Arm® Neoverse™ V2 Core

Version 1.0

Telemetry Specification

Non-Confidential

Issue 01
109528_0100_01_en

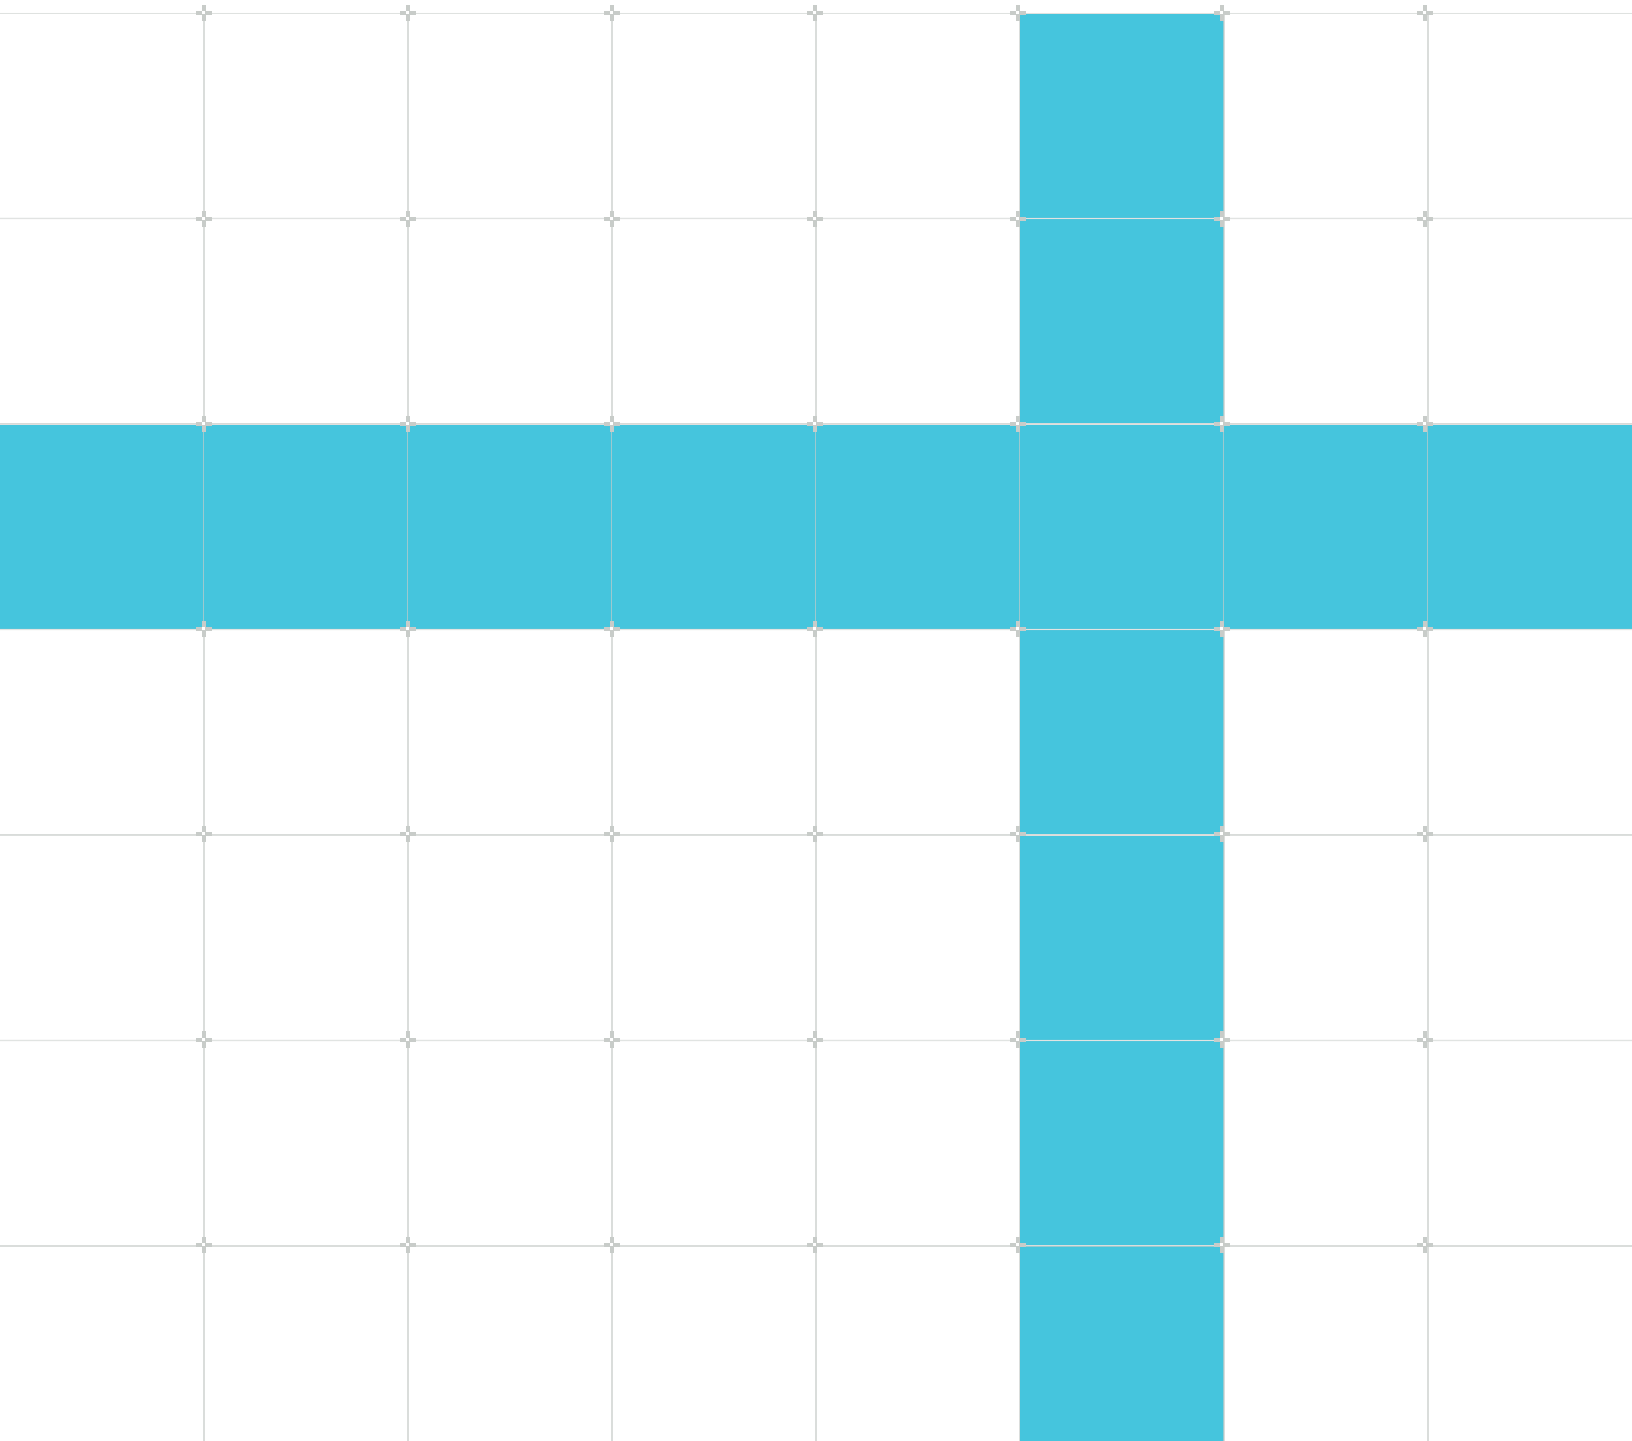

Arm® Neoverse™ V2 Core

Telemetry Specification

Release information

Document history

Issue	Date	Confidentiality	Change
0100-01	15 May 2024	Non-Confidential	Initial release

Proprietary Notice

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.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

PRE-1121-V1.0

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by Arm and the party that Arm delivered this document to.

Unrestricted Access is an Arm internal classification.

Product Status

The information in this document is Final, that is for a developed product.

Feedback

Arm welcomes feedback on this product and its documentation. To provide feedback on the product, create a ticket on https://support.developer.arm.com

To provide feedback on the document, fill the following survey: https://developer.arm.com/documentation-feedback-survey.

Inclusive language commitment

We believe that this document contains no offensive language. To report offensive language in this document, email terms@arm.com.

Contents

1. Introduction

1.1 Conventions

The following subsections describe conventions used in Arm documents.

Glossary

The Arm Glossary is a list of terms used in Arm documentation, together with definitions for those terms. The Arm Glossary does not contain terms that are industry standard unless the Arm meaning differs from the generally accepted meaning.

See the Arm Glossary for more information: developer.arm.com/glossary.

Typographic conventions

Arm documentation uses typographical conventions to convey specific meaning.

