Virtual Machine Control Register value
cntp_ctl	0x400	Bits64	Counter-timer Physical Timer Control Register value
cntp_cval	0x408	Bits64	Counter-timer Physical Timer CompareValue Register value
cntv_ctl	0x410	Bits64	Counter-timer Virtual Timer Control Register value

Name	Byte offset	Type	Description
cntv_cval	0x418	Bits64	Counter-timer Virtual Timer CompareValue Register value

Unused bits of the RsiPlaneExit structure SBZ.

The RsiPlaneExit structure is used in the following types:

- [RsiPlaneRun](#)

B5.4.17 RsiPlaneExitReason type

The RsiPlaneExitReason enumeration represents the reason for a Plane exit.

The RsiPlaneExitReason enumeration is a [concrete type](#).

The width of the RsiPlaneExitReason enumeration is 8 bits.

The values of the RsiPlaneExitReason enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_EXIT_SYNC	Plane exit due to synchronous exception

Unused encodings for the RsiPlaneExitReason enumeration are reserved for use by future versions of this specification.

The RsiPlaneExitReason enumeration is used in the following types:

- [RsiPlaneExit](#)

B5.4.18 RsiPlaneRun type

The RsiPlaneRun structure contains fields used to share information between RMM and P0 during Plane entry and Plane exit.

The RsiPlaneRun structure is a [concrete type](#).

The width of the RsiPlaneRun structure is 4096 (0x1000) bytes.

The members of the RsiPlaneRun structure are shown in the following table.

Name	Byte offset	Type	Description
enter	0x0	RsiPlaneEnter	Entry information
exit	0x800	RsiPlaneExit	Exit information


B5.4.19 RsiRdevValidateIoFlags type

The RsiRdevValidateIoFlags fieldset contains flags provided when requesting validation of a device memory mapping.

The RsiRdevValidateIoFlags fieldset is a [concrete type](#).

The width of the RsiRdevValidateIoFlags fieldset is 64 bits.

The fields of the RsiRdevValidateIoFlags fieldset are shown in the following diagram.



The fields of the RsiRdevValidateIoFlags fieldset are shown in the following table.

Name	Bits	Description	Value
share	0	Whether the device memory mapping is shared.	RsiDevMemShared
coh	1	Whether the output address of the device memory mapping is within the system coherent memory space.	RsiDevMemCoherent
	63:2	Reserved	SBZ

B5.4.20 RsiRealmConfig type

The RsiRealmConfig structure contains realm configuration.

The RsiRealmConfig structure is a [concrete type](#).

The width of the RsiRealmConfig structure is 4096 (0x1000) bytes.

See also:

- [B5.3.24 RSI_REALM_CONFIG command](#)

The members of the RsiRealmConfig structure are shown in the following table.

Name	Byte offset	Type	Description
ipa_width	0x0	UInt64	IPA width in bits
hash_algo	0x8	RsiHashAlgorithm	Hash algorithm
num_aux_planes	0x10	UInt64	Number of auxiliary Planes
gicv3_vtr	0x18	Bits64	GICv3 VGIC Type Register value
rpv	0x200	Bits512	Realm Personalization Value

Unused bits of the RsiRealmConfig structure MBZ.

B5.4.21 RsiResponse type

The RsiResponse enumeration represents whether the Host accepted or rejected a Realm request.

The RsiResponse enumeration is a [concrete type](#).

The width of the RsiResponse enumeration is 1 bits.

The values of the RsiResponse enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_ACCEPT	Host accepted the Realm request.
1	RSI_REJECT	Host rejected the Realm request.

B5.4.22 RsiRipas type

The RsiRipas enumeration represents realm IPA state.

The RsiRipas enumeration is a [concrete type](#).

The width of the RsiRipas enumeration is 8 bits.

See also:

- [A5.4 RIPAS change](#)
- [B5.3.5 RSI_IPA_STATE_GET command](#)
- [B5.3.6 RSI_IPA_STATE_SET command](#)

The values of the RsiRipas enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_EMPTY	Address where no Realm resources are mapped.
1	RSI_RAM	Address where private code or data owned by the Realm is mapped.
2	RSI_DESTROYED	Address which is inaccessible to the Realm due to an action taken by the Host.
3	RSI_DEV	Address where memory of an assigned Realm device is mapped.

Unused encodings for the RsiRipas enumeration are reserved for use by future versions of this specification.

B5.4.23 RsiRipasChangeDestroyed type

The RsiRipasChangeDestroyed enumeration represents whether a RIPAS change from DESTROYED should be permitted.

The RsiRipasChangeDestroyed enumeration is a [concrete type](#).

The width of the RsiRipasChangeDestroyed enumeration is 1 bits.

The values of the RsiRipasChangeDestroyed enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_NO_CHANGE_DESTROYED	A RIPAS change from DESTROYED should not be permitted.
1	RSI_CHANGE_DESTROYED	A RIPAS change from DESTROYED should be permitted.

The RsiRipasChangeDestroyed enumeration is used in the following types:

- [RsiRipasChangeFlags](#)


B5.4.24 RsiRipasChangeFlags type

The RsiRipasChangeFlags fieldset contains flags provided by the Realm when requesting a RIPAS change.

The RsiRipasChangeFlags fieldset is a [concrete type](#).

The width of the RsiRipasChangeFlags fieldset is 64 bits.

The fields of the RsiRipasChangeFlags fieldset are shown in the following diagram.



The fields of the RsiRipasChangeFlags fieldset are shown in the following table.

Name	Bits	Description	Value
destroyed	0	Whether a RIPAS change from DESTROYED should be permitted	RsiRipasChangeDestroyed
	63:1	Reserved	SBZ

B5.4.25 RsiTrap type

The RsiTrap enumeration represents whether a trap is enabled.

The RsiTrap enumeration is a [concrete type](#).

The width of the RsiTrap enumeration is 1 bits.

The values of the RsiTrap enumeration are shown in the following table.

Encoding	Name	Description
0	RSI_NO_TRAP	Trap is disabled.
1	RSI_TRAP	Trap is enabled.

The RsiTrap enumeration is used in the following types:

- [RsiPlaneEnterFlags](#)

DRAFT

Chapter B6

Power State Control Interface

This section describes how Power State Control Interface (PSCI) function execution by a Realm execution of SMC instructions is handled.

B6.1 PSCI overview

I_{GBVWX}

In this section,

- `rec` refers to the currently executing REC
- `exit` refer to the `RmiRecExit` object which was provided to the `RMI_REC_ENTER` command
- `target_rec` refers to the REC object identified by an MPIDR value passed to a PSCI function.

I_{GHKCJ}

The RMM provides a trusted implementation of parts of the PSCI ABI. This section describes the checks performed by the RMM when a Realm executes a PSCI command, and the internal RMM state changes which result from a successful PSCI command execution. Successful execution by the RMM of some PSCI commands results in a *REC exit due to PSCI*, which allows the Host to perform further processing of the command.

I_{XHDQF}

The HVC conduit for PSCI is not supported for Realms.

See also:

- [Arm Power State Coordination Interface \(PSCI\) \[22\]](#)
- [A2.3.2 REC attributes](#)
- [A4.3.7 REC exit due to PSCI](#)
- [A4.5 Host call](#)
- [D1.4 PSCI flows](#)

B6.2 PSCI version

R_{TFCVF}

The RMM must support version ≥ 1.1 of the Power State Control Interface.

See also:

- [B6.3.8 PSCI_VERSION command](#)

B6.3 PSCI commands

The following table summarizes the FIDs of commands in the PSCI interface.

FID	Command
0x84000000	PSCI_VERSION
...	
0x84000002	PSCI_CPU_OFF
...	
0x84000008	PSCI_SYSTEM_OFF
0x84000009	PSCI_SYSTEM_RESET
0x8400000A	PSCI_FEATURES
...	
0xC4000001	PSCI_CPU_SUSPEND
...	
0xC4000003	PSCI_CPU_ON
0xC4000004	PSCI_AFFINITY_INFO

B6.3.1 PSCI_AFFINITY_INFO command

Query status of a VPE.

This command causes a REC exit due to PSCI. In response, the Host should provide the target REC (identified by `target_affinity`) by calling `RMI_PSCI_COMPLETE`.

See also:

- [A2.3.2 REC attributes](#)
- [A4.3.7 REC exit due to PSCI](#)
- [B4.3.23 RMI_PSCI_COMPLETE command](#)
- [B6.3.2 PSCI_CPU_OFF command](#)
- [B6.3.3 PSCI_CPU_ON command](#)

B6.3.1.1 Interface

B6.3.1.1.1 Input values

Name	Register	Bits	Type	Description
<code>fid</code>	X0	63:0	UInt64	FID, value 0xC4000004
<code>target_affinity</code>	X1	63:0	Bits64	This parameter contains a copy of the affinity fields of the MPIDR register
<code>lowest_affinity_level</code>	X2	31:0	UInt32	Denotes the lowest affinity level field that is valid in the <code>target_affinity</code> parameter

The following unused bits of `PSCI_AFFINITY_INFO` input values SBZ: X2[63:32].

B6.3.1.1.2 Context

The `PSCI_AFFINITY_INFO` command operates on the following context.

Name	Type	Value	Before	Description
<code>target_rec</code>	RmmRec	<code>RecFromMpidr(target_affinity)</code>	false	Target REC

B6.3.1.1.3 Output values

Name	Register	Bits	Type	Description
<code>result</code>	X0	63:0	PsciReturnCode	Command return code

B6.3.1.2 Failure conditions

ID	Condition
<code>target_bound</code>	pre: <code>lowest_affinity_level != 0</code> post: <code>result == PSCI_INVALID_PARAMETERS</code>

ID	Condition
target_match	pre: !MpidrIsUsed(target_affinity) post: result == PSCI_INVALID_PARAMETERS

B6.3.1.2.1 Failure condition ordering

The PSCI_AFFINITY_INFO command does not have any failure condition orderings.

B6.3.1.3 Success conditions

ID	Condition
runnable	pre: target_rec.flags.runnable == RUNNABLE post: result == PSCI_SUCCESS
not_runnable	pre: target_rec.flags.runnable == NOT_RUNNABLE post: result == PSCI_OFF

B6.3.1.4 Footprint

The PSCI_AFFINITY_INFO command does not have any footprint.

B6.3.2 PSCI_CPU_OFF command

Power down the calling core.

This command causes a REC exit due to PSCI.

See also:

- [A2.3.2 REC attributes](#)
- [A4.3.7 REC exit due to PSCI](#)
- [B6.3.3 PSCI_CPU_ON command](#)
- [B6.3.4 PSCI_CPU_SUSPEND command](#)

B6.3.2.1 Interface

B6.3.2.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0x84000002

B6.3.2.1.2 Context

The PSCI_CPU_OFF command operates on the following context.

Name	Type	Value	Before	Description
rec	RmmRec	CurrentRec()	false	Current REC

B6.3.2.1.3 Output values

The PSCI_CPU_OFF command does not have any output values.

Following execution of PSCI_CPU_OFF, control does not return to the caller.

B6.3.2.2 Failure conditions

The PSCI_CPU_OFF command does not have any failure conditions.

B6.3.2.3 Success conditions

The PSCI_CPU_OFF command does not have any success conditions.

Following execution of PSCI_CPU_OFF, control does not return to the caller.

B6.3.2.4 Footprint

The PSCI_CPU_OFF command does not have any footprint.

B6.3.3 PSCI_CPU_ON command

Power up a core.

This command causes a REC exit due to PSCI. In response, the Host should provide the target REC (identified by `target_cpu`) by calling `RMI_PSCI_COMPLETE`.

See also:

- [A2.3.2 REC attributes](#)
- [A4.3.7 REC exit due to PSCI](#)
- [B4.3.23 RMI_PSCI_COMPLETE command](#)
- [B6.3.2 PSCI_CPU_OFF command](#)
- [B6.3.4 PSCI_CPU_SUSPEND command](#)
- [D1.4.1 PSCI_CPU_ON flow](#)

B6.3.3.1 Interface

B6.3.3.1.1 Input values

Name	Register	Bits	Type	Description
<code>fid</code>	X0	63:0	UInt64	FID, value <code>0xC4000003</code>
<code>target_cpu</code>	X1	63:0	Bits64	This parameter contains a copy of the affinity fields of the MPIDR register
<code>entry_point_address</code>	X2	63:0	Address	Address at which the core must resume execution
<code>context_id</code>	X3	31:0	UInt32	This parameter is only meaningful to the caller (must be present in X0 of the target PE upon first entry to Non-Secure exception level)

The following unused bits of PSCI_CPU_ON input values SBZ: X3[63:32].

B6.3.3.1.2 Context

The PSCI_CPU_ON command operates on the following context.

Name	Type	Value	Before	Description
<code>realm</code>	RmmRealm	<code>CurrentRealm()</code>	false	Current Realm
<code>target_rec</code>	RmmRec	<code>RecFromMpidr(target_cpu)</code>	false	Target REC

B6.3.3.1.3 Output values

Name	Register	Bits	Type	Description
<code>result</code>	X0	63:0	PsciReturnCode	Command return code

B6.3.3.2 Failure conditions

ID	Condition
entry	pre: !AddrIsProtected(entry_point_address, realm) post: result == PSCI_INVALID_ADDRESS
mpidr	pre: !MpidrIsUsed(target_cpu) post: result == PSCI_INVALID_PARAMETERS
runnable	pre: target_rec.flags.runnable == RUNNABLE post: result == PSCI_ALREADY_ON

B6.3.3.2.1 Failure condition ordering

The PSCI_CPU_ON command does not have any failure condition orderings.

B6.3.3.3 Success conditions

ID	Condition
entry	target_rec.pc == ToBits64(UInt(entry_point_address))
runnable	target_rec.flags.runnable == RUNNABLE

B6.3.3.4 Footprint

ID	Value
runnable	target_rec.flags.runnable

B6.3.4 PSCI_CPU_SUSPEND command

Suspend execution on the calling VPE.

This command causes a REC exit due to PSCI.

See also:

- [A4.3.7 REC exit due to PSCI](#)
- [B6.3.2 PSCI_CPU_OFF command](#)
- [B6.3.3 PSCI_CPU_ON command](#)

B6.3.4.1 Interface

B6.3.4.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0xC4000001
power_state	X1	31:0	UInt32	Identifier for a specific local state
entry_point_address	X2	63:0	Address	Address at which the core must resume execution
context_id	X3	63:0	UInt64	This parameter is only meaningful to the caller (must be present in X0 upon first entry to Non-Secure exception level)

The following unused bits of PSCI_CPU_SUSPEND input values SBZ: X1[63:32].

The RMM treats all target power states as suspend requests, and therefore the `entry_point_address` and `context_id` arguments are ignored.

B6.3.4.1.2 Output values

The PSCI_CPU_SUSPEND command does not have any output values.

Following execution of PSCI_CPU_SUSPEND, control does not return to the caller.

B6.3.4.2 Failure conditions

The PSCI_CPU_SUSPEND command does not have any failure conditions.

B6.3.4.3 Success conditions

The PSCI_CPU_SUSPEND command does not have any success conditions.

Following execution of PSCI_CPU_SUSPEND, control does not return to the caller.

B6.3.4.4 Footprint

The PSCI_CPU_SUSPEND command does not have any footprint.

B6.3.5 PSCI_FEATURES command

Query whether a specific PSCI feature is implemented.

See also:

- [B6.3.1 PSCI_AFFINITY_INFO command](#)
- [B6.3.2 PSCI_CPU_OFF command](#)
- [B6.3.3 PSCI_CPU_ON command](#)
- [B6.3.4 PSCI_CPU_SUSPEND command](#)
- [B6.3.6 PSCI_SYSTEM_OFF command](#)
- [B6.3.7 PSCI_SYSTEM_RESET command](#)

B6.3.5.1 Interface

B6.3.5.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0x8400000A
psci_func_id	X1	31:0	UInt32	Function ID for a PSCI Function

The following unused bits of PSCI_FEATURES input values SBZ: X1[63:32].

B6.3.5.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	PsciReturnCode	Command return code

B6.3.5.2 Failure conditions

The PSCI_FEATURES command does not have any failure conditions.

B6.3.5.3 Success conditions

ID	Condition
func_ok	pre: psci_func_id is a supported PSCI function. post: result == PSCI_SUCCESS
func_not_ok	pre: psci_func_id is not a supported PSCI function. post: result == PSCI_NOT_SUPPORTED

B6.3.5.4 Footprint

The PSCI_FEATURES command does not have any footprint.

B6.3.6 PSCI_SYSTEM_OFF command

Shut down the system.

This command causes a REC exit due to PSCI.

See also:

- [A2.3.2 REC attributes](#)
- [A4.3.7 REC exit due to PSCI](#)
- [B6.3.7 PSCI_SYSTEM_RESET command](#)

B6.3.6.1 Interface

B6.3.6.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0x84000008

B6.3.6.1.2 Context

The PSCI_SYSTEM_OFF command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm

B6.3.6.1.3 Output values

The PSCI_SYSTEM_OFF command does not have any output values.

Following execution of PSCI_SYSTEM_OFF, control does not return to the caller.

B6.3.6.2 Failure conditions

The PSCI_SYSTEM_OFF command does not have any failure conditions.

B6.3.6.3 Success conditions

ID	Condition
state	realm.state == REALM_SYSTEM_OFF

Following execution of PSCI_SYSTEM_OFF, control does not return to the caller.

B6.3.6.4 Footprint

The PSCI_SYSTEM_OFF command does not have any footprint.

B6.3.7 PSCI_SYSTEM_RESET command

Shut down the system.

This command causes a REC exit due to PSCI.

See also:

- [A2.3.2 REC attributes](#)
- [A4.3.7 REC exit due to PSCI](#)
- [B6.3.6 PSCI_SYSTEM_OFF command](#)

B6.3.7.1 Interface

B6.3.7.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0x84000009

B6.3.7.1.2 Context

The PSCI_SYSTEM_RESET command operates on the following context.

Name	Type	Value	Before	Description
realm	RmmRealm	CurrentRealm()	false	Current Realm

B6.3.7.1.3 Output values

The PSCI_SYSTEM_RESET command does not have any output values.

Following execution of PSCI_SYSTEM_RESET, control does not return to the caller.

B6.3.7.2 Failure conditions

The PSCI_SYSTEM_RESET command does not have any failure conditions.

B6.3.7.3 Success conditions

ID	Condition
state	realm.state == REALM_SYSTEM_OFF

Following execution of PSCI_SYSTEM_RESET, control does not return to the caller.

B6.3.7.4 Footprint

The PSCI_SYSTEM_RESET command does not have any footprint.

B6.3.8 PSCI_VERSION command

Query the version of PSCI implemented.

B6.3.8.1 Interface

B6.3.8.1.1 Input values

Name	Register	Bits	Type	Description
fid	X0	63:0	UInt64	FID, value 0x84000000

B6.3.8.1.2 Output values

Name	Register	Bits	Type	Description
result	X0	63:0	PsciInterfaceVersion	Interface version

See also:

- [B6.2 PSCI version](#)

B6.3.8.2 Failure conditions

The PSCI_VERSION command does not have any failure conditions.

B6.3.8.3 Success conditions

The PSCI_VERSION command does not have any success conditions.

B6.3.8.4 Footprint

The PSCI_VERSION command does not have any footprint.

B6.4 PSCI types

This section defines types which are used in the PSCI interface.


B6.4.1 PscilInterfaceVersion type

The PsciInterfaceVersion fieldset contains an PSCI interface version.

The PsciInterfaceVersion fieldset is a **concrete type**.

The width of the PsciInterfaceVersion fieldset is 64 bits.

The fields of the PsciInterfaceVersion fieldset are shown in the following diagram.



The fields of the PsciInterfaceVersion fieldset are shown in the following table.

Name	Bits	Description	Value
minor	15:0	Interface minor version number (the value y in interface version x.y)	UInt16
major	30:16	Interface major version number (the value x in interface version x.y)	UInt15
	63:31	Reserved	MBZ

B6.4.2 PsciReturnCode type

The `PsciReturnCode` enumeration represents the return code of a PSCI command.

The PsciReturnCode enumeration is a **concrete type**.

The width of the PsciReturnCode enumeration is 64 bits.

The values of the PsciReturnCode enumeration are shown in the following table.

Encoding	Name	Description
-9	PSCI_INVALID_ADDRESS	Refer to PSCI specification
-8	PSCI_DISABLED	Refer to PSCI specification
-7	PSCI_NOT_PRESENT	Refer to PSCI specification
-6	PSCI_INTERNAL_FAILURE	Refer to PSCI specification
-5	PSCI_ON_PENDING	Refer to PSCI specification
-4	PSCI_ALREADY_ON	Refer to PSCI specification
-3	PSCI_DENIED	Refer to PSCI specification
-2	PSCI_INVALID_PARAMETERS	Refer to PSCI specification
-1	PSCI_NOT_SUPPORTED	Refer to PSCI specification

Encoding	Name	Description
0	PSCI_SUCCESS	Refer to PSCI specification
1	PSCI_OFF	Refer to PSCI specification

Unused encodings for the PsciReturnCode enumeration are reserved for use by future versions of this specification.

DRAFT

DRAFT

Part C

Constants and types

Chapter C1

RMM constants

This section describes constants which are used in the definition of RMM commands or RMM abstract state.

C1.1 RMM_GRANULE_SIZE

Size of a Granule in bytes.

The value of RMM_GRANULE_SIZE is 0x1000.

C1.2 RMM_NUM_PERM_OVERLAY_INDICES

Number of permission overlay indices.

The value of RMM_NUM_PERM_OVERLAY_INDICES is 15.

C1.3 RMM_RTT_BLOCK_LEVEL

RTT level of a block entry.

The value of RMM_RTT_BLOCK_LEVEL is 2.

C1.4 RMM_RTT_PAGE_LEVEL

RTT level of a page entry.

The value of RMM_RTT_PAGE_LEVEL is 3.

C1.5 RMM_RTT_TREE_PRIMARY

Index of primary RTT tree.

The value of RMM_RTT_TREE_PRIMARY is 0.

DRAFT

Chapter C2

RMM types

This section describes types which are used to model the abstract state of the RMM.

C2.1 RmmAddressRange type

The RmmAddressRange structure contains address range.

The RmmAddressRange structure is an [abstract type](#).

The members of the RmmAddressRange structure are shown in the following table.

Name	Type	Description
base	Address	Base of address range (inclusive)
top	Address	Top of address range (exclusive)

The RmmAddressRange structure is used in the following types:

- [RmmPdev](#)

C2.2 RmmBoolean type

The RmmBoolean enumeration represents whether a feature is enabled.

The RmmBoolean enumeration is an [abstract type](#).

The values of the RmmBoolean enumeration are shown in the following table.

Name	Description
RMM_FALSE	False
RMM_TRUE	True

The RmmBoolean enumeration is used in the following types:

- [RmmRec](#)
- [RmmRttWalkNotAligned](#)


C2.3 RmmDataFlags type

The RmmDataFlags fieldset contains flags provided by the Host during DATA Granule creation.

The RmmDataFlags fieldset is a [concrete type](#).

The width of the RmmDataFlags fieldset is 64 bits.

The fields of the RmmDataFlags fieldset are shown in the following diagram.



The fields of the RmmDataFlags fieldset are shown in the following table.

Name	Bits	Description	Value
measure	0	Whether to measure DATA Granule contents	RmmDataMeasureContent
	63:1	Reserved	SBZ

The RmmDataFlags fieldset is used in the following types:

- [RmmMeasurementDescriptorData](#)

C2.4 RmmDataMeasureContent type

The RmmDataMeasureContent enumeration represents whether to measure DATA Granule contents.

The RmmDataMeasureContent enumeration is a [concrete type](#).

The width of the RmmDataMeasureContent enumeration is 1 bits.

The values of the RmmDataMeasureContent enumeration are shown in the following table.

Encoding	Name	Description
0	NO_MEASURE_CONTENT	Do not measure DATA Granule contents.
1	MEASURE_CONTENT	Measure DATA Granule contents.

The RmmDataMeasureContent enumeration is used in the following types:

- [RmmDataFlags](#)

C2.5 RmmDevCommState type

The RmmDevCommState enumeration represents the state of communication between an RMM device object and a device.

The RmmDevCommState enumeration is an [abstract type](#).

The values of the RmmDevCommState enumeration are shown in the following table.

Name	Description
DEV_COMM_ACTIVE	The RMM has initiated a device transaction. One or more device requests associated with this device transaction have been sent from the RMM to the device. The RMM has not received all the expected device responses associated with this device transaction.
DEV_COMM_ERROR	The RMM encountered an error during communication with the device.
DEV_COMM_IDLE	The RMM is not communicating with the device.
DEV_COMM_PENDING	The RMM has a device request which is ready to be sent to the device.

The RmmDevCommState enumeration is used in the following types:

- [RmmVdev](#)
- [RmmPdev](#)

C2.6 RmmDevMemShared type

The RmmDevMemShared enumeration represents whether device memory is shared.

The RmmDevMemShared enumeration is an [abstract type](#).

The values of the RmmDevMemShared enumeration are shown in the following table.

Name	Description
DEV_MEM_PRIVATE	Device memory is private
DEV_MEM_SHARED	Device memory is shared

The RmmDevMemShared enumeration is used in the following types:

- [RmmRec](#)

C2.7 RmmFeature type

The RmmFeature enumeration represents whether a feature is enabled.

The RmmFeature enumeration is an [abstract type](#).

See also:

- [Chapter A3 Feature discovery and configuration](#)

The values of the RmmFeature enumeration are shown in the following table.

Name	Description
FEATURE_FALSE	<ul style="list-style-type: none">• During discovery: Feature is not supported.• During selection: Feature is not enabled.
FEATURE_TRUE	<ul style="list-style-type: none">• During discovery: Feature is supported.• During selection: Feature is enabled.

The RmmFeature enumeration is used in the following types:

- [RmmFeatures](#)
- [RmmRealm](#)

C2.8 RmmFeatures type

The RmmFeatures structure contains features supported by RMM implementation.

The RmmFeatures structure is an [abstract type](#).

See also:

- [Chapter A3 Feature discovery and configuration](#)

The members of the RmmFeatures structure are shown in the following table.