Convention	Use
italic	Citations.
bold	Interface elements, such as menu names. Terms in descriptive lists, where appropriate.
`monospace`	Text that you can enter at the keyboard, such as commands, file and program names, and source code.
`monospace underline`	A permitted abbreviation for a command or option. You can enter the underlined text instead of the full command or option name.
`<and>`	Encloses replaceable terms for assembler syntax where they appear in code or code fragments. For example: `MRC p15, 0, <Rd>, <CRn>, <CRm>, <Opcode_2>`
SMALL CAPITALS	Terms that have specific technical meanings as defined in the *Arm® Glossary*. For example, **IMPLEMENTATION DEFINED**, **IMPLEMENTATION SPECIFIC**, **UNKNOWN**, and **UNPREDICTABLE**.

We recommend the following. If you do not follow these recommendations your system might not work.

Your system requires the following. If you do not follow these requirements your system will not work.

You are at risk of causing permanent damage to your system or your equipment, or harming yourself.

This information is important and needs your attention.

A useful tip that might make it easier, better or faster to perform a task.

A reminder of something important that relates to the information you are reading.

1.2 Useful resources

This document contains information that is specific to this product. See the following resources for other useful information.

Access to Arm documents depends on their confidentiality:

- Non-Confidential documents are available at developer.arm.com/documentation. Each document link in the following tables goes to the online version of the document.
- Confidential documents are available to licensees only through the product package.

Arm product resources	Document ID	Confidentiality
Arm® CPU Telemetry Solution Topdown Methodology Specification	109542	Non-Confidential
Arm® Neoverse® V2 Core Technical Reference Manual	102375	Non-Confidential
Arm® Telemetry on Arm Developer	–	Non-Confidential
Arm® Telemetry Solution GitLab repository	–	Non-Confidential

Arm architecture and specifications	Document ID	Confidentiality
Arm® Architecture Reference Manual for A-profile architecture	DDI 0487	Non-Confidential

1.3 Other information

See the Arm website for other relevant information.

- Arm® Developer.
- Arm® Documentation.
- Technical Support.
- Arm® Glossary.

2. Overview of the Neoverse V2 core Telemetry methodology

The *Arm® Neoverse V2 Core Telemetry Specification* describes the Topdown methodology, derived metrics, and Performance Monitoring Unit (PMU) events supported by the Arm Neoverse V2 core, also known as the processor.

This specification is applicable to all releases of the product.

This specification implements the framework provided by the *Arm® CPU Telemetry Solution Topdown Methodology Specification*, which is referred to as the Architecture Specification. The reader is expected to read this document in conjunction with the Architecture Specification.

Arm Telemetry framework

This specification outlines the telemetry features implemented for the Arm Neoverse V2 core and follows the Arm Telemetry framework for CPUs defined in the Architecture Specification.

The following list provides a brief description of the Telemetry framework:

Events

Hardware performance monitoring events implemented by the core that contain raw data read from the registers or memory buffers.

Metrics

Derived mathematical relationships between events that provide insight into the system behavior. They are developed to abstract hardware details of the events from consumers of the telemetry data.

Metric groups

Group of metrics that can be analyzed together to investigate a bottleneck scenario or a specific resource in a given system.

Methodology

Actionable guidance, such as Arm Topdown methodology, to explain how to consume the different metrics and events for a specific usage model. Decision tree with a group of metrics that can be analyzed hierarchically to investigate a bottleneck scenario or a specific resource in a given system.

Tool support for profiling and monitoring

This specification is also available in a machine-readable format (JSON) to be consumed by profiling and monitoring tools. The JSON schema implements the Arm Telemetry framework from the Architecture Specification.

The JSON for the core is published in the open source *Arm® Telemetry Solution GitLab repository*.

2.1 Documentation and resources

Arm products include a set of documents.

The documentation and resources for the Neoverse V2 core consist of:

- *Arm® Telemetry on Arm Developer*
- *Arm® Telemetry Solution GitLab repository*
- *Arm® Neoverse® V2 Core Technical Reference Manual*

3. Telemetry features of Neoverse V2 core

The Neoverse V2 core is a super pipelined superscalar processor that has an in-order frontend and out-of-order backend.

The following figure shows the microarchitecture details of the Neoverse V2 core.

Figure 3-1: Neoverse V2 microarchitecture

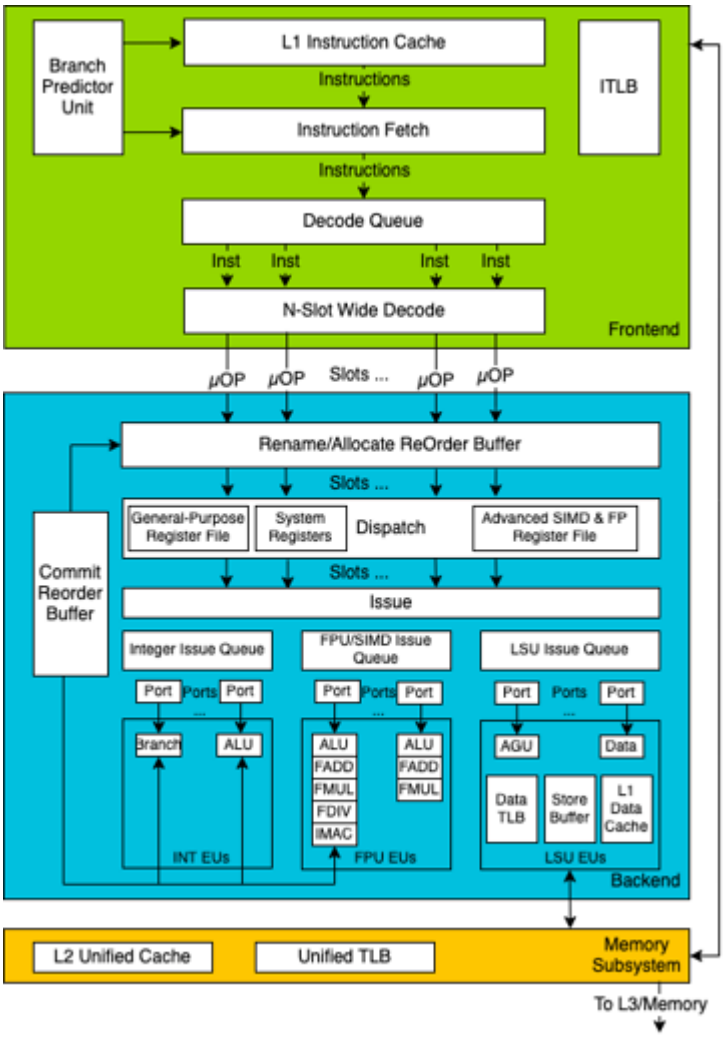

The frontend of the core pipeline is comprised of the instruction fetch and decode units. The frontend also includes a branch predictor unit that predicts branch target addresses and fetches instructions ahead of the pipeline. This unit helps to hide latencies caused by control flow bubbles in the pipeline. The fetch unit can fetch multiple instructions for each cycle whose bandwidth is specific to a microarchitecture design, which gets stored in a decode queue. The decode queue sends multiple instructions per cycle for decoding, whose bandwidth is determined by the number of decode slots available. The decode unit decomposes the Arm architecture instructions into micro-operations, also known as micro-ops or (µops). The decode unit decodes more than one

micro-operation for each cycle, which are then fed to the rename unit for organization for out-of-order execution in the backend of the core. The bandwidth is determined by the number of renamed slots available in the microarchitecture. From a microarchitecture standpoint, the rename unit is considered the boundary between the frontend and backend of the core.

The backend of the core has a scheduler that orchestrates the operations to be executed when the issue queue associated with the operation can accept the operation. The issue queue sends operations for execution when the execution unit is free and the source operands are ready. Once the execution is complete, the results are sent to the commit Reorder Buffer (ROB) from where the instructions are retired when the speculated execution is confirmed. The backend of the core executes the operations out-of-order and stores results with the help of the reorder buffer. The dispatch unit tracks dependencies between operations and determines the operand availability for the execution of operations. Register renaming is occurs at this stage to mitigate data dependency hazards.

In the dispatch unit, issue queues are employed for:

- Queuing the micro-operations (µops) to assigned ports
- Managing dependencies between operations
- Tracking operand availability for execution

Each execution port supports different categories of operations. After the execution of operations, the ROB is updated with execution results. Completed operations are retired architecturally in the right program order. Operations are flushed when the predicted program flow changes due to mispredictions or exceptions.

The Memory subsystem of the core handles the execution of load and store operations which rely heavily on the memory hierarchy levels. The Neoverse V2 core has dedicated cache levels, L1 and L2 for each core, where the L2 cache is shared between the L1 data cache and the L1 instruction cache. The Load Store Unit controls the data flow between the caches and to memory. The core has multiple load/store units, which can handle both read and write operations. L1 and L2 caches are set-associative. The size of the cache is configurable for each implementation and determines the number of sets in each way. The private L2 cache of the core connects to the rest of the system through an AMBA® 5 CHI interface.

Neoverse V2 core system configurations

All systems with the Arm® Neoverse™ Coherent Mesh Network support a shared system-level cache. Understanding the cache hierarchy and configuration of the system being analyzed is crucial in deriving insights from the cache effectiveness Performance Monitoring Unit (PMU) events.

It is always best to check with the Silicon Provider for details on the system configuration for the underlying system, including the cache sizes.

PMU capabilities of Neoverse V2

The Neoverse V2 core implements version 3.5 of the Performance Monitors Extension, `FEAT_PMUv3p5`, and Arm v8.4 debug architecture, `FEAT_Debugv8p4`.

For more information, see *Arm® Architecture Reference Manual for A-profile architecture*.

The Neoverse V2 PMU has six configurable counter registers and one dedicated function counter to count CPU cycles.

4. CPU performance analysis methodology

The Arm Topdown methodology for performance analysis and microarchitecture exploration is conducted in two stages.

Stage 1

The first stage is to perform Topdown analysis. It uses hierarchical pipeline stall-related metrics to detect and identify the performance bottleneck in the CPU. For more information, see Stage 1: Topdown analysis.

Stage 2

The second stage is to conduct microarchitecture exploration to further analyze bottlenecked CPU resources. It uses a set of CPU resource effectiveness metrics. For more information, see Stage 2: Microarchitecture Exploration.

For more information, see *Arm® CPU Telemetry Solution Topdown Methodology Specification*.

4.1 Topdown methodology for the Neoverse V2 core

Topdown analysis helps with hotspot detection and is the first stage completed in the Topdown methodology. Microarchitecture exploration is the second stage completed in the Topdown methodology and conducts a micro-architectural analysis of the bottlenecking CPU resource.

The following figure shows the Topdown methodology for the Neoverse V2 core covering both the Stage 1 metrics and the Stage 2 metric groups that can be used during your analysis.

Figure 4-1: Neoverse V2 Topdown methodology overview

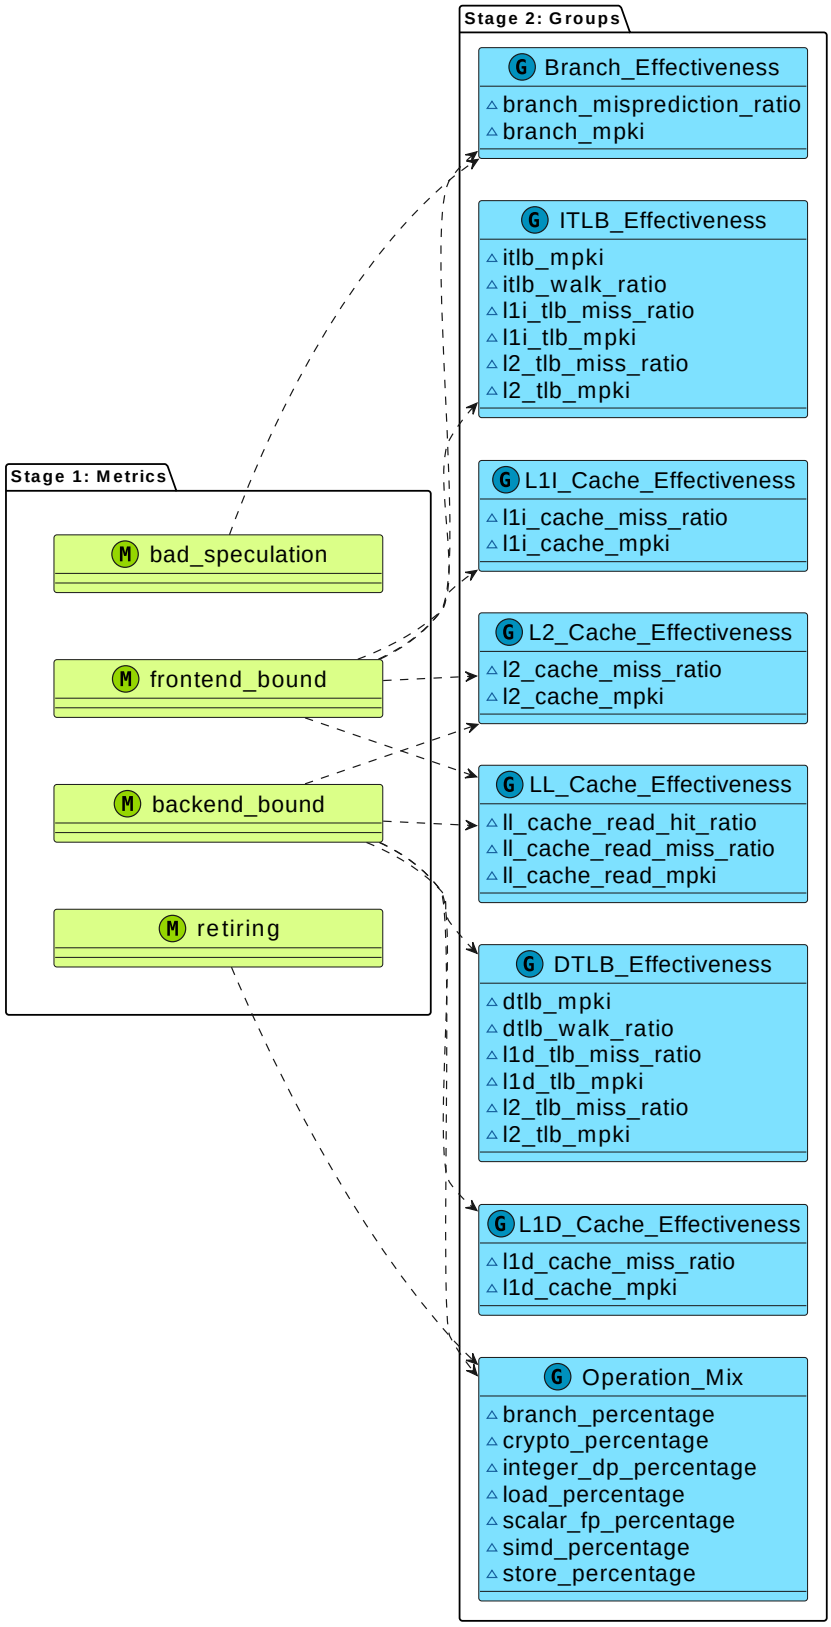

Arm recommends collecting all the metrics that are in Stage 1 and Stage 2 for workload characterization. For further analysis, Arm recommends a set of set of microarchitecture exploration metric groups against some of the hotspots detected in Stage 1. All the Stage 2 metrics can be used to derive further insights into the overall microarchitecture behavior during the execution of the application under investigation and can be used independently to Stage 1.

4.2 Stage 1: Topdown analysis

The Neoverse V2 core supports the four key metrics for Topdown analysis level 1 that are slot-based, which is a measurement of the efficiency of pipeline slots.

The four metrics in the first level as part of the Topdown Level 1 metric group is defined by the *Arm® CPU Telemetry Solution Topdown Methodology Specification*, as follows:

frontend_bound
: This metric is the percentage of total slots that were stalled due to resource constraints in the frontend unit of the processor.

backend_bound
: This metric is the percentage of total slots that were stalled due to resource constraints in the backend unit of the processor.

bad_speculation
: This metric is the percentage of total slots that executed operations but did not retire due to a pipeline flush caused by mis-speculation. It indicates the cycles that were used but were inefficient executing the wrong instructions. It also includes cycles spent recovering from the pipeline flush, which requires an instruction pipeline refill from the correct instruction location.

retiring
: This metric is the percentage of total slots that retired operations. This indicates the cycles that were used and efficient.

For more information on the Topdown Level 1 metric group and its corresponding metrics, see Topdown_L1.

4.3 Stage 2: Microarchitecture exploration

When the execution pipeline bottleneck region is identified from Stage 1, the next step is to investigate the CPU resources for further analysis.

As described in the Architecture Specification:

- A relatively high frontend_bound metric shows that execution cycles are wasted due to pipeline stalls in the in-order frontend division of the processor. There are many reasons why frontend stalls can occur, such as inefficiency in the branch prediction unit, fetch latency due to instruction cache misses, and translation delays caused by Instruction TLB walks.

- A relatively high backend_bound metric shows that execution cycles are wasted due to pipeline stalls in the backend of the processor. There are many reasons why backend stalls can occur, such as inefficiency in execution units, data cache misses, and translation delays caused by data TLB walks.
- A relatively high bad_speculation metric shows the pipeline stalls caused by flushes or machine clears that break the pipeline needing a control flow change. Branch mispredictions are one of the major causes for these stalls, as well as exceptions.

The Neoverse V2 core supports the following microarchitecture exploration metric groups that can be used for Stage 2 analysis. These metric groups support further analysis of the bottlenecking CPU resources following the Stage 1 analysis hotspot.