Name	Type	Description
max_ipa_width	UInt64	Maximum IPA width
feat_lpa2	RmmFeature	Whether LPA2 is supported
feat_sve	RmmFeature	Whether SVE is supported
max_sve_vl	UInt64	Maximum SVE vector length
num_bps	UInt64	Number of breakpoints available
num_wps	UInt64	Number of watchpoints available
feat_pmu	RmmFeature	Number of watchpoints available
pmu_num_ctrs	UInt64	Number of PMU counters available
feat_sha_256	RmmFeature	Whether SHA-256 is supported
feat_sha_512	RmmFeature	Whether SHA-512 is supported
feat_da	RmmFeature	Whether Realm device assignment is supported

Name	Type	Description
max_num_aux_planes	UInt64	Maximum number of auxiliary Planes
plane_rtt	RmmPlaneRttFeature	RTT usage models supported for multi-Plane Realms
max_mecid	Bits64	Maximum supported MECID
max_recs_order	UInt64	Order of the maximum number of RECs which can be created per Realm
gicv3_num_lrs	UInt64	Number of GICv3 List Registers which are available.

C2.9 RmmGptEntry type

The RmmGptEntry enumeration represents granule Protection Table entry.

The RmmGptEntry enumeration is an [abstract type](#).

See also:

- [B3.30 GranuleAccessPermitted function](#)

The values of the RmmGptEntry enumeration are shown in the following table.

Name	Description
GPT_AAP	Access permitted via any PAS.
GPT_NS	Access permitted via Non-secure PAS only.
GPT_REALM	Access permitted via Realm PAS only.
GPT_ROOT	Access permitted via Root PAS only.
GPT_SECURE	Access permitted via Secure PAS only.

The RmmGptEntry enumeration is used in the following types:

- [RmmGranule](#)

C2.10 RmmGranule type

The RmmGranule structure contains attributes of a Granule.

The RmmGranule structure is an [abstract type](#).

The members of the RmmGranule structure are shown in the following table.

Name	Type	Description
gpt	RmmGptEntry	GPT entry
state	RmmGranuleState	Lifecycle state

C2.11 RmmGranuleState type

The RmmGranuleState enumeration represents the state of a granule.

The RmmGranuleState enumeration is an [abstract type](#).

The values of the RmmGranuleState enumeration are shown in the following table.

Name	Description
DATA	Realm code or data.
DELEGATED	Delegated for use by the RMM.
DEV_DELEGATED_PRIVATE	Device memory, delegated to the RMM and accessible via Realm PAS only.
DEV_DELEGATED_SHARED	Device memory, delegated to the RMM and accessible via any PAS.
DEV_PRIVATE	Device memory, mapped into a Realm and inaccessible by other requestors.
DEV_SHARED	Device memory, mapped into a Realm and also accessible by other requestors.
DEV_UNDELEGATED	Device memory, not delegated for use by the RMM.
PDEV	Physical device.
PDEV_AUX	Physical device auxiliary Granule.
RD	Realm Descriptor.
REC	Realm Execution Context.
REC_AUX	Realm Execution Context auxiliary Granule.
RTT	Realm Translation Table.
UNDELEGATED	Not delegated for use by the RMM.
VDEV	Virtual device.
VDEV_AUX	Virtual device auxiliary Granule.

The RmmGranuleState enumeration is used in the following types:

- [RmmGranule](#)

C2.12 RmmHashAlgorithm type

The RmmHashAlgorithm enumeration represents hash algorithm.

The RmmHashAlgorithm enumeration is an [abstract type](#).

The values of the RmmHashAlgorithm enumeration are shown in the following table.

Name	Description
HASH_SHA_256	SHA-256 (Secure Hash Standard (SHS) [19])
HASH_SHA_512	SHA-512 (Secure Hash Standard (SHS) [19])

The RmmHashAlgorithm enumeration is used in the following types:

- [RmmRealm](#)

- [RmmPdev](#)

C2.13 RmmHipas type

The RmmHipas enumeration represents host IPA state.

The RmmHipas enumeration is an [abstract type](#).

The values of the RmmHipas enumeration are shown in the following table.

Name	Description
HIPAS_ASSIGNED	Protected IPA which is associated with a DATA Granule.
HIPAS_ASSIGNED_DEV_PRIVATE	Protected IPA which is associated with a DEV_PRIVATE Granule.
HIPAS_ASSIGNED_DEV_SHARED	Protected IPA which is associated with a DEV_SHARED Granule.
HIPAS_ASSIGNED_NS	Unprotected IPA which is associated with an NS Granule.
HIPAS_UNASSIGNED	Protected IPA which is not associated with any Granule.
HIPAS_UNASSIGNED_NS	Unprotected IPA which is not associated with any Granule.

C2.14 RmmLfaPolicy type

The RmmLfaPolicy enumeration represents a Live Firmware Activation policy.

The RmmLfaPolicy enumeration is an [abstract type](#).

The values of the RmmLfaPolicy enumeration are shown in the following table.

Name	Description
LFA_ALLOW	LFA is permitted.
LFA_DISALLOW	LFA is not permitted.

The RmmLfaPolicy enumeration is used in the following types:

- [RmmRealm](#)

C2.15 RmmMeasurementDescriptorData type

The RmmMeasurementDescriptorData structure contains data structure used to calculate the contribution to the RIM of a DATA Granule.

The RmmMeasurementDescriptorData structure is a [concrete type](#).

The width of the RmmMeasurementDescriptorData structure is 256 (0x100) bytes.

See also:

- [B4.3.1.4 RMI_DATA_CREATE extension of RIM](#)

The members of the RmmMeasurementDescriptorData structure are shown in the following table.

Name	Byte offset	Type	Description
desc_type	0x0	Bits8	Measurement descriptor type, value 0x0
len	0x8	UInt64	Length of this data structure in bytes
rim	0x10	RmmRealmMeasurement	Current RIM value
ipa	0x50	Address	IPA at which the DATA Granule is mapped in the Realm
flags	0x58	RmmDataFlags	Flags provided by Host
content	0x60	RmmRealmMeasurement	Hash of contents of DATA Granule, or zero if flags indicate DATA Granule contents are unmeasured

Unused bits of the RmmMeasurementDescriptorData structure MBZ.

C2.16 RmmMeasurementDescriptorRec type

The RmmMeasurementDescriptorRec structure contains data structure used to calculate the contribution to the RIM of a REC.

The RmmMeasurementDescriptorRec structure is a [concrete type](#).

The width of the RmmMeasurementDescriptorRec structure is 256 (0x100) bytes.

See also:

- [B4.3.28.4 RMI_REC_CREATE extension of RIM](#)

The members of the RmmMeasurementDescriptorRec structure are shown in the following table.

Name	Byte offset	Type	Description
desc_type	0x0	Bits8	Measurement descriptor type, value 0x1
len	0x8	UInt64	Length of this data structure in bytes
rim	0x10	RmmRealmMeasurement	Current RIM value
content	0x50	RmmRealmMeasurement	Hash of 4KB page which contains REC parameters data structure

Unused bits of the RmmMeasurementDescriptorRec structure MBZ.

C2.17 RmmMeasurementDescriptorRipas type

The RmmMeasurementDescriptorRipas structure contains data structure used to calculate the contribution to the RIM of a RIPAS change.

The RmmMeasurementDescriptorRipas structure is a [concrete type](#).

The width of the RmmMeasurementDescriptorRipas structure is 256 (0x100) bytes.

See also:

- [B4.3.41.4 RMI_RTT_INIT_RIPAS extension of RIM](#)

The members of the RmmMeasurementDescriptorRipas structure are shown in the following table.

Name	Byte offset	Type	Description
desc_type	0x0	Bits8	Measurement descriptor type, value 0x2
len	0x8	UInt64	Length of this data structure in bytes
rim	0x10	RmmRealmMeasurement	Current RIM value
base	0x50	Address	Base IPA of the RIPAS change
top	0x58	Address	Top IPA of the RIPAS change

Unused bits of the RmmMeasurementDescriptorRipas structure MBZ.

C2.18 RmmMecPolicy type

The RmmMecPolicy enumeration represents a MEC policy.

The RmmMecPolicy enumeration is an [abstract type](#).

The values of the RmmMecPolicy enumeration are shown in the following table.

Name	Description
MEC_POLICY_PRIVATE	The MEC protects memory owned by a single Realm. A MEC with this policy may be referred to as a <i>Private MEC</i> .
MEC_POLICY_SHARED	The MEC protects memory owned by multiple Realms. A MEC with this policy may be referred to as a <i>Shared MEC</i> .

The RmmMecPolicy enumeration is used in the following types:

- [RmmRealm](#)

C2.19 RmmMecState type

The RmmMecState enumeration represents state of a MEC.

The RmmMecState enumeration is an [abstract type](#).

The values of the RmmMecState enumeration are shown in the following table.

Name	Description
MEC_STATE_PRIVATE_ASSIGNED	A Private MEC which is assigned to a Realm.
MEC_STATE_PRIVATE_UNASSIGNED	A Private MEC which is not assigned to a Realm.
MEC_STATE_SHARED	A Shared MEC.

C2.20 RmmMemPermLocked type

The RmmMemPermLocked enumeration represents whether a memory permission value is locked.

The RmmMemPermLocked enumeration is an [abstract type](#).

The values of the RmmMemPermLocked enumeration are shown in the following table.

Name	Description
MEM_PERM_LOCKED	Memory permission value is locked
MEM_PERM_UNLOCKED	Memory permission value is unlocked

The RmmMemPermLocked enumeration is used in the following types:

- [RmmRealm](#)

C2.21 RmmMemPerms type

The RmmMemPerms structure contains memory permissions.

The RmmMemPerms structure is an [abstract type](#).

The members of the RmmMemPerms structure are shown in the following table.

Name	Type	Description
values	Bits64 [16]	Mapping from memory permission index to memory permission label Values use architectural encodings.

The RmmMemPerms structure is used in the following types:

- [RmmRealm](#)

C2.22 RmmPdev type

The RmmPdev structure contains attributes of a PDEV.

The RmmPdev structure is an [abstract type](#).

The members of the RmmPdev structure are shown in the following table.

Name	Type	Description
pdev_id	Bits64	Device identifier
prot_config	RmmPdevProtConfig	Configuration of protection between system and device
segment_id	Bits16	Segment identifier PCIe Segment identifier of the Root Port and endpoint.

Name	Type	Description
root_id	Bits16	Root Port identifier Physical PCIe routing identifier of the Root Port to which the endpoint is connected.
cert_id	UInt64	Certificate identifier
rid_base	UInt64	Base of requester ID range (inclusive)
rid_top	UInt64	Top of requester ID range (exclusive)
hash_algo	RmmHashAlgorithm	Algorithm used to generate device digests
ide_sid	UInt64	IDE stream ID
iocoh_num_addr_range	UInt64	Number of IO-coherent address ranges
iocoh_addr_range	RmmAddressRange[16]	IO-coherent address range
fcoh_num_addr_range	UInt64	Number of fully-coherent address ranges
fcoh_addr_range	RmmAddressRange[4]	Fully-coherent address range
aux	Address[32]	Addresses of auxiliary Granules
num_aux	UInt64	Number of auxiliary Granules
state	RmmPdevState	Lifecycle state
comm_state	RmmDevCommState	Device communication state
num_vdevs	UInt64	Number of VDEVs associated with this PDEV whose state is not VDEV_STOPPED

C2.23 RmmPdevProtConfig type

The RmmPdevProtConfig enumeration represents configuration of protection between system and device.

The RmmPdevProtConfig enumeration is an [abstract type](#).

The values of the RmmPdevProtConfig enumeration are shown in the following table.

Name	Description
PDEV_FCOH_E2E_IDE	Fully-coherent device with end-to-end protection provided by IDE.
PDEV_FCOH_E2E_SYS	Fully-coherent device with end-to-end protection provided by system construction.
PDEV_IOCOH_E2E_IDE	IO-coherent device with end-to-end protection provided by IDE.
PDEV_IOCOH_E2E_SYS	IO-coherent device with end-to-end protection provided by system construction.

The RmmPdevProtConfig enumeration is used in the following types:

- [RmmPdev](#)

C2.24 RmmPdevState type

The RmmPdevState enumeration represents the state of a PDEV.

The RmmPdevState enumeration is an [abstract type](#).

The values of the RmmPdevState enumeration are shown in the following table.

Name	Description
PDEV_COMMUNICATING	The RMM is communicating with the device.
PDEV_ERROR	Device has reported a fatal error.
PDEV_HAS_KEY	RMM has device public key.
PDEV_IDE_RESETTING	The PDEV's IDE link is being reset.
PDEV_NEEDS_KEY	RMM needs device public key.
PDEV_NEW	Initial state of the device.
PDEV_READY	Secure connection between the RMM and the device has been established. Physical link between the device and memory is secured. Ready for creation of VDEV instances.
PDEV_STOPPED	Secure connection between the RMM and the device has been terminated.
PDEV_STOPPING	The RMM is communicating with the device to terminate the secure connection between the RMM and the device.

The RmmPdevState enumeration is used in the following types:

- [RmmPdev](#)

C2.25 RmmPhysicalAddressSpace type

The RmmPhysicalAddressSpace enumeration represents the PAS of a Granule.

The RmmPhysicalAddressSpace enumeration is an [abstract type](#).