As a common step in Stage 2, the two metric groups MPKI and Miss rate are recommended for a quick behavioral analysis of the CPU components that could be the most probable bottlenecks.

MPKI – Misses Per Kilo Instructions

Misses Per Kilo Instructions is a set of metrics that can be derived to normalize the misses in CPU components, such as branches, caches, and TLBs against the total instructions executed. This is an industry-standard metric that also helps with comparison across different implementations of the Arm architecture, as instructions retired should count the same on all AArch64-based microarchitectures.

MPKI lists all the Misses Per Kilo Instructions metrics that can be derived for the Neoverse V2 core and their formulae.

Miss Ratio

The Miss Ratio metric group provides a set of metrics that calculate the ratio of the misses in the CPU components, such as branches, caches and TLBs against the total accesses in those components. These metrics provide insights on the efficiency of each CPU component in the pipeline and help to root cause issues.

Miss_Ratio lists all the Miss Ratio metrics that can be derived for the Neoverse V2 core and their formulae.

Operation Mix

The Neoverse V2 microarchitecture has a variety of execution units that process different types of operations.

The execution units that are shown in Figure 3-1: Neoverse V2 microarchitecture on page 12 process the following types of operations:

- Branch
- Single-cycle integers
- Multicycle integers
- Load/store unit with address generation
- Advanced floating-point/SIMD operations

The PMU events listed in Spec_Operation count the operations that are issued to these execution units.

Operation_Mix lists all the Speculative Operation Mix metrics that can be derived for the Neoverse V2 core and their formulae.

These metrics use events that count speculatively issued operations at the issue stage, which provide an estimate of the execution unit utilization, but not the retired instruction mix of the program. To derive the utilization of each operation type, the percentage of each type of operation is calculated as a fraction of the total operations issued, which is counted by the event INST_SPEC.

The Neoverse V2 core does not support retired events for counting the architectural instruction mix. The core supports events to further break down the branch operations into immediate, indirect, and return branches, counted by events BR_IMMED_SPEC, BR_INDIRECT_SPEC, and BR_RETURN_SPEC respectively. BR_RETURN_SPEC is a subset of BR_INDIRECT_SPEC, as returns are also counted as indirect branches. The sum of the BR_IMMED_SPEC and BR_INDIRECT_SPEC branch operation events can compute the total branches executed.

Branch Effectiveness and Branch Mix

Branch mispredictions are costly in a deeply pipelined CPU, causing pipeline flushes and wasted cycles. As a general rule, workloads typically contain, on average, one branch in every six instructions.

Though modern CPUs have optimized branch prediction units, there are many use cases such as ray tracing and decision tree algorithms that are branch heavy and hard to predict. In some of these applications, there can be hundreds of unique branch paths to take and the target may be input data dependent.

Branch prediction performance can be evaluated using two PMU events, BR_MIS_PRED_RETIRED and BR_RETIRED. BR_MIS_PRED_RETIRED provides an account of the total branches that were executed but mispredicted. This means that the direction of the code path was wrong and the following operations in the path were wasted, causing a pipeline flush. BR_RETIRED counts the total branches architecturally executed by the CPU.

There are two performance metrics that can be derived for a high-level evaluation of the branch execution performance regarding the overall program execution:

branch_mpki metrics

Provides total branch mispredictions per kilo instructions

branch_misprediction_ratio metrics

Provides an indication of the ratio of branches that were mispredicted to overall branches

Branch_Effectiveness lists all the Branch Effectiveness metrics that can be derived for the Neoverse V2 core and their formulae.

Branch prediction units work differently depending on the branch type. The following list describes the three main sub-units that work for different branch types:

- *Branch History Table* (BHT) stores the history of conditional branches, taken or not.

- *Branch Target Buffer* (BTB) stores the target address for indirect branches.
- *Return Address Stack* (RAS) stores the function return branches.

The Neoverse V2 core supports the following three events that respectively categorizes the immediate, indirect, and return branches executed:

- BR_IMMED_SPEC
- BR_RETURN_SPEC
- BR_INDIRECT_SPEC

Getting a breakdown of the branch type helps to investigate each of these sub-units within the branch prediction unit.

TLB/MMU Effectiveness

Arm recommends checking the virtual memory system performance which affects the instruction fetch performance in the frontend and memory access performance on the data side.

The processor needs to translate a virtual address to physical address for any instruction/data memory access before it accesses the respective cache. A program's view of memory is virtual address, but the processor works with the physical address when accessing cache or memory.

Virtual to physical mappings are defined in the page translation tables which reside in system memory. Accessing these tables requires one or more memory accesses that take many cycles to complete, which is referred to as a translation table walk. However, to make these translations faster, *Translation Lookaside Buffers* (TLBs) cache translation table walks, greatly reducing the number of accesses to system memory.

Neoverse V2 implements a two-level TLB hierarchy. The first level contains separate, dedicated TLBs for the instruction and data (load/store) address translations. Total accesses to these TLBs are counted by L1I_TLB and L1D_TLB respectively. The second level contains a unified L2 TLB that is shared by both instruction-side and data-side accesses. There are corresponding REFILL counters for these TLB levels. Some performance metrics that can be derived for a high-level evaluation of the TLB execution performance are the l<n>_tlb_mpki and l<n>_tlb_miss_rate metrics, where <n> stands for each level of TLB instruction and data side.

Accesses that cause a translation table walk due to misses in the instruction side and data side TLBs are counted by the events ITLB_WALK and DTLB_WALK, respectively. To evaluate the TLB effectiveness and cost of latency caused by translation table walks specifically, dtlb_mpki, dtlb_walk_ratio, itlb_mpki, and itlb_walk_ratio are the key metrics that can be derived. itlb_mpki and dtlb_mpki provide the rate of TLB Walks per kilo instructions for instruction and data accesses respectively. These derived metrics help to evaluate and correlate the TLB efficiency with respect to the total instructions.

dtlb_walk_ratio provides the ratio of DTLB Walks to the overall TLB lookups made by the program. This is the same as DTLB_WALK/MEM_ACCESS as every MEM_ACCESS causes a L1D_TLB access. itlb_walk_ratio provides a percentage of ITLB walks to the overall TLB lookups initiated from the instruction side.

ITLB_Effectiveness and DTLB_Effectiveness list all the TLB Effectiveness metrics that can be derived for Neoverse V2 and their formulae.

Cache Effectiveness

The Neoverse V2 implements a multi-level cache hierarchy.

The first level (L1) includes a dedicated cache for instructions and a separate dedicated cache for data accesses. The second level (L2) is a unified L2 cache that is shared between code and data. Further down the hierarchy, the system could have an optional shared *System Level Cache* (SLC) in the interconnect. Arm recommends checking the cache configurations with the platform provider.

The Neoverse V2 core supports hierarchical PMU events for all the cache hierarchy levels. For each level of caches, there are total access counts and refill counts. AArch64 does not support cache MISS counters, but only REFILLs. A cache miss could lead to multiple cache line refills if the access is on a cache line boundary or multiple cache misses could be satisfied by a single REFILL.

Some performance metrics that can be derived for a high-level evaluation of the cache execution behavior are the l<n>_cache_mpki and l<n>_cache_miss_ratio metrics, where <n> stands for each level of instruction and data caches.

L1I_Cache_Effectiveness, L1D_Cache_Effectiveness, and L2_Cache_Effectiveness list all the L1 and L2 Cache Effectiveness metrics that can be derived for Neoverse V2 and their formulae.

Core Memory Traffic

The MEM_ACCESS event counts the total number of memory operations that were issued by the *Load Store Unit* (LSU) of the core. As these operations are looked up in the L1D_CACHE first, both the events L1D_CACHE and MEM_ACCESS count at the same rate.

Neoverse V2 also supports two additional events, MEM_ACCESS_RD and MEM_ACCESS_WR that can provide the read and write traffic breakdown respectively. Note that these events are not the same as LD_SPEC and ST_SPEC as they count memory operations speculatively issued but not necessarily executed.

Last Level Cache Counter Usage

On systems that support a shared SLC in the interconnect, LL_CACHE_RD counts the total accesses to the SLC. In a system that has the SLC configured to count LL_CACHE_RD events, LL_CACHE_RD counts total SLC accesses made by the core, and LL_CACHE_MISS_RD counts accesses missed at the SLC.

The last level cache read miss metrics ll_cache_read_mpki and ll_cache_miss_ratio can be derived to study the last level read behavior. Another useful metric to measure is the SLC hit percentage for the read traffic is the SLC Read Hit Ratio denoted as ll_cache_read_hit_ratio. Last level cache events do not have a write variant in Neoverse V2, since SLC is only used as an eviction cache for the core. In addition, all the writes complete early at the interconnect when the transaction is acknowledged but not necessarily completed.

LL_Cache_Effectiveness lists all the Last Level Cache Effectiveness metrics that can be derived for Neoverse V2 and their formulae.

Remote Cache Access

For systems with multiple sockets or SoCs, Neoverse V2 supports the REMOTE_ACCESS event, which counts the memory transactions that were completed by a subordinate source from another chip.