See also:

- [B3.30 GranuleAccessPermitted function](#)

The values of the RmmPhysicalAddressSpace enumeration are shown in the following table.

Name	Description
PAS_NS	Non-secure PAS.
PAS_REALM	Realm PAS.
PAS_ROOT	Root PAS.
PAS_SECURE	Secure PAS.

C2.26 RmmPlaneRttFeature type

The RmmPlaneRttFeature enumeration represents RTT usage models supported for multi-Plane Realms.

The RmmPlaneRttFeature enumeration is an [abstract type](#).

See also:

- [A3.11 Support for auxiliary Planes](#)

The values of the RmmPlaneRttFeature enumeration are shown in the following table.

Name	Description
PLANE_RTT_AUX	A multi-Plane Realm uses auxiliary RTTs
PLANE_RTT_AUX_SINGLE	A multi-Plane Realm can be configured to either use auxiliary RTTs, or a single RTT
PLANE_RTT_SINGLE	A multi-Plane Realm uses a single RTT

The RmmPlaneRttFeature enumeration is used in the following types:

- [RmmFeatures](#)

C2.27 RmmRdev type

The RmmRdev structure contains attributes of an RDEV.

The RmmRdev structure is an [abstract type](#).

The members of the RmmRdev structure are shown in the following table.

Name	Type	Description
state	RmmRdevState	Lifecycle state
operation	RmmRdevOperation	Operation being performed by the RDEV
vdev_ptr	Address	PA of VDEV associated with this RDEV

C2.28 RmmRdevOperation type

The RmmRdevOperation enumeration represents operation being performed by an RDEV.

The RmmRdevOperation enumeration is an [abstract type](#).

The values of the RmmRdevOperation enumeration are shown in the following table.

Name	Description
RDEV_OP_GET_INTERFACE_REPORT	RDEV is handling an RSI_RDEV_GET_INTERFACE_REPORT request.
RDEV_OP_GET_MEASUREMENTS	RDEV is handling an RSI_RDEV_GET_MEASUREMENTS request.

Name	Description
RDEV_OP_LOCK	RDEV is handling an RSI_RDEV_LOCK request.
RDEV_OP_NONE	No operation is being performed.
RDEV_OP_START	RDEV is handling an RSI_RDEV_START request.

The RmmRdevOperation enumeration is used in the following types:

- [RmmRdev](#)

C2.29 RmmRdevState type

The RmmRdevState enumeration represents the state of an RDEV.

The RmmRdevState enumeration is an [abstract type](#).

The values of the RmmRdevState enumeration are shown in the following table.

Name	Description
RDEV_ERROR	Device interface has reported a fatal error.
RDEV_LOCKED	Device interface is locked.
RDEV_LOCKED_BUSY	Device interface is locked and is handling an interruptible Realm device operation.
RDEV_NEW	Device interface is unlocked.
RDEV_NEW_BUSY	Device interface is unlocked and is handling an interruptible Realm device operation.
RDEV_STARTED	Device interface is started.
RDEV_STARTED_BUSY	Device interface is started and is handling an interruptible Realm device operation.
RDEV_STOPPED	Device interface is stopped.
RDEV_STOPPING	Device interface is stopping.

The RmmRdevState enumeration is used in the following types:

- [RmmRdev](#)

C2.30 RmmRealm type

The RmmRealm structure contains attributes of a Realm.

The RmmRealm structure is an [abstract type](#).

See also:

- [A2.1 Realm](#)

The members of the RmmRealm structure are shown in the following table.

Name	Type	Description
feat_lpa2	RmmFeature	Whether LPA2 is enabled for this Realm
ipa_width	UInt8	IPA width in bits
measurements	RmmRealmMeasurement [5]	Realm measurements
hash_algo	RmmHashAlgorithm	Algorithm used to compute Realm measurements
rec_index	UInt64	Index of next REC to be created
rtt_base	Address [4]	Realm Translation Table base addresses If rtt_tree_pp is FEATURE_FALSE then only the first entry is valid. If rtt_tree_pp is FEATURE_TRUE then only the first (num_aux_planes + 1) entries are valid.
rtt_level_start	Int64	RTT starting level
rtt_num_start	UInt64	Number of physically contiguous starting level RTTs
state	RmmRealmState	Lifecycle state
vmid	Bits16 [4]	Virtual Machine Identifiers If rtt_tree_pp is FEATURE_FALSE then only the first entry is valid. If rtt_tree_pp is FEATURE_TRUE then only the first (num_aux_planes + 1) entries are valid.
rpv	Bits512	Realm Personalization Value
feat_da	RmmFeature	Whether Realm device assignment is enabled for this Realm
rtt_tree_pp	RmmFeature	Whether this Realm has an RTT per Plane
num_aux_planes	UInt64	Number of auxiliary Planes
overlay_perms	RmmMemPerms [4]	Memory overlay permissions
overlay_locked	RmmMemPermLocked [16]	Whether memory overlay value is locked
lfa_policy	RmmLfaPolicy	Live Firmware Activation policy for components within the Realm's TCB
mecid	Bits64	Memory Encryption Context Identifier
mec_policy	RmmMecPolicy	MEC policy
num_recs	UInt64	Number of RECs owned by this Realm
num_vdevs	UInt64	Number of VDEVs which have been assigned to this Realm

C2.31 RmmRealmMeasurement type

The RmmRealmMeasurement type is realm measurement.

The RmmRealmMeasurement type is a [concrete type](#).

The width of the RmmRealmMeasurement type is 512 bits.

C2.32 RmmRealmState type

The RmmRealmState enumeration represents the state of a Realm.

The RmmRealmState enumeration is an [abstract type](#).

The values of the RmmRealmState enumeration are shown in the following table.

Name	Description
REALM_ACTIVE	Eligible for execution.
REALM_NEW	Under construction. Not eligible for execution.
REALM_SYSTEM_OFF	System has been turned off. Not eligible for execution.

The RmmRealmState enumeration is used in the following types:

- [RmmRealm](#)

C2.33 RmmRec type

The RmmRec structure contains attributes of a REC.

The RmmRec structure is an [abstract type](#).

See also:

- [A2.3 Realm Execution Context](#)

The members of the RmmRec structure are shown in the following table.

Name	Type	Description
attest_state	RmmRecAttestState	Attestation token generation state
attest_challenge	Bits512	Challenge for under-construction attestation token
aux	Address [16]	Addresses of auxiliary Granules
emulatable_abort	RmmRecEmulatableAbort	Whether the most recent exit from this REC was due to an Emulatable Data Abort
flags	RmmRecFlags	Flags which control REC behavior
gprs	Bits64 [32]	General-purpose register values
mpidr	Bits64	MPIDR value
owner	Address	PA of RD of Realm which owns this REC
pc	Bits64	Program counter value
pending	RmmRecPending	Whether a REC operation is pending
vdev_id	Bits64	Virtual device ID
inst_id	UInt64	Device instance ID
inst_id_valid	RmmBoolean	Whether device instance ID is valid

Name	Type	Description
state	RmmRecState	Lifecycle state
sysregs	RmmSystemRegisters	EL1 and EL0 system register values
ripas_addr	Address	Next address to be processed in RIPAS change
ripas_top	Address	Top address of pending RIPAS change
ripas_value	RmmRipas	RIPAS value of pending RIPAS change
ripas_destroyed	RmmRipasChangeDestroyed	Whether a RIPAS change from DESTROYED should be permitted
ripas_response	RmmRecResponse	Host response to RIPAS change request
ripas_dev_pa	Address	Base PA of device memory region, if RIPAS change is pending due to execution of RSI_RDEV_VALIDATE_MAPPING
ripas_dev_shared	RmmDevMemShared	Value of shared bit, if RIPAS change is pending due to execution of RSI_RDEV_VALIDATE_MAPPING
s2ap_addr	Address	Next address to be processed in S2AP change
s2ap_top	Address	Top address of pending S2AP change
s2ap_overlay	UInt3	Overlay index of pending S2AP change
s2ap_response	RmmRecResponse	Host response to S2AP change request
gic_owner	UInt64	Index of Plane which is the GIC owner

C2.34 RmmRecAttestState type

The RmmRecAttestState enumeration represents whether an attestation token generation operation is ongoing on this REC.

The RmmRecAttestState enumeration is an [abstract type](#).

The values of the RmmRecAttestState enumeration are shown in the following table.

Name	Description
ATTEST_IN_PROGRESS	An attestation token generation operation is in progress.
NO_ATTEST_IN_PROGRESS	No attestation token generation operation is in progress.

The RmmRecAttestState enumeration is used in the following types:

- [RmmRec](#)

C2.35 RmmRecEmulatableAbort type

The RmmRecEmulatableAbort enumeration represents whether the most recent exit from a REC was due to an Emulatable Data Abort.

The RmmRecEmulatableAbort enumeration is an [abstract type](#).

The values of the RmmRecEmulatableAbort enumeration are shown in the following table.

Name	Description
EMULATABLE_ABORT	The most recent exit from a REC was due to an Emulatable Data Abort.
NOT_EMULATABLE_ABORT	The most recent exit from a REC was not due to an Emulatable Data Abort.

The RmmRecEmulatableAbort enumeration is used in the following types:

- [RmmRec](#)

C2.36 RmmRecFlags type

The RmmRecFlags structure contains REC flags.

The RmmRecFlags structure is an [abstract type](#).

The members of the RmmRecFlags structure are shown in the following table.

Name	Type	Description
runnable	RmmRecRunnable	Whether the REC is eligible to run

The RmmRecFlags structure is used in the following types:

- [RmmRec](#)

C2.37 RmmRecPending type

The RmmRecPending enumeration represents whether a REC operation is pending.

The RmmRecPending enumeration is an [abstract type](#).

The values of the RmmRecPending enumeration are shown in the following table.

Name	Description
REC_PENDING_HOST_CALL	A Host call is pending.
REC_PENDING_NONE	No operation is pending.
REC_PENDING_PSCI	A PSCI operation is pending.
REC_PENDING_VDEV_REQUEST	A VDEV request is pending.

The RmmRecPending enumeration is used in the following types:

- [RmmRec](#)

C2.38 RmmRecResponse type

The RmmRecResponse enumeration represents whether the Host accepted or rejected a Realm request.

The RmmRecResponse enumeration is an [abstract type](#).

The values of the RmmRecResponse enumeration are shown in the following table.

Name	Description
ACCEPT	Host accepted the Realm request.
REJECT	Host rejected the Realm request.

The RmmRecResponse enumeration is used in the following types:

- [RmmRec](#)

C2.39 RmmRecRunnable type

The RmmRecRunnable enumeration represents whether a REC is eligible for execution.

The RmmRecRunnable enumeration is an [abstract type](#).

The values of the RmmRecRunnable enumeration are shown in the following table.

Name	Description
NOT_RUNNABLE	Not eligible for execution.
RUNNABLE	Eligible for execution.

The RmmRecRunnable enumeration is used in the following types:

- [RmmRecFlags](#)

C2.40 RmmRecState type

The RmmRecState enumeration represents the state of a REC.

The RmmRecState enumeration is an [abstract type](#).

The values of the RmmRecState enumeration are shown in the following table.

Name	Description
REC_READY	REC is not currently running.
REC_RUNNING	REC is currently running.

The RmmRecState enumeration is used in the following types:

- [RmmRec](#)

C2.41 RmmRipas type

The RmmRipas enumeration represents realm IPA state.

The RmmRipas enumeration is an [abstract type](#).

The values of the RmmRipas enumeration are shown in the following table.

Name	Description
DESTROYED	Address which is inaccessible to the Realm due to an action taken by the Host.
DEV	Address where memory of an assigned Realm device is mapped.
EMPTY	Address where no Realm resources are mapped.
RAM	Address where private code or data owned by the Realm is mapped.

The RmmRipas enumeration is used in the following types:

- [RmmRec](#)
- [RmmRttEntry](#)

C2.42 RmmRipasChangeDestroyed type

The RmmRipasChangeDestroyed enumeration represents whether a RIPAS change from DESTROYED should be permitted.

The RmmRipasChangeDestroyed enumeration is an [abstract type](#).

The values of the RmmRipasChangeDestroyed enumeration are shown in the following table.

Name	Description
CHANGE_DESTROYED	A RIPAS change from DESTROYED should be permitted.
NO_CHANGE_DESTROYED	A RIPAS change from DESTROYED should not be permitted.

The RmmRipasChangeDestroyed enumeration is used in the following types:

- [RmmRec](#)

C2.43 RmmRtt type

The RmmRtt structure contains an RTT.

The RmmRtt structure is an [abstract type](#).

The members of the RmmRtt structure are shown in the following table.

Name	Type	Description
entries	RmmRttEntry [512]	Entries

C2.44 RmmRttEntry type

The RmmRttEntry structure contains attributes of an RTT Entry.

The RmmRttEntry structure is an [abstract type](#).

See also:

- [A5.5 Realm Translation Table](#)

The members of the RmmRttEntry structure are shown in the following table.

Name	Type	Description
addr	Address	Output address
ripas	RmmRipas	RIPAS
state	RmmRttEntryState	State
MemAttr	Bits3	MemAttr
s2ap_base	UInt3	S2AP base permission index
s2ap_overlay	UInt3	S2AP overlay permission index

The RmmRttEntry structure is used in the following types:

- [RmmRttWalkResult](#)
- [RmmRtt](#)

C2.45 RmmRttEntryState type

The RmmRttEntryState enumeration represents the state of an RTTE.

The RmmRttEntryState enumeration is an [abstract type](#).

The values of the RmmRttEntryState enumeration are shown in the following table.

Name	Description
ASSIGNED	This RTTE is identified by a Protected IPA. The output address of this RTTE points to a DATA Granule.
ASSIGNED_DEV_PRIVATE	This RTTE is identified by a Protected IPA. The output address of this RTTE points to a DEV_PRIVATE Granule.
ASSIGNED_DEV_SHARED	This RTTE is identified by a Protected IPA. The output address of this RTTE points to a DEV_SHARED Granule.
ASSIGNED_NS	This RTTE is identified by an Unprotected IPA. The output address of this RTTE points to an NS Granule.
AUX_DESTROYED	An auxiliary RTT was destroyed while a corresponding primary RTT entry was live.
TABLE	The output address of this RTTE points to the next-level RTT.
UNASSIGNED	This RTTE is identified by a Protected IPA. This RTTE is not associated with any Granule.
UNASSIGNED_NS	This RTTE is identified by an Unprotected IPA. This RTTE is not associated with any Granule.

The RmmRttEntryState enumeration is used in the following types:

- [RmmRttEntry](#)

C2.46 RmmRttWalkNotAligned type

The RmmRttWalkNotAligned structure contains result of an RTT walk which is not aligned to the requested level.

The RmmRttWalkNotAligned structure is an [abstract type](#).

The members of the RmmRttWalkNotAligned structure are shown in the following table.

Name	Type	Description
valid	RmmBoolean	TRUE if an RTT walk was performed whose result is not aligned to the requested level
index	UInt64	RTT index
addr	Address	Address
walk	RmmRttWalkResult	Walk result

C2.47 RmmRttWalkResult type

The RmmRttWalkResult structure contains result of an RTT walk.

The RmmRttWalkResult structure is an [abstract type](#).

See also:

- [A5.5.10 RTT walk](#)

The members of the RmmRttWalkResult structure are shown in the following table.

Name	Type	Description
level	Int8	RTT level reached by the walk
rtt_addr	Address	Address of RTT reached by the walk
rtte	RmmRttEntry	RTTE reached by the walk

The RmmRttWalkResult structure is used in the following types:

- [RmmRttWalkNotAligned](#)

C2.48 RmmSystemRegisters type

The RmmSystemRegisters structure contains EL0 and EL1 system registers.

The RmmSystemRegisters structure is an [abstract type](#).

The RmmSystemRegisters structure is used in the following types:

- [RmmRec](#)

C2.49 RmmVdev type

The RmmVdev structure contains attributes of a VDEV.

The RmmVdev structure is an [abstract type](#).

The members of the RmmVdev structure are shown in the following table.

Name	Type	Description
vdev_id	Bits64	Virtual device identifier
tdi_id	Bits64	TDI identifier
inst_id	UInt64	Instance identifier
pdev	Address	PA of parent PDEV
realm	Address	PA of RD of Realm which owns this REC
state	RmmVdevState	Lifecycle state
comm_state	RmmDevCommState	Device communication state
aux	Address [32]	Addresses of auxiliary Granules
num_aux	UInt64	Number of auxiliary Granules

C2.50 RmmVdevState type

The RmmVdevState enumeration represents the state of a VDEV.

The RmmVdevState enumeration is an [abstract type](#).

The values of the RmmVdevState enumeration are shown in the following table.

Name	Description
VDEV_COMMUNICATING	The RMM is communicating with the VDEV.
VDEV_ERROR	Device interface has reported a fatal error.
VDEV_READY	No device transaction is associated with the VDEV.
VDEV_STOPPED	Device interface is stopped.
VDEV_STOPPING	The RMM is communicating with the VDEV to stop the device interface.

The RmmVdevState enumeration is used in the following types:

- [RmmVdev](#)

Chapter C3

Generic types

This section defines types which are shared between RMM interfaces and descriptions of RMM abstract state.

See also:

- [B4.4 RMI types](#)
- [B5.4 RSI types](#)
- [B6.4 PSCI types](#)
- [Chapter C2 RMM types](#)

C3.1 Address type

The Address type is an address.

The Address type is a [concrete type](#).

The width of the Address type is 64 bits.

C3.2 BitsN type

The BitsN type is an N-bit field.

The BitsN type is a [concrete type](#).

The width of the BitsN type is N bits.

C3.3 IntN type

The IntN type is an signed N-bit integer.

The IntN type is a [concrete type](#).

The width of the IntN type is N bits.

C3.4 UIntN type

The UIntN type is an unsigned N-bit integer.

The UIntN type is a [concrete type](#).

The width of the UIntN type is N bits.

DRAFT

DRAFT

Part D
Usage

Chapter D1

Flows

This section presents flows which explain how the RMM architecture can be used by the Host, and by Realm software.

Note that parts of the sequences below are for illustration only. For example, in the Realm creation flows, the `RMI_GRANULE_DELEGATE` and `RMI_GRANULE_UNDELEGATE` commands are called immediately before or after the `RMI_X_CREATE` and `RMI_X_DESTROY` commands respectively. An alternative flow would be for the Host to maintain a pool of Granules in the `DELEGATED` state, from which RMM data structures and Realm data can be allocated on demand.


D1.1 Granule delegation flows

D1.1.1 Granule delegation flow

The following diagram shows how the GPT entry of a Granule is changed to GPT_REALM.

See [Arm Architecture Reference Manual Supplement, The Realm Management Extension \(RME\), for Armv9-A \[2\]](#) for example software flows for the operations performed by the Monitor in this flow.

It is anticipated that the Monitor software will be required to use synchronization mechanisms to serialize access to the GPT.



See also:


- [A2.2.1 Granule attributes](#)
- [B4.3.7 RMI_GRANULE_DELEGATE command](#)
- [D1.1.2 Granule undelegation flow](#)

D1.1.2 Granule undelegation flow

The following diagram shows how the GPT entry of a Granule is changed from GPT_REALM.

See [Arm Architecture Reference Manual Supplement, The Realm Management Extension \(RME\), for Armv9-A \[2\]](#) for example software flows for the operations performed by the Monitor in this flow.

It is anticipated that the Monitor software will be required to use synchronization mechanisms to serialize access to the GPT.



See also:

- [A2.2.1 Granule attributes](#)
- [B4.3.10 RMI_GRANULE_UNDELEGATE command](#)
- [D1.1.1 Granule delegation flow](#)

D1.2 Realm lifecycle flows

This section contains flows which relate to the Realm lifecycle.

See also:

- [A2.1.5 Realm lifecycle](#)

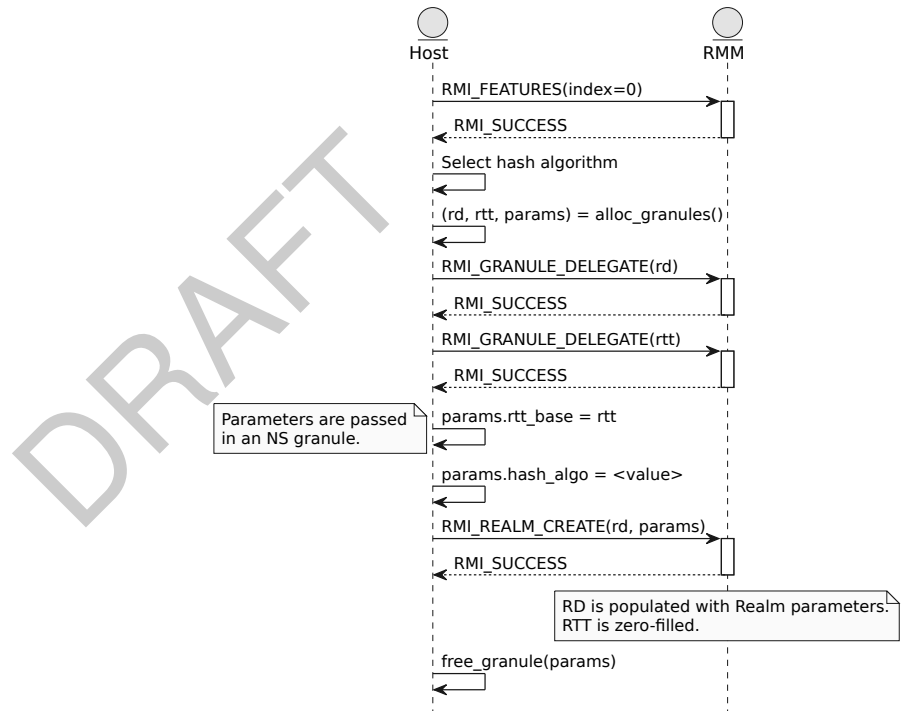
D1.2.1 Realm creation flow

The following diagram shows the flow for creating a Realm.

To create a Realm, the Host must allocate and delegate two Granules:

- `rd` to store the Realm Descriptor
- `rtt` which will be the starting level Realm Translation Table (RTT)

The Host also provides an NS Granule (`params`) containing Realm creation parameters.



See also:


- [B4.3.7 RMI_GRANULE_DELEGATE command](#)
- [B4.3.25 RMI_REALM_CREATE command](#)
- [D1.2.5 Realm destruction flow](#)

D1.2.2 Realm Translation Table creation flow

The following diagram shows the flow for populating the Realm Translation Tables (RTTs).

The starting level Realm Translation Tables (RTTs) are provided at Realm creation time.

Subsequent levels of RTT are added using the `RMI_RTT_CREATE` command. This can be performed when the state of the Realm is `REALM_NEW` or `REALM_ACTIVE`.



See also:

- [Chapter A5 Realm memory management](#)
- [B4.3.38 RMI_RTT_CREATE command](#)
- [D1.2.1 Realm creation flow](#)
- [D1.2.3 Initialize memory of New Realm flow](#)

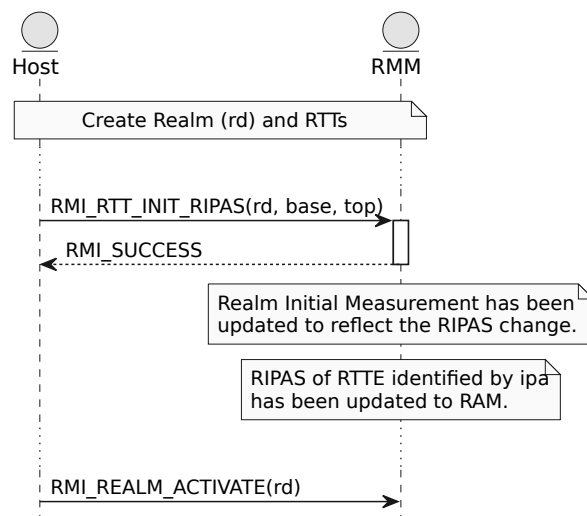
D1.2.3 Initialize memory of New Realm flow

Immediately following Realm creation, every page in the Protected IPA space has its RIPAS set to EMPTY. There are two ways in which the Host can set the RIPAS of a given page of Protected IPA space to RAM:

1. Change the RIPAS by executing RMI_RTT_INIT_RIPAS, but do not populate the contents of the page. The RIM is extended to reflect the RIPAS change.
2. Both change the RIPAS and populate the page with contents provided by the Host, by executing RMI_DATA_CREATE. The RIM is extended to reflect the contents added by the Host.

Once the Host has performed either of these actions for a given page of Protected IPA space, that page cannot be further modified prior to Realm activation.

The following diagram shows the flow for initializing the RIPAS without providing contents.

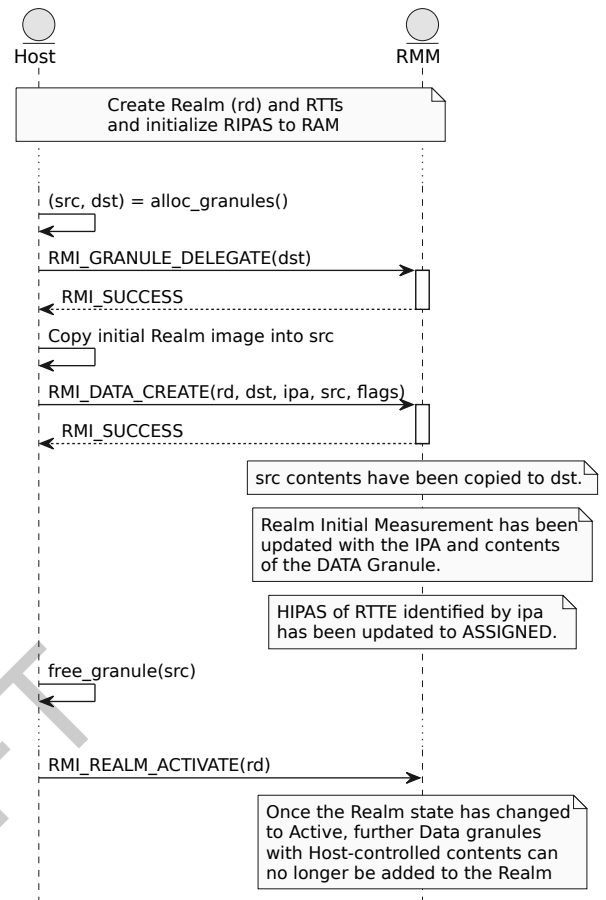


The following diagram shows the flow for populating the page with contents provided by the Host.