5. Neoverse V2 Telemetry cheat-sheets and lookup tables

The cheat-sheets and lookup tables enable you to find and access metrics and events in different ways.

Cheat-sheets

Both metrics and events are listed by metric groups.

Lookup tables

Metrics are listed alphabetically, with the related events, and metric groups.

Events are listed by code number, with the related metrics, metric groups, and functional groups.

5.1 Metrics cheat sheet for Neoverse V2

Metrics are listed in their respective metric groups. Some metrics are used in more than one metric group.

Neoverse V2 specification provides the following types of metrics:

- Total implemented Common metrics: 35

Topdown Level 1 (4)	Cycle Accounting (2)	General (1)
• backend_bound • bad_speculation • frontend_bound • retiring	• backend_stalled_cycles • frontend_stalled_cycles	• ipc

Misses Per Kilo Instructions (10)	Miss Ratio (10)	Branch Effectiveness (2)
• branch_mpki • dtlb_mpki • itlb_mpki • l1d_cache_mpki • l1d_tlb_mpki • l1i_cache_mpki • l1i_tlb_mpki • l2_cache_mpki • l2_tlb_mpki • ll_cache_read_mpki	• branch_misprediction_ratio • dtlb_walk_ratio • itlb_walk_ratio • l1d_cache_miss_ratio • l1d_tlb_miss_ratio • l1i_cache_miss_ratio • l1i_tlb_miss_ratio • l2_cache_miss_ratio • l2_tlb_miss_ratio • ll_cache_read_miss_ratio	• branch_misprediction_ratio • branch_mpki

Instruction TLB Effectiveness (6)	Data TLB Effectiveness (6)	L1 Instruction Cache Effectiveness (2)
• itlb_mpki	• dtlb_mpki	• l1i_cache_miss_ratio
• itlb_walk_ratio	• dtlb_walk_ratio	• l1i_cache_mpki
• l1i_tlb_miss_ratio	• l1d_tlb_miss_ratio	
• l1i_tlb_mpki	• l1d_tlb_mpki	
• l2_tlb_miss_ratio	• l2_tlb_miss_ratio	
• l2_tlb_mpki	• l2_tlb_mpki	

L1 Data Cache Effectiveness (2)	L2 Unified Cache Effectiveness (2)	Last Level Cache Effectiveness (3)
• l1d_cache_miss_ratio	• l2_cache_miss_ratio	• ll_cache_read_hit_ratio
• l1d_cache_mpki	• l2_cache_mpki	• ll_cache_read_miss_ratio
		• ll_cache_read_mpki

Speculative Operation Mix (7)
• branch_percentage
• crypto_percentage
• integer_dp_percentage
• load_percentage
• scalar_fp_percentage
• simd_percentage
• store_percentage

5.2 PMU events cheat sheet for Neoverse V2

Events are listed in their respective metric groups. Some events are not used in the Methodology, therefore are not shown in the cheat sheet.

Neoverse V2 specification provides the following types of PMU events:

- Total implemented Common events: 155
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Topdown Level 1 (7)	Cycle Accounting (3)	General (2)
• BR_MIS_PRED	• CPU_CYCLES	• CPU_CYCLES
• CPU_CYCLES	• STALL_BACKEND	• INST_RETIRED
• OP_RETIRED	• STALL_FRONTEND	
• OP_SPEC		
• STALL_SLOT		
• STALL_SLOT_BACKEND		
• STALL_SLOT_FRONTEND		

Misses Per Kilo Instructions (11)	Miss Ratio (18)	Branch Effectiveness (3)
• BR_MIS_PRED_RETIRED • DTLB_WALK • INST_RETIRED • ITLB_WALK • L1D_CACHE_REFILL • L1D_TLB_REFILL • L1I_CACHE_REFILL • L1I_TLB_REFILL • L2D_CACHE_REFILL • L2D_TLB_REFILL • LL_CACHE_MISS_RD	• BR_MIS_PRED_RETIRED • BR_RETIRED • DTLB_WALK • ITLB_WALK • L1D_CACHE • L1D_CACHE_REFILL • L1D_TLB • L1D_TLB_REFILL • L1I_CACHE • L1I_CACHE_REFILL • L1I_TLB • L1I_TLB_REFILL • L2D_CACHE • L2D_CACHE_REFILL • L2D_TLB • L2D_TLB_REFILL • LL_CACHE_MISS_RD • LL_CACHE_RD	• BR_MIS_PRED_RETIRED • BR_RETIRED • INST_RETIRED

Instruction TLB Effectiveness (6)	Data TLB Effectiveness (6)	L1 Instruction Cache Effectiveness (3)
• INST_RETIRED • ITLB_WALK • L1I_TLB • L1I_TLB_REFILL • L2D_TLB • L2D_TLB_REFILL	• DTLB_WALK • INST_RETIRED • L1D_TLB • L1D_TLB_REFILL • L2D_TLB • L2D_TLB_REFILL	• INST_RETIRED • L1I_CACHE • L1I_CACHE_REFILL

L1 Data Cache Effectiveness (3)	L2 Unified Cache Effectiveness (3)	Last Level Cache Effectiveness (3)
• INST_RETIRED • L1D_CACHE • L1D_CACHE_REFILL	• INST_RETIRED • L2D_CACHE • L2D_CACHE_REFILL	• INST_RETIRED • LL_CACHE_MISS_RD • LL_CACHE_RD

Speculative Operation Mix (9)
• ASE_SPEC • BR_IMMED_SPEC • BR_INDIRECT_SPEC • CRYPTO_SPEC • DP_SPEC • INST_SPEC • LD_SPEC • ST_SPEC • VFP_SPEC

5.3 Metrics lookup table for Neoverse V2

All metrics are listed alphabetically, with the related events, and metric groups. Some metrics are used in more than one metric group, in that case they are listed multiple times so that you can jump to the most relevant metric group for your requirements.

Table 5-11: Metrics listed by name, with related events and metric groups

Metric Name	Formula from Events	Metric Groups
backend_bound	100 * (STALL_SLOT_BACKEND / (CPU_CYCLES * 8) - BR_MIS_PRED * 3 / CPU_CYCLES)	• Topdown_L1
backend_stalled_cycles	STALL_BACKEND / CPU_CYCLES * 100	• Cycle_Accounting
bad_speculation	100 * ((1 - OP_RETIRED / OP_SPEC) * (1 - STALL_SLOT / (CPU_CYCLES * 8)) + BR_MIS_PRED * 4 / CPU_CYCLES)	• Topdown_L1
• branch_misprediction_ratio in Branch_Effectiveness • branch_misprediction_ratio in Miss_Ratio	BR_MIS_PRED_RETIRED / BR_RETIRED	• Branch_Effectiveness • Miss_Ratio
• branch_mpki in Branch_Effectiveness • branch_mpki in MPKI	BR_MIS_PRED_RETIRED / INST_RETIRED * 1000	• Branch_Effectiveness • MPKI
branch_percentage	(BR_IMMED_SPEC + BR_INDIRECT_SPEC) / INST_SPEC * 100	• Operation_Mix
crypto_percentage	CRYPTO_SPEC / INST_SPEC * 100	• Operation_Mix
• dtlb_mpki in DTLB_Effectiveness • dtlb_mpki in MPKI	DTLB_WALK / INST_RETIRED * 1000	• DTLB_Effectiveness • MPKI
• dtlb_walk_ratio in DTLB_Effectiveness • dtlb_walk_ratio in Miss_Ratio	DTLB_WALK / L1D_TLB	• DTLB_Effectiveness • Miss_Ratio
frontend_bound	100 * (STALL_SLOT_FRONTEND / (CPU_CYCLES * 8) - BR_MIS_PRED / CPU_CYCLES)	• Topdown_L1
frontend_stalled_cycles	STALL_FRONTEND / CPU_CYCLES * 100	• Cycle_Accounting
integer_dp_percentage	DP_SPEC / INST_SPEC * 100	• Operation_Mix