To do this, the Host must:

- Delegate a destination Granule (*dst*).
- Provide an NS Granule (*src*), whose contents will be copied into the destination Granule.
- Specify the Protected IPA *ipa* at which the *dst* Granule should be mapped in the Realm's IPA space.
- Ensure that the level 3 RTT which contains the RTTE identified by the Protected IPA has been created.

Once the Data Granule has been created, the *src* Granule can be reallocated by the Host.



See also:

- [A2.2.1 Granule attributes](#)
- [A5.2.2 Realm IPA state](#)
- [A7.1.1 Realm Initial Measurement](#)
- [B4.3.1 RMI_DATA_CREATE command](#)
- [B4.3.7 RMI_GRANULE_DELEGATE command](#)
- [B4.3.41 RMI_RTT_INIT_RIPAS command](#)
- [D1.2.1 Realm creation flow](#)
- [D1.2.2 Realm Translation Table creation flow](#)
- [D1.2.5 Realm destruction flow](#)


D1.2.4 REC creation flow

The following diagram shows the flow for creating a REC during Realm creation.

To create a REC, the Host must:

- Delegate a destination Granule (*rec*).
- Query the number of auxiliary Granules required, by calling `RMI_REC_AUX_COUNT`
- Delegate the required number of auxiliary Granules (*aux*)
- Provide auxiliary Granule addresses, register values and REC activation status in an NS Granule (*params*).

Once the REC has been created, the *params* Granule can be reallocated by the Host.



See also:

- [B4.3.7 RMI_GRANULE_DELEGATE command](#)
- [B4.3.27 RMI_REC_AUX_COUNT command](#)
- [B4.3.28 RMI_REC_CREATE command](#)
- [D1.2.1 Realm creation flow](#)
- [D1.2.5 Realm destruction flow](#)

D1.2.5 Realm destruction flow

The following diagram shows the flow for destroying a Realm.

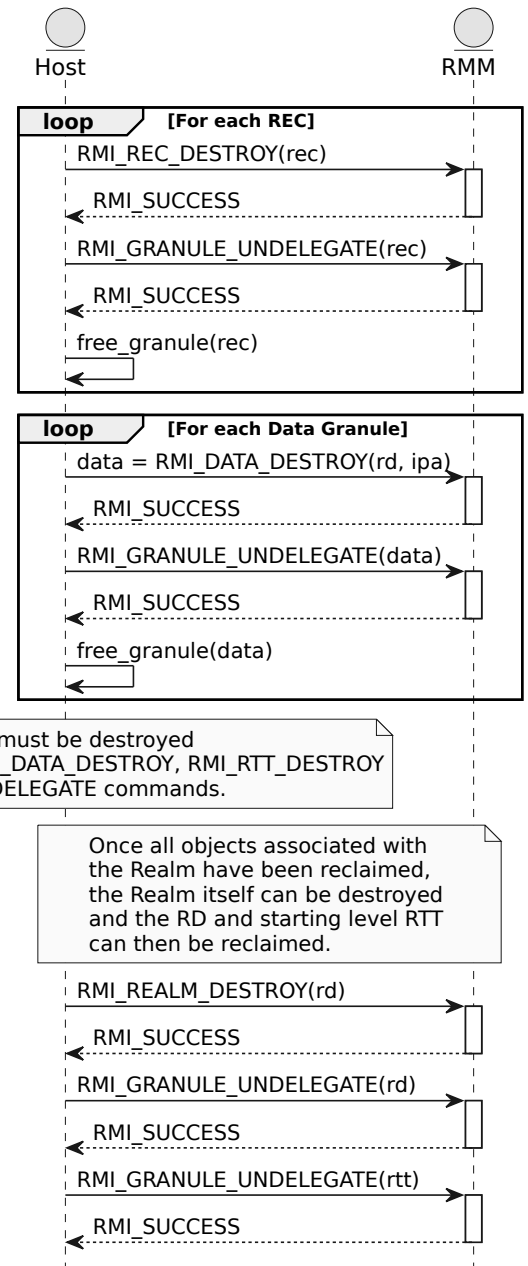
To destroy a Realm, the Host must first make the Realm non-live. This is done by destroying (in any order) the objects which are associated with the Realm:

- Data Granules
- RECs
- RTTs

Finally, the Realm itself can be destroyed.

Once each of these objects has been destroyed, the corresponding Granules can be undelegated and reallocated by the Host.

DRAFT



See also:

- [A2.1.4 Realm liveness](#)
- [B4.3.3 RMI_DATA_DESTROY command](#)
- [B4.3.10 RMI_GRANULE_UNDELEGATE command](#)
- [B4.3.26 RMI_REALM_DESTROY command](#)
- [B4.3.29 RMI_REC_DESTROY command](#)
- [D1.2.1 Realm creation flow](#)

D1.3 Realm exception model flows

This section contains flows which relate to the Realm exception model.

See also:

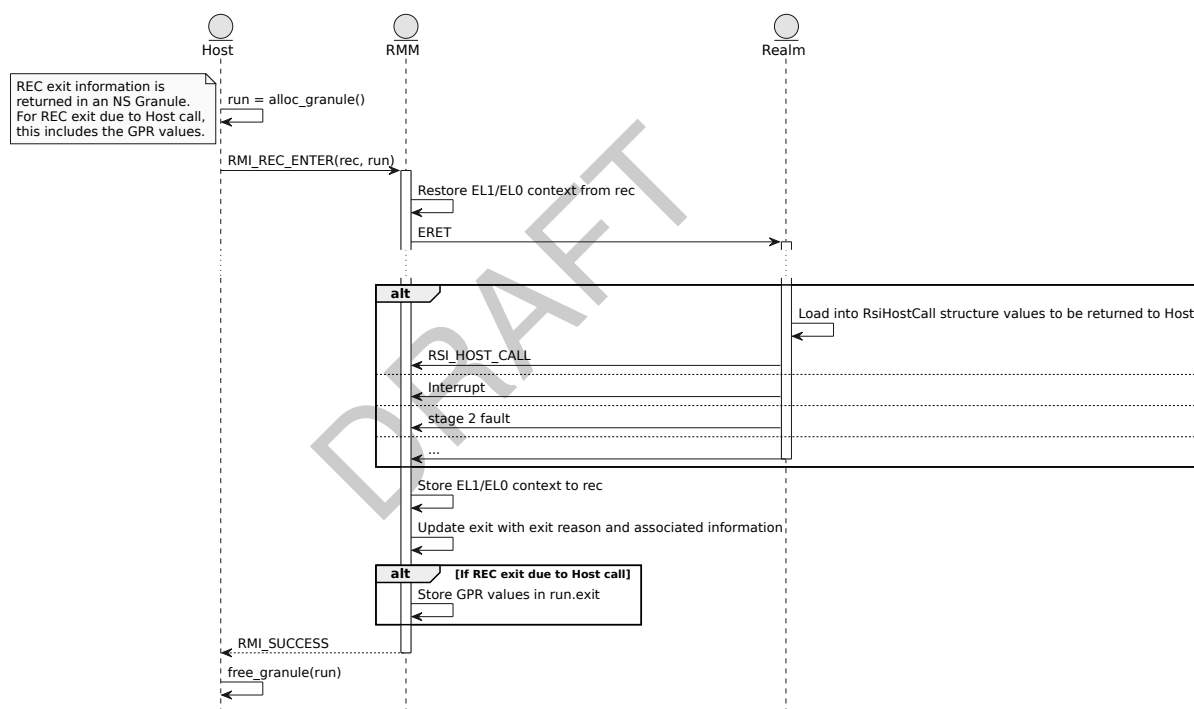
- [Chapter A4 Realm exception model](#)

D1.3.1 Realm entry and exit flow

The following diagram shows how a Realm is executed, and illustrates the different reasons for exiting the Realm and returning control to the Host.

A REC is entered using the RMI_REC_ENTER command. The parameters to this command include:

- an *RmiRecEnter* object, which is a data structure used to pass values from the Host to the RMM on REC entry
- an *RmiRecExit* object, which is a data structure used to pass values from the RMM to the Host on REC exit



See also:

- [Chapter A4 Realm exception model](#)
- [D1.3.2 Host call flow](#)
- [D1.3.3 REC exit due to Data Abort fault flow](#)
- [D1.3.4 MMIO emulation flow](#)


D1.3.2 Host call flow

The following diagram shows how software executing inside the Realm can voluntarily yield control back to the Host by making a Host call.

A REC is entered using the RMI_REC_ENTER command. The parameters to this command include:

- an *RmiRecEnter* object, which is a data structure used to pass values from the Host to the RMM on REC entry
- an *RmiRecExit* object, which is a data structure used to pass values from the RMM to the Host on REC exit

On execution of RSI_HOST_CALL, arguments are copied from the RsiHostCall object in Realm memory into the RmiRecExit object in NS memory. On the subsequent RMI_REC_ENTER, return values are copied from the RmiRecEnter object in NS memory into the RsiHostCall object in Realm memory.



See also:

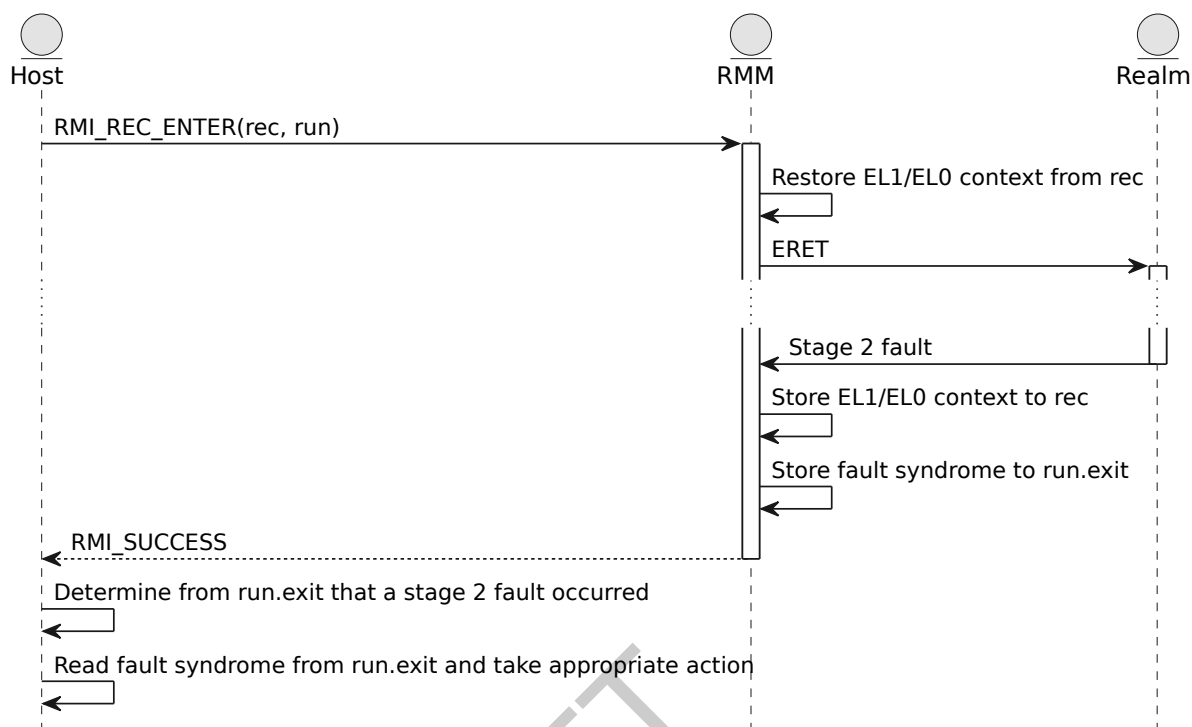
- [A4.5 Host call](#)

D1.3.3 REC exit due to Data Abort fault flow

The following diagram shows how a Data Abort due to a Realm access is taken to the Host.

A REC is entered using the RMI_REC_ENTER command. The parameters to this command include:

- an *RmiRecEnter* object, which is a data structure used to pass values from the Host to the RMM on REC entry
- an *RmiRecExit* object, which is a data structure used to pass values from the RMM to the Host on REC exit




See also:

- [Chapter A4 Realm exception model](#)

D1.3.4 MMIO emulation flow

The following diagram shows how an MMIO access by a Realm can be emulated by the Host.



See also:


- [Chapter A4 Realm exception model](#)

D1.4 PSCI flows

D1.4.1 PSCI_CPU_ON flow

The following diagram shows how one Realm VPE can set the “runnable” flag in another Realm VPE by executing PSCI_CPU_ON.

DRAFT



See also:

- [B4.3.23 RMI_PSCI_COMPLETE command](#)
- [B6.3.3 PSCI_CPU_ON command](#)

DRAFT

D1.5 Realm memory management flows

This section contains flows which relate to management of Realm memory.

See also:

- [Chapter A5 Realm memory management](#)

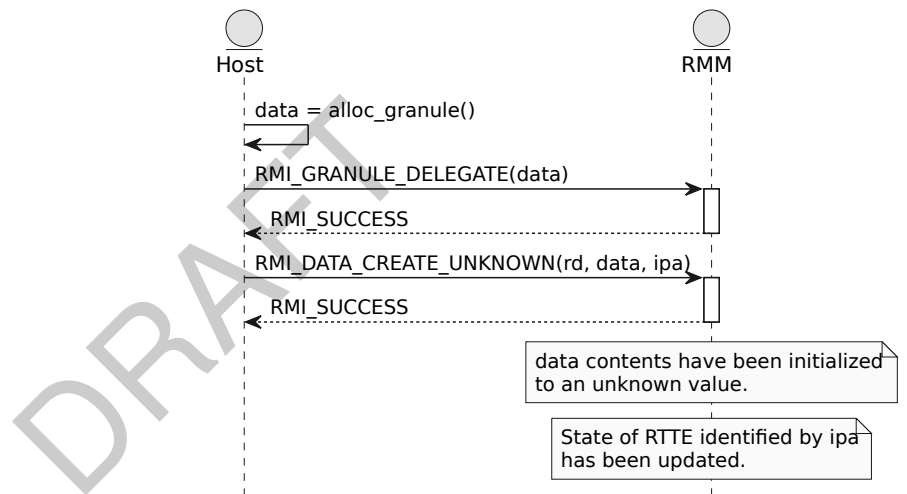
D1.5.1 Add memory to Active Realm flow

The following diagram shows the flow for adding memory to a Realm whose state is REALM_ACTIVE.

To add memory to a Realm whose state is REALM_ACTIVE, the Host must:

- Delegate a destination Granule (*dst*).
- Specify the Protected IPA at which the *dst* Granule will be mapped in the Realm's IPA space.
- Ensure that the level 3 RTT which contains the RTTE identified by the Protected IPA has been created.

Once a given Protected IPA has been populated with unknown content, it cannot be repopulated.




See also:

- [A2.1.5 Realm lifecycle](#)
- [Chapter A5 Realm memory management](#)
- [B4.3.2 RMI_DATA_CREATE_UNKNOWN command](#)
- [B4.3.7 RMI_GRANULE_DELEGATE command](#)

D1.5.2 NS memory flow

The following diagram describes how NS memory can be mapped into a Realm.



See also:

- [Chapter A5 Realm memory management](#)
- [B4.3.42 RMI_RTT_MAP_UNPROTECTED command](#)

- [B4.3.46 RMI_RTT_UNMAP_UNPROTECTED command](#)

D1.5.3 RIPAS change flow

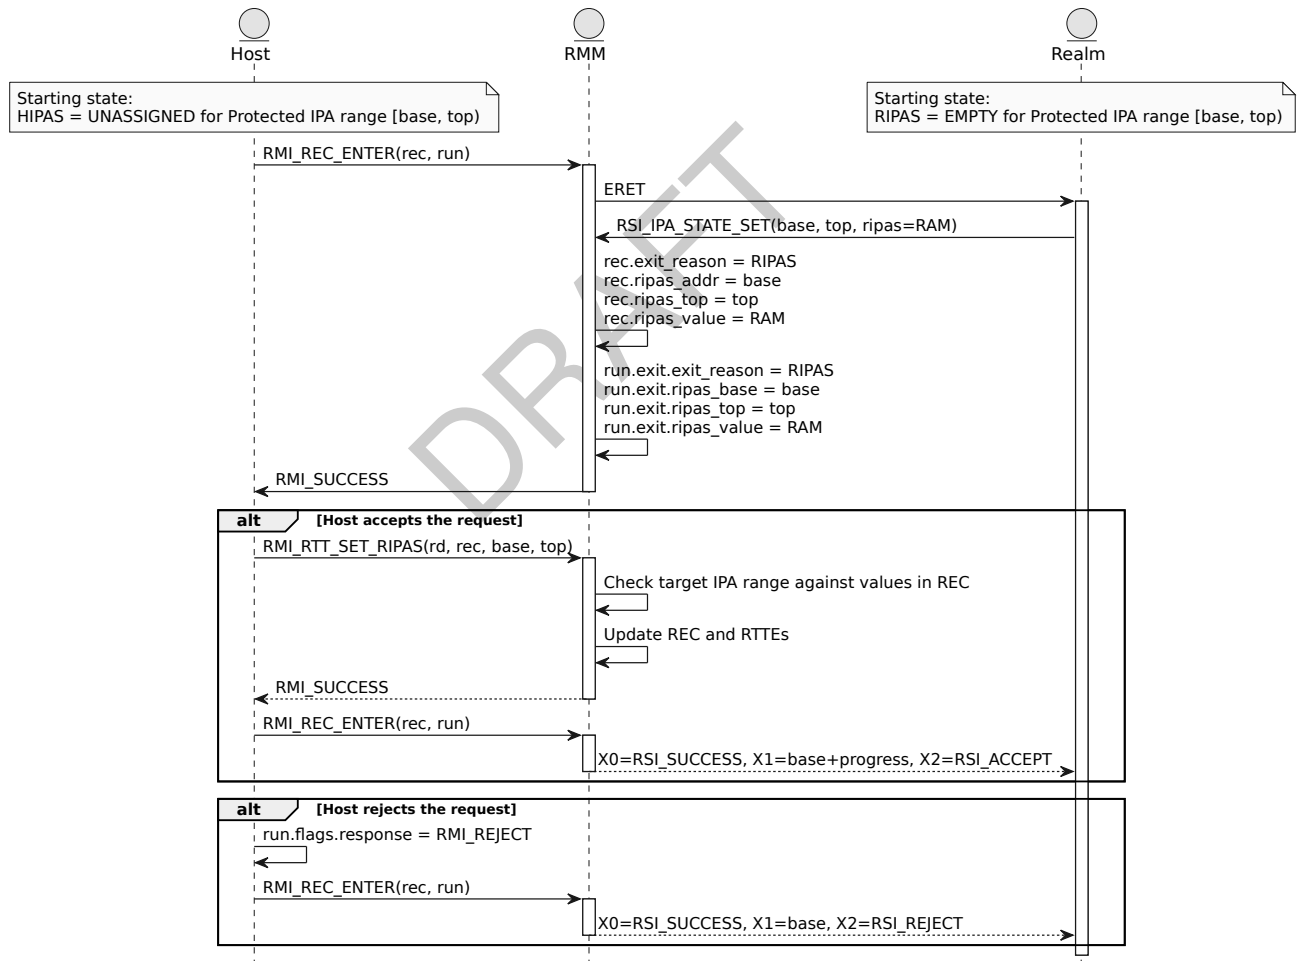
The following diagram describes how a Realm requests a RIPAS change, and how that request is handled by the Host.

- The Realm calls `RSI_IPA_STATE_SET` to request a RIPAS change for IPA range `[base, top)`.
- This causes a REC exit due to RIPAS change pending.

On taking a REC exit due to RIPAS change pending, the Host does the following:

- Reads the region base and top addresses from the `RmiRecExit` object.
- Applies the requested RIPAS change to an IPA range starting from the base of the target region, and extending no further than the top of the target region.
- Calls `RMI_REC_ENTER` to re-enter the REC.

The Realm observes in `X1` the top of the region for which the RIPAS change was applied.



See also:

- [A5.4 RIPAS change](#)
- [B4.3.30 RMI_REC_ENTER command](#)
- [B4.3.44 RMI_RTT_SET_RIPAS command](#)
- [B5.3.6 RSI_IPA_STATE_SET command](#)
- [D2.2 Realm shared memory protocol flow](#)

D1.5.4 S2AP change flow

The following diagram describes how a Realm requests a S2AP change, and how that request is handled by the Host.

- The Realm calls RSI_MEM_SET_PERM_INDEX to request an S2AP change for IPA range [base, top).
- This causes a REC exit due to S2AP change pending.

On taking a REC exit due to S2AP change pending, the Host does the following:

- Reads the region base and top addresses from the RmiRecExit object.
- Applies the requested S2AP change to an IPA range starting from the base of the target region, and extending no further than the top of the target region.
- Calls RMI_REC_ENTER to re-enter the REC.

The Realm observes in X1 the top of the region for which the S2AP change was applied.



See also:

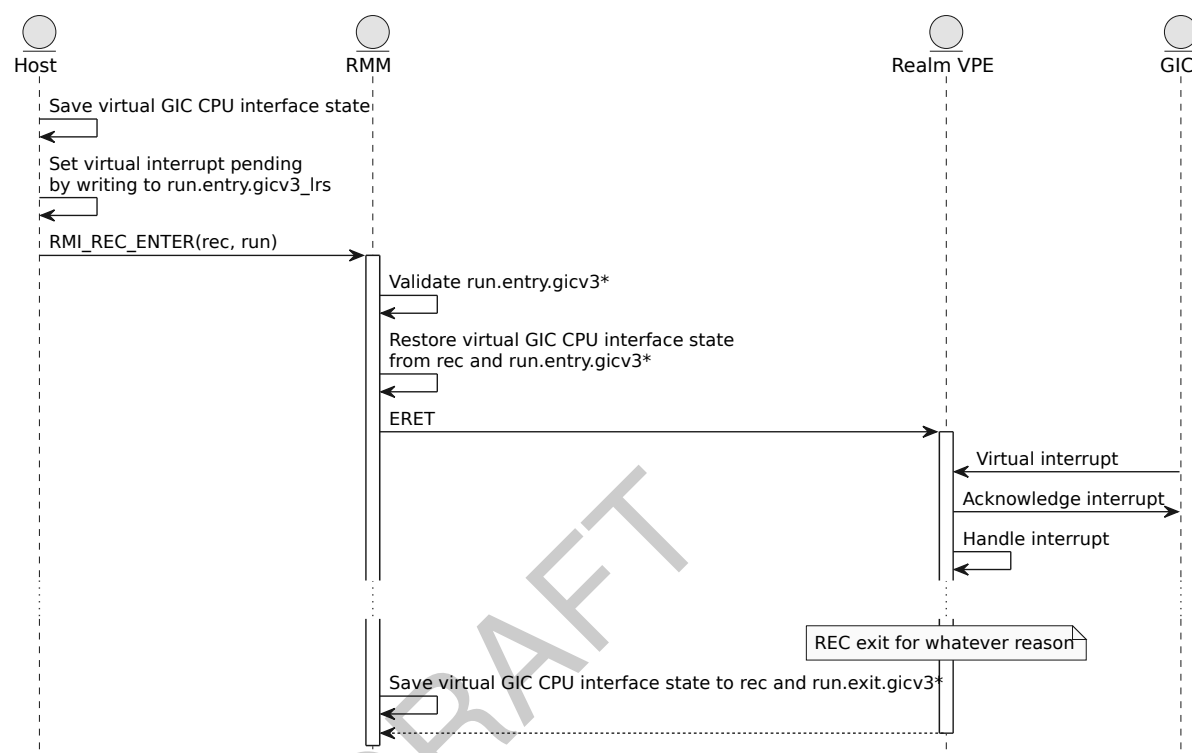
- [A10.3.2.4 Stage 2 access permissions change](#)
- [B4.3.30 RMI_REC_ENTER command](#)
- [B4.3.45 RMI_RTT_SET_S2AP command](#)
- [B5.3.10 RSI_MEM_SET_PERM_INDEX command](#)

DRAFT

D1.6 Realm interrupts and timers flows

D1.6.1 Interrupt flow

The following diagram shows how a virtual interrupt is injected into a Realm by the Host.