Metric Name	Formula from Events	Metric Groups
ipc	INST_RETIRED / CPU_CYCLES	• General
• itlb_mpki in ITLB_Effectiveness • itlb_mpki in MPKI	ITLB_WALK / INST_RETIRED * 1000	• ITLB_Effectiveness • MPKI
• itlb_walk_ratio in ITLB_Effectiveness • itlb_walk_ratio in Miss_Ratio	ITLB_WALK / L1I_TLB	• ITLB_Effectiveness • Miss_Ratio
• l1d_cache_miss_ratio in L1D_Cache_Effectiveness • l1d_cache_miss_ratio in Miss_Ratio	L1D_CACHE_REFILL / L1D_CACHE	• L1D_Cache_Effectiveness • Miss_Ratio
• l1d_cache_mpki in L1D_Cache_Effectiveness • l1d_cache_mpki in MPKI	L1D_CACHE_REFILL / INST_RETIRED * 1000	• L1D_Cache_Effectiveness • MPKI
• l1d_tlb_miss_ratio in DTLB_Effectiveness • l1d_tlb_miss_ratio in Miss_Ratio	L1D_TLB_REFILL / L1D_TLB	• DTLB_Effectiveness • Miss_Ratio
• l1d_tlb_mpki in DTLB_Effectiveness • l1d_tlb_mpki in MPKI	L1D_TLB_REFILL / INST_RETIRED * 1000	• DTLB_Effectiveness • MPKI
• l1i_cache_miss_ratio in L1I_Cache_Effectiveness • l1i_cache_miss_ratio in Miss_Ratio	L1I_CACHE_REFILL / L1I_CACHE	• L1I_Cache_Effectiveness • Miss_Ratio
• l1i_cache_mpki in L1I_Cache_Effectiveness • l1i_cache_mpki in MPKI	L1I_CACHE_REFILL / INST_RETIRED * 1000	• L1I_Cache_Effectiveness • MPKI
• l1i_tlb_miss_ratio in ITLB_Effectiveness • l1i_tlb_miss_ratio in Miss_Ratio	L1I_TLB_REFILL / L1I_TLB	• ITLB_Effectiveness • Miss_Ratio
• l1i_tlb_mpki in ITLB_Effectiveness • l1i_tlb_mpki in MPKI	L1I_TLB_REFILL / INST_RETIRED * 1000	• ITLB_Effectiveness • MPKI
• l2_cache_miss_ratio in L2_Cache_Effectiveness • l2_cache_miss_ratio in Miss_Ratio	L2D_CACHE_REFILL / L2D_CACHE	• L2_Cache_Effectiveness • Miss_Ratio
• l2_cache_mpki in L2_Cache_Effectiveness • l2_cache_mpki in MPKI	L2D_CACHE_REFILL / INST_RETIRED * 1000	• L2_Cache_Effectiveness • MPKI
• l2_tlb_miss_ratio in ITLB_Effectiveness • l2_tlb_miss_ratio in DTLB_Effectiveness • l2_tlb_miss_ratio in Miss_Ratio	L2D_TLB_REFILL / L2D_TLB	• ITLB_Effectiveness • DTLB_Effectiveness • Miss_Ratio
• l2_tlb_mpki in ITLB_Effectiveness • l2_tlb_mpki in DTLB_Effectiveness • l2_tlb_mpki in MPKI	L2D_TLB_REFILL / INST_RETIRED * 1000	• ITLB_Effectiveness • DTLB_Effectiveness • MPKI
ll_cache_read_hit_ratio	(LL_CACHE_RD - LL_CACHE_MISS_RD) / LL_CACHE_RD	• LL_Cache_Effectiveness
• ll_cache_read_miss_ratio in LL_Cache_Effectiveness • ll_cache_read_miss_ratio in Miss_Ratio	LL_CACHE_MISS_RD / LL_CACHE_RD	• LL_Cache_Effectiveness • Miss_Ratio

Metric Name	Formula from Events	Metric Groups
• ll_cache_read_mpki in LL_Cache_Effectiveness • ll_cache_read_mpki in MPKI	LL_CACHE_MISS_RD / INST_RETIRED * 1000	• LL_Cache_Effectiveness • MPKI
load_percentage	LD_SPEC / INST_SPEC * 100	• Operation_Mix
retiring	100 * (OP_RETIRED / OP_SPEC * (1 - STALL_SLOT / (CPU_CYCLES * 8)))	• Topdown_L1
scalar_fp_percentage	VFP_SPEC / INST_SPEC * 100	• Operation_Mix
simd_percentage	ASE_SPEC / INST_SPEC * 100	• Operation_Mix
store_percentage	ST_SPEC / INST_SPEC * 100	• Operation_Mix

5.4 PMU events lookup table for Neoverse V2

All events are listed in event code order, with the related metrics, metric groups, and functional groups. Some events are not used in the Methodology, however, they are all listed for completeness.

Summary of Events:

- Total Possible Common events: 734
- Total implemented Common events: 155
  - Common : Architectural-defined events: 91
  - Common : Implementation-defiend events: 64
- Total Implemented Product ImpDef events: 0
- PMU Only Events : 0
- ETE Only Events : 0

Table 5-12: Events listed by Event Code, with related Metrics, Metric Groups, and Functional Groups

Code, Mnemonic	Metrics	Metric Groups	Functional Groups
0x0000, SW_INCR	-	-	• Retired
0x0001, L1I_CACHE_REFILL	• l1i_cache_mpki in L1I_Cache_Effectiveness • l1i_cache_mpki in MPKI • l1i_cache_miss_ratio in L1I_Cache_Effectiveness • l1i_cache_miss_ratio in Miss_Ratio	• L1I_Cache_Effectiveness • MPKI • Miss_Ratio	• L1I_Cache

Code, Mnemonic	Metrics	Metric Groups	Functional Groups
0x0002, L1I_TLB_REFILL	• l1i_tlb_mpki in ITLB_Effectiveness • l1i_tlb_mpki in MPKI • l1i_tlb_miss_ratio in ITLB_Effectiveness • l1i_tlb_miss_ratio in Miss_Ratio	• Miss_Ratio • MPKI • ITLB_Effectiveness	• TLB
0x0003, L1D_CACHE_REFILL	• l1d_cache_mpki in L1D_Cache_Effectiveness • l1d_cache_mpki in MPKI • l1d_cache_miss_ratio in L1D_Cache_Effectiveness • l1d_cache_miss_ratio in Miss_Ratio	• L1D_Cache_Effectiveness • MPKI • Miss_Ratio	• L1D_Cache
0x0004, L1D_CACHE	• l1d_cache_miss_ratio in L1D_Cache_Effectiveness • l1d_cache_miss_ratio in Miss_Ratio	• L1D_Cache_Effectiveness • Miss_Ratio	• L1D_Cache
0x0005, L1D_TLB_REFILL	• l1d_tlb_mpki in DTLB_Effectiveness • l1d_tlb_mpki in MPKI • l1d_tlb_miss_ratio in DTLB_Effectiveness • l1d_tlb_miss_ratio in Miss_Ratio	• DTLB_Effectiveness • MPKI • Miss_Ratio	• TLB

Code, Mnemonic	Metrics	Metric Groups	Functional Groups
0x0008, INST_RETIRED	• ipc • branch_mpki in Branch_Effectiveness • branch_mpki in MPKI • itlb_mpki in ITLB_Effectiveness • itlb_mpki in MPKI • l1i_tlb_mpki in ITLB_Effectiveness • l1i_tlb_mpki in MPKI • dtlb_mpki in DTLB_Effectiveness • dtlb_mpki in MPKI • l1d_tlb_mpki in DTLB_Effectiveness • l1d_tlb_mpki in MPKI • l2_tlb_mpki in ITLB_Effectiveness • l2_tlb_mpki in DTLB_Effectiveness • l2_tlb_mpki in MPKI • l1i_cache_mpki in L1I_Cache_Effectiveness • l1i_cache_mpki in MPKI • l1d_cache_mpki in L1D_Cache_Effectiveness • l1d_cache_mpki in MPKI • l2_cache_mpki in L2_Cache_Effectiveness • l2_cache_mpki in MPKI • ll_cache_read_mpki in LL_Cache_Effectiveness • ll_cache_read_mpki in MPKI	• LL_Cache_Effectiveness • MPKI • L1D_Cache_Effectiveness • L2_Cache_Effectiveness • Branch_Effectiveness • General • ITLB_Effectiveness • L1I_Cache_Effectiveness • DTLB_Effectiveness	• Retired
0x0009, EXC_TAKEN	-	-	• Exception
0x000A, EXC_RETURN	-	-	• Exception
0x000B, CID_WRITE_RETIRED	-	-	• Retired
0x0010, BR_MIS_PRED	• frontend_bound • backend_bound • bad_speculation	• Topdown_L1	• Spec_Operation

Code, Mnemonic	Metrics	Metric Groups	Functional Groups
0x0011, CPU_CYCLES	• frontend_stalled_cycles • backend_stalled_cycles • frontend_bound • backend_bound • retiring • bad_speculation • ipc	• Topdown_L1 • Cycle_Accounting • General	• General
0x0012, BR_PRED	-	-	• Spec_Operation
0x0013, MEM_ACCESS	-	-	• Memory
0x0014, L1I_CACHE	• l1i_cache_miss_ratio in L1I_Cache_Effectiveness • l1i_cache_miss_ratio in Miss_Ratio	• L1I_Cache_Effectiveness • Miss_Ratio	• L1I_Cache
0x0015, L1D_CACHE_WB	-	-	• L1D_Cache
0x0016, L2D_CACHE	• l2_cache_miss_ratio in L2_Cache_Effectiveness • l2_cache_miss_ratio in Miss_Ratio	• L2_Cache_Effectiveness • Miss_Ratio	• L2_Cache
0x0017, L2D_CACHE_REFILL	• l2_cache_mpki in L2_Cache_Effectiveness • l2_cache_mpki in MPKI • l2_cache_miss_ratio in L2_Cache_Effectiveness • l2_cache_miss_ratio in Miss_Ratio	• Miss_Ratio • MPKI • L2_Cache_Effectiveness	• L2_Cache
0x0018, L2D_CACHE_WB	-	-	• L2_Cache
0x0019, BUS_ACCESS	-	-	• Bus
0x001A, MEMORY_ERROR	-	-	• Memory
0x001B, INST_SPEC	• load_percentage • store_percentage • integer_dp_percentage • simd_percentage • scalar_fp_percentage • branch_percentage • crypto_percentage	• Operation_Mix	• Spec_Operation
0x001C, TTBR_WRITE_RETIRED	-	-	• Retired
0x001D, BUS_CYCLES	-	-	• Bus
0x001E, CHAIN	-	-	• Chain
0x0020, L2D_CACHE_ALLOCATE	-	-	• L2_Cache