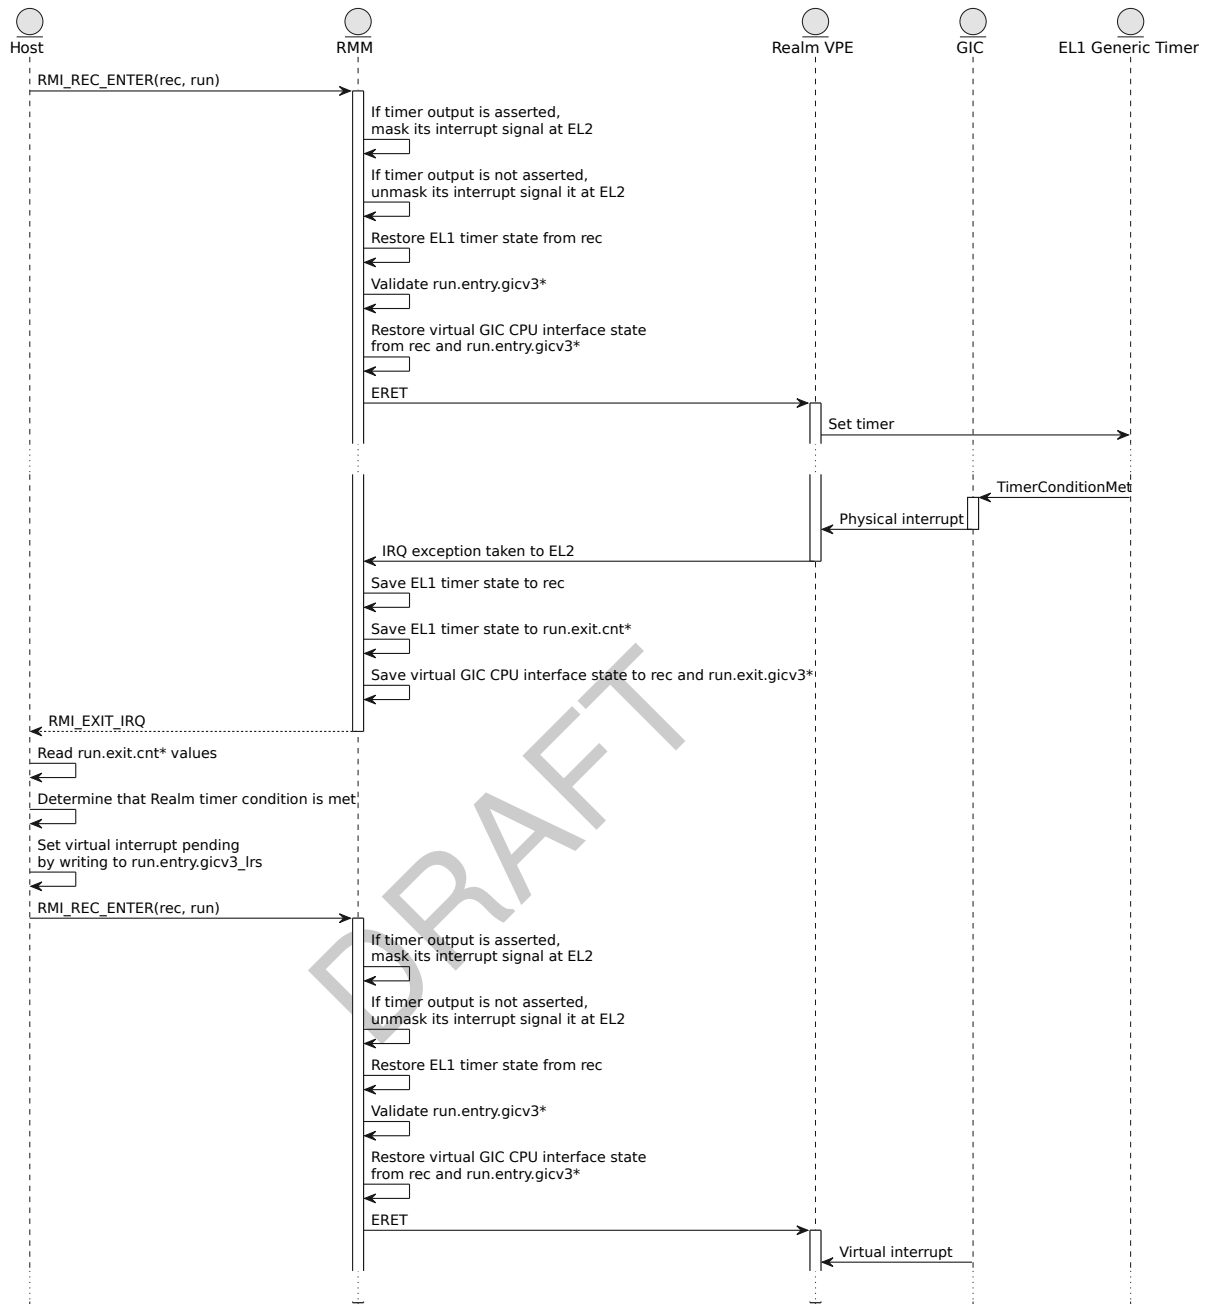


See also:

- [A6.1 Realm interrupts](#)

D1.6.2 Timer interrupt delivery flow

The following diagram shows how a timer interrupt is delivered to and handled by a Realm.



See also:

- [A6.2 Realm timers](#)

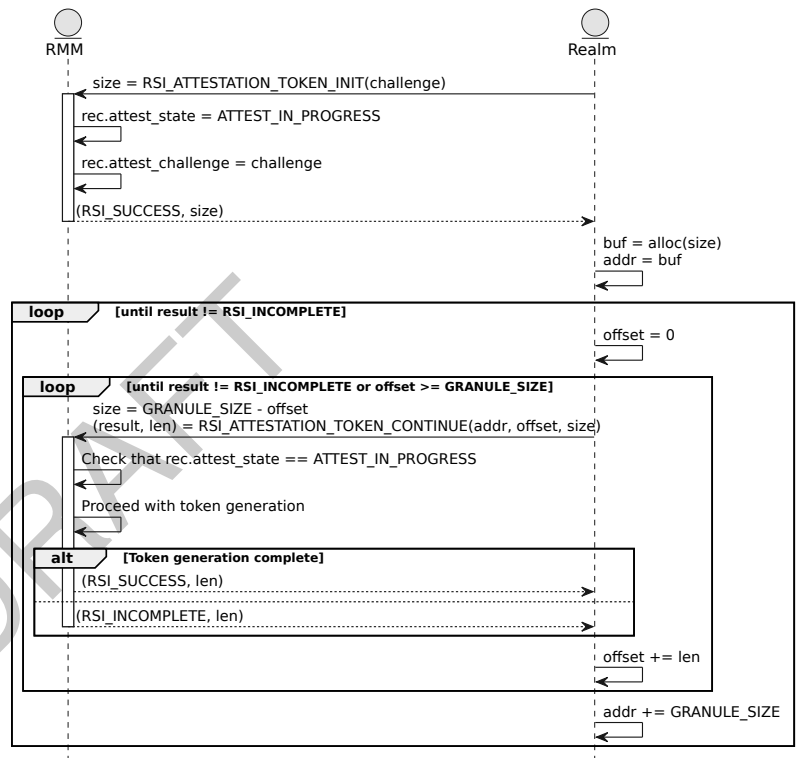
D1.7 Realm attestation flows

D1.7.1 Attestation token generation flow

The following diagram shows the flow for a Realm to obtain an attestation token.

The Realm first calls `RSI_ATTESTATION_TOKEN_INIT`, providing a challenge value. The output values include an upper bound on the attestation token size.

The Realm then calls `RSI_ATTESTATION_TOKEN_CONTINUE`, providing the address of a buffer where the next part of the attestation token will be written. This command is called in a loop, until the result is not `RSI_INCOMPLETE`.



See also:

- [A7.2.2 Attestation token generation](#)
- [B5.3.1 RSI_ATTESTATION_TOKEN_CONTINUE command](#)
- [B5.3.2 RSI_ATTESTATION_TOKEN_INIT command](#)

D1.7.2 Handling interrupts during attestation token generation flow


The following diagram shows how interrupts are handled during generation of an attestation token.

If the RMM detects that a physical interrupt is pending during execution of `RSI_ATTESTATION_TOKEN_CONTINUE`, it saves the execution context to the REC object, and performs a REC exit due to IRQ.

During handling of the IRQ, the Host may signal a virtual interrupt to the REC.

On the next entry to the REC, if a virtual interrupt is pending, it is taken to the REC's exception vector.

Whether or not a virtual interrupt was taken, on return to the original thread, the REC determines that `X0` is `RSI_INCOMPLETE`, and therefore calls `RSI_ATTESTATION_TOKEN_CONTINUE` again.



See also:

- [A4.3.5 REC exit due to IRQ](#)
- [A6.1 Realm interrupts](#)
- [A7.2.2 Attestation token generation](#)
- [B5.3.1 RSI_ATTESTATION_TOKEN_CONTINUE command](#)
- [B5.3.2 RSI_ATTESTATION_TOKEN_INIT command](#)
- [D1.3.1 Realm entry and exit flow](#)

D1.8 Realm device assignment flows

See [Chapter A9 Realm device assignment](#).

DRAFT

Chapter D2

Realm shared memory protocol

This section describes a protocol for management of memory which is shared between a Realm and the Host. This protocol makes use of the primitives described in this specification. However, the protocol itself is not part of the RMM architecture. Use of this protocol is subject to a contract between the Realm and Host software agents.

See also:

- [Chapter A5 Realm memory management](#)

D2.1 Realm shared memory protocol description

The Host agrees to provide the Realm with a certain amount of memory. This memory is referred to below as the Realm's "memory footprint".

The memory footprint is described to the Realm, for example via firmware tables. The Realm can choose, at any point during its execution, how much of its memory footprint is protected (accessible only to the Realm) and how much is shared with the Host.

Realm software treats the most significant IPA bit as a "protection attribute" bit. This means that for every Protected IPA (in which the most significant bit is '0'), there exists a corresponding Unprotected IPA alias, which is generated by setting the most significant bit to '1'.

The choice of whether a given page is protected or shared at a given time is expressed by setting the RIPAS of the Protected IPA:

- If the RIPAS of the Protected IPA is RAM, the page is protected and access to the Unprotected IPA alias causes a Synchronous External Abort taken to the Realm.
- If the RIPAS of the Protected IPA is EMPTY, the page is shared and access to the Unprotected IPA alias does not cause a Synchronous External Abort taken to the Realm.

The initial RIPAS for every page in the Realm's memory footprint is described to the Realm, for example via firmware tables. The Host agrees that during Realm execution, it will accept a RIPAS change request on any page within the Realm's memory footprint.

See also:

- [A5.2.1 Realm IPA space](#)
- [A5.2.2 Realm IPA state](#)
- [A5.4 RIPAS change](#)

D2.2 Realm shared memory protocol flow

The following diagram illustrates how the protocol is used to set up and tear down a shared memory buffer.

Chapter D2. Realm shared memory protocol
D2.2. Realm shared memory protocol flow




Figure D2.1: Realm shared memory protocol flow

See also:

- [D1.5.3 RIPAS change flow](#)

Glossary

ASL

Arm Specification Language
Language used to express pseudocode implementations. Formal language definition can be found in [Arm Specification Language Reference Manual](#) [17].

CBOR

Concise Binary Object Representation

CCA

Confidential Compute Architecture

CCA platform

All hardware and firmware components which are involved in delivering the CCA security guarantee. See [Arm CCA Security model](#) [4].

CDDL

Concise Data Definition Language

COSE

CBOR Object Signing and Encryption

DOE

Data Object Exchange See [PCI Express 6.0 specification](#) [14]

DSM

Device Security Manager See [PCI Express 6.0 specification](#) [14]

EAT

Entity Attestation Token

ECAM

Enhanced Configuration Access Mechanism See [PCI Express 6.0 specification](#) [14]

FAL

Firmware Activity Log

FID

Function Identifier

GIC

Generic Interrupt Controller
See [Arm Generic Interrupt Controller \(GIC\) Architecture Specification version 3 and version 4](#) [6]

GPF

Granule Protection Fault

GPT

Granule Protection Table
Table which determines the Physical Address Space of each Granule.

HIPAS

Host IPA state

Host

Software executing in Non-secure Security state which manages resources used by Realms

IAK

Initial Attestation Key Key used to sign the CCA platform attestation token.

IDE

Integrity and Data Encryption
See [PCI Express 6.0 specification \[14\]](#)

IPA

Intermediate Physical Address
Address space visible to software executing at EL1 in the Realm.

IPI

Inter-processor interrupt

IRI

Interrupt Routing Infrastructure
A subset of the components which make up the GIC.

ITS

Interrupt Translation Service
A service provided by the GIC.

LFA

Live Firmware Activation

MBZ

Must Be Zero

MEC

Memory Encryption Context

MECID

Memory Encryption Context Identifier

MMIO

Memory-mapped I/O

MPIDR

Multiprocessor Affinity Register

NS

Non-secure

PAS

Physical Address Space

PDEV

Physical Device
Object which represents a communication channel between the RMM and a physical device, for example a PCIe device.

PE

Processing Element

PMU

Performance Monitor Unit

PSCI

Power State Control Interface
See [Arm Power State Coordination Interface \(PSCI\) \[22\]](#)

RAK

Realm Attestation Key Key used to sign the Realm attestation token.

RD

Realm Descriptor
Object which stores attributes of a Realm.

RDEV

Realm Device
Object which represents Realm view of an assigned device

Realm

A protected execution environment

REC

Realm Execution Context
Object which stores PE state associated with a thread of execution within a Realm.

REM

Realm Extensible Measurement Measurement value which can be extended during the lifetime of a Realm.

RHA

Realm Hash Algorithm

RIM

Realm Initial Measurement Measurement of the state of a Realm at the time of activation.

RIPAS

Realm IPA state

RME

Realm Management Extension

RMI

Realm Management Interface The ABI exposed by the RMM for use by the Host.

RMM

Realm Management Monitor

RNVS

	Root Non-volatile Storage
RPV	
	Realm Personalization Value
RSI	
	Realm Services Interface The ABI exposed by the RMM for use by the Realm.
RTT	
	Realm Translation Table Object which describes the IPA space of a Realm.
RTTE	
	Realm Translation Table Entry
SBZ	
	Should Be Zero
SEA	
	Synchronous External Abort
SGI	
	Software Generated Interrupt
SMCCC	
	SMC Calling Convention See Arm SMC Calling Convention [16]
SPDM	
	Security Protocol and Data Model See Security Protocol and Data Model (SPDM) [18] and Secured Messages using SPDM Specification version 1.1.0 [15]
SPM	
	Secure Partition Manager
TA	
	Trusted Application
TOS	
	Trusted OS
TSM	
	Trusted Security Manager See Chapter A9 Realm device assignment
VDEV	
	Virtual Device Object which represents the binding between a device function and a Realm.
VMM	
	Virtual Machine Monitor
VMSA	
	Virtual Memory System Architecture

VPE

Virtual Processing Element

Wiping

An operation which changes the value of a memory location from X to Y , such that the value X cannot be determined from the value Y

DRAFT