Code, Mnemonic	Metrics	Metric Groups	Functional Groups
0x0021, BR_RETIRED	• branch_misprediction_ratio in Branch_Effectiveness • branch_misprediction_ratio in Miss_Ratio	• Branch_Effectiveness • Miss_Ratio	• Retired
0x0022, BR_MIS_PRED_RETIRED	• branch_mpki in Branch_Effectiveness • branch_mpki in MPKI • branch_misprediction_ratio in Branch_Effectiveness • branch_misprediction_ratio in Miss_Ratio	• Branch_Effectiveness • MPKI • Miss_Ratio	• Retired
0x0023, STALL_FRONTEND	• frontend_stalled_cycles	• Cycle_Accounting	• Stall
0x0024, STALL_BACKEND	• backend_stalled_cycles	• Cycle_Accounting	• Stall
0x0025, L1D_TLB	• dtlb_walk_ratio in DTLB_Effectiveness • dtlb_walk_ratio in Miss_Ratio • l1d_tlb_miss_ratio in DTLB_Effectiveness • l1d_tlb_miss_ratio in Miss_Ratio	• DTLB_Effectiveness • Miss_Ratio	• TLB
0x0026, L1I_TLB	• itlb_walk_ratio in ITLB_Effectiveness • itlb_walk_ratio in Miss_Ratio • l1i_tlb_miss_ratio in ITLB_Effectiveness • l1i_tlb_miss_ratio in Miss_Ratio	• ITLB_Effectiveness • Miss_Ratio	• TLB
0x0029, L3D_CACHE_ALLOCATE	-	-	• L3_Cache
0x002A, L3D_CACHE_REFILL	-	-	• L3_Cache
0x002B, L3D_CACHE	-	-	• L3_Cache
0x002D, L2D_TLB_REFILL	• l2_tlb_mpki in ITLB_Effectiveness • l2_tlb_mpki in DTLB_Effectiveness • l2_tlb_mpki in MPKI • l2_tlb_miss_ratio in ITLB_Effectiveness • l2_tlb_miss_ratio in DTLB_Effectiveness • l2_tlb_miss_ratio in Miss_Ratio	• MPKI • ITLB_Effectiveness • DTLB_Effectiveness • Miss_Ratio	• TLB

Code, Mnemonic	Metrics	Metric Groups	Functional Groups
0x002F, L2D_TLB	• l2_tlb_miss_ratio in ITLB_Effectiveness • l2_tlb_miss_ratio in DTLB_Effectiveness • l2_tlb_miss_ratio in Miss_Ratio	• Miss_Ratio • DTLB_Effectiveness • ITLB_Effectiveness	• TLB
0x0031, REMOTE_ACCESS	-	-	• Memory
0x0034, DTLB_WALK	• dtlb_mpki in DTLB_Effectiveness • dtlb_mpki in MPKI • dtlb_walk_ratio in DTLB_Effectiveness • dtlb_walk_ratio in Miss_Ratio	• DTLB_Effectiveness • MPKI • Miss_Ratio	• TLB
0x0035, ITLB_WALK	• itlb_mpki in ITLB_Effectiveness • itlb_mpki in MPKI • itlb_walk_ratio in ITLB_Effectiveness • itlb_walk_ratio in Miss_Ratio	• Miss_Ratio • MPKI • ITLB_Effectiveness	• TLB
0x0036, LL_CACHE_RD	• ll_cache_read_miss_ratio in LL_Cache_Effectiveness • ll_cache_read_miss_ratio in Miss_Ratio • ll_cache_read_hit_ratio	• LL_Cache_Effectiveness • Miss_Ratio	• LL_Cache
0x0037, LL_CACHE_MISS_RD	• ll_cache_read_mpki in LL_Cache_Effectiveness • ll_cache_read_mpki in MPKI • ll_cache_read_miss_ratio in LL_Cache_Effectiveness • ll_cache_read_miss_ratio in Miss_Ratio • ll_cache_read_hit_ratio	• LL_Cache_Effectiveness • MPKI • Miss_Ratio	• LL_Cache
0x0039, L1D_CACHE_LMISS_RD	-	-	• L1D_Cache
0x003A, OP_RETIRED	• retiring • bad_speculation	• Topdown_L1	• Retired
0x003B, OP_SPEC	• retiring • bad_speculation	• Topdown_L1	• Spec_Operation
0x003C, STALL	-	-	• Stall
0x003D, STALL_SLOT_BACKEND	• backend_bound	• Topdown_L1	• Stall
0x003E, STALL_SLOT_FRONTEND	• frontend_bound	• Topdown_L1	• Stall
0x003F, STALL_SLOT	• retiring • bad_speculation	• Topdown_L1	• Stall

Code, Mnemonic	Metrics	Metric Groups	Functional Groups
0x0040, L1D_CACHE_RD	-	-	• L1D_Cache
0x0041, L1D_CACHE_WR	-	-	• L1D_Cache
0x0042, L1D_CACHE_REFILL_RD	-	-	• L1D_Cache
0x0043, L1D_CACHE_REFILL_WR	-	-	• L1D_Cache
0x0044, L1D_CACHE_REFILL_INNER	-	-	• L1D_Cache
0x0045, L1D_CACHE_REFILL_OUTER	-	-	• L1D_Cache
0x0046, L1D_CACHE_WB_VICTIM	-	-	• L1D_Cache
0x0047, L1D_CACHE_WB_CLEAN	-	-	• L1D_Cache
0x0048, L1D_CACHE_INVAL	-	-	• L1D_Cache
0x004C, L1D_TLB_REFILL_RD	-	-	• TLB
0x004D, L1D_TLB_REFILL_WR	-	-	• TLB
0x004E, L1D_TLB_RD	-	-	• TLB
0x004F, L1D_TLB_WR	-	-	• TLB
0x0050, L2D_CACHE_RD	-	-	• L2_Cache
0x0051, L2D_CACHE_WR	-	-	• L2_Cache
0x0052, L2D_CACHE_REFILL_RD	-	-	• L2_Cache
0x0053, L2D_CACHE_REFILL_WR	-	-	• L2_Cache
0x0056, L2D_CACHE_WB_VICTIM	-	-	• L2_Cache
0x0057, L2D_CACHE_WB_CLEAN	-	-	• L2_Cache
0x0058, L2D_CACHE_INVAL	-	-	• L2_Cache
0x005C, L2D_TLB_REFILL_RD	-	-	• TLB
0x005D, L2D_TLB_REFILL_WR	-	-	• TLB
0x005E, L2D_TLB_RD	-	-	• TLB
0x005F, L2D_TLB_WR	-	-	• TLB
0x0060, BUS_ACCESS_RD	-	-	• Bus
0x0061, BUS_ACCESS_WR	-	-	• Bus
0x0066, MEM_ACCESS_RD	-	-	• Memory
0x0067, MEM_ACCESS_WR	-	-	• Memory
0x0068, UNALIGNED_LD_SPEC	-	-	• Spec_Operation
0x0069, UNALIGNED_ST_SPEC	-	-	• Spec_Operation
0x006A, UNALIGNED_LDST_SPEC	-	-	• Spec_Operation
0x006C, LDREX_SPEC	-	-	• Spec_Operation
0x006D, STREX_PASS_SPEC	-	-	• Spec_Operation

Code, Mnemonic	Metrics	Metric Groups	Functional Groups
0x006E, STREX_FAIL_SPEC	-	-	• Spec_Operation
0x006F, STREX_SPEC	-	-	• Spec_Operation
0x0070, LD_SPEC	• load_percentage	• Operation_Mix	• Spec_Operation
0x0071, ST_SPEC	• store_percentage	• Operation_Mix	• Spec_Operation
0x0073, DP_SPEC	• integer_dp_percentage	• Operation_Mix	• Spec_Operation
0x0074, ASE_SPEC	• simd_percentage	• Operation_Mix	• Spec_Operation
0x0075, VFP_SPEC	• scalar_fp_percentage	• Operation_Mix	• Spec_Operation
0x0076, PC_WRITE_SPEC	-	-	• Spec_Operation
0x0077, CRYPTO_SPEC	• crypto_percentage	• Operation_Mix	• Spec_Operation
0x0078, BR_IMMED_SPEC	• branch_percentage	• Operation_Mix	• Spec_Operation
0x0079, BR_RETURN_SPEC	-	-	• Spec_Operation
0x007A, BR_INDIRECT_SPEC	• branch_percentage	• Operation_Mix	• Spec_Operation
0x007C, ISB_SPEC	-	-	• Spec_Operation
0x007D, DSB_SPEC	-	-	• Spec_Operation
0x007E, DMB_SPEC	-	-	• Spec_Operation
0x0081, EXC_UNDEF	-	-	• Exception
0x0082, EXC_SVC	-	-	• Exception
0x0083, EXC_PABORT	-	-	• Exception
0x0084, EXC_DABORT	-	-	• Exception
0x0086, EXC_IRQ	-	-	• Exception
0x0087, EXC_FIQ	-	-	• Exception
0x0088, EXC_SMC	-	-	• Exception
0x008A, EXC_HVC	-	-	• Exception
0x008B, EXC_TRAP_PABORT	-	-	• Exception
0x008C, EXC_TRAP_DABORT	-	-	• Exception
0x008D, EXC_TRAP_OTHER	-	-	• Exception
0x008E, EXC_TRAP_IRQ	-	-	• Exception
0x008F, EXC_TRAP_FIQ	-	-	• Exception
0x0090, RC_LD_SPEC	-	-	• Spec_Operation
0x0091, RC_ST_SPEC	-	-	• Spec_Operation
0x00A0, L3D_CACHE_RD	-	-	• L3_Cache
0x4000, SAMPLE_POP	-	-	• SPE
0x4001, SAMPLE_FEED	-	-	• SPE
0x4002, SAMPLE_FILTRATE	-	-	• SPE
0x4003, SAMPLE_COLLISION	-	-	• SPE
0x4004, CNT_CYCLES	-	-	• General
0x4005, STALL_BACKEND_MEM	-	-	• Stall
0x4006, L1I_CACHE_LMISS	-	-	• L1I_Cache
0x4009, L2D_CACHE_LMISS_RD	-	-	• L2_Cache

Code, Mnemonic	Metrics	Metric Groups	Functional Groups
0x400B, L3D_CACHE_LMISS_RD	-	-	• L3_Cache
0x400C, TRB_WRAP	-	-	• TRACE
0x4010, TRCEXTOUT0	-	-	• TRACE
0x4011, TRCEXTOUT1	-	-	• TRACE
0x4012, TRCEXTOUT2	-	-	• TRACE
0x4013, TRCEXTOUT3	-	-	• TRACE
0x4018, CTI_TRIGOUT4	-	-	• TRACE
0x4019, CTI_TRIGOUT5	-	-	• TRACE
0x401A, CTI_TRIGOUT6	-	-	• TRACE
0x401B, CTI_TRIGOUT7	-	-	• TRACE
0x4020, LDST_ALIGN_LAT	-	-	• Memory
0x4021, LD_ALIGN_LAT	-	-	• Memory
0x4022, ST_ALIGN_LAT	-	-	• Memory
0x4024, MEM_ACCESS_CHECKED	-	-	• Memory
0x4025, MEM_ACCESS_CHECKED_RD	-	-	• Memory
0x4026, MEM_ACCESS_CHECKED_WR	-	-	• Memory
0x8005, ASE_INST_SPEC	-	-	• Spec_Operation
0x8006, SVE_INST_SPEC	-	-	• SVE
0x8014, FP_HP_SPEC	-	-	• FP_Operation
0x8018, FP_SP_SPEC	-	-	• FP_Operation
0x801C, FP_DP_SPEC	-	-	• FP_Operation
0x8074, SVE_PRED_SPEC	-	-	• SVE
0x8075, SVE_PRED_EMPTY_SPEC	-	-	• SVE
0x8076, SVE_PRED_FULL_SPEC	-	-	• SVE
0x8077, SVE_PRED_PARTIAL_SPEC	-	-	• SVE
0x8079, SVE_PRED_NOT_FULL_SPEC	-	-	• SVE
0x80BC, SVE_LDFF_SPEC	-	-	• SVE
0x80BD, SVE_LDFF_FAULT_SPEC	-	-	• SVE
0x80C0, FP_SCALE_OPS_SPEC	-	-	• FP_Operation
0x80C1, FP_FIXED_OPS_SPEC	-	-	• FP_Operation
0x80E3, ASE_SVE_INT8_SPEC	-	-	• SVE
0x80E7, ASE_SVE_INT16_SPEC	-	-	• SVE
0x80EB, ASE_SVE_INT32_SPEC	-	-	• SVE
0x80EF, ASE_SVE_INT64_SPEC	-	-	• SVE

6. Metrics by metric group in Neoverse V2

Metrics are measured using different combinations of PMU events. They are organized into groups that can be analyzed together for a use case. To calculate the metrics, two or more PMU counters are programmed with the events listed for the metric. The counters are read at the same time to determine the metric value.

Summary:

- Total metrics: 35

Metrics for Neoverse V2 are grouped into the following metric groups:

- Topdown_L1, Topdown Level 1 (4 metrics)
- Cycle_Accounting, Cycle Accounting (2 metrics)
- General, General (1 metrics)
- MPKI, Misses Per Kilo Instructions (10 metrics)
- Miss_Ratio, Miss Ratio (10 metrics)
- Branch_Effectiveness, Branch Effectiveness (2 metrics)
- ITLB_Effectiveness, Instruction TLB Effectiveness (6 metrics)
- DTLB_Effectiveness, Data TLB Effectiveness (6 metrics)
- L1I_Cache_Effectiveness, L1 Instruction Cache Effectiveness (2 metrics)
- L1D_Cache_Effectiveness, L1 Data Cache Effectiveness (2 metrics)
- L2_Cache_Effectiveness, L2 Unified Cache Effectiveness (2 metrics)
- LL_Cache_Effectiveness, Last Level Cache Effectiveness (3 metrics)
- Operation_Mix, Speculative Operation Mix (7 metrics)

6.1 Topdown_L1 metrics for Neoverse V2

Topdown Level 1. This metric group contains the first set of metrics to begin topdown analysis of application performance, which provide the percentage distribution of processor pipeline utilization.

Summary of metrics in Topdown_L1:

- Total metrics: 4

Table 6-1: Topdown_L1 metrics summary

Metric	Name	Description
backend_bound	Backend Bound	This metric is the percentage of total slots that were stalled due to resource constraints in the...
bad_speculation	Bad Speculation	This metric is the percentage of total slots that executed operations and didn't retire due to a...
frontend_bound	Frontend Bound	This metric is the percentage of total slots that were stalled due to resource constraints in the...
retiring	Retiring	This metric is the percentage of total slots that retired operations, which indicates cycles that...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

backend_bound, Backend Bound, metric

This metric is the percentage of total slots that were stalled due to resource constraints in the backend of the processor.

Units

This unit is expressed as percent of slots.

Formula

100 * (STALL_SLOT_BACKEND / (CPU_CYCLES * 8) - BR_MIS_PRED * 3 / CPU_CYCLES)

Related telemetry artifacts

Events

BR_MIS_PRED
CPU_CYCLES
STALL_SLOT_BACKEND

Metric group

Topdown_L1

Methodology

Stage 1

bad_speculation, Bad Speculation, metric

This metric is the percentage of total slots that executed operations and didn't retire due to a pipeline flush. This indicates cycles that were utilized but inefficiently.

Units

This unit is expressed as percent of slots.

Formula

100 * ((1 - OP_RETIRED / OP_SPEC) * (1 - STALL_SLOT / (CPU_CYCLES * 8)) + BR_MIS_PRED * 4 / CPU_CYCLES)

Related telemetry artifacts

Events

BR_MIS_PRED
CPU_CYCLES
OP_RETIRED
OP_SPEC
STALL_SLOT

Metric group

Topdown_L1

Methodology

Stage 1

frontend_bound, Frontend Bound, metric

This metric is the percentage of total slots that were stalled due to resource constraints in the frontend of the processor.

Units

This unit is expressed as percent of slots.

Formula

100 * (STALL_SLOT_FRONTEND / (CPU_CYCLES * 8) - BR_MIS_PRED / CPU_CYCLES)

Related telemetry artifacts

Events

BR_MIS_PRED
CPU_CYCLES
STALL_SLOT_FRONTEND

Metric group

Topdown_L1

Methodology

Stage 1

retiring, Retiring, metric

This metric is the percentage of total slots that retired operations, which indicates cycles that were utilized efficiently.

Units

This unit is expressed as percent of slots.

Formula

100 * (OP_RETIRED / OP_SPEC * (1 - STALL_SLOT / (CPU_CYCLES * 8)))

Related telemetry artifacts

Events

CPU_CYCLES
OP_RETIRED
OP_SPEC
STALL_SLOT

Metric group

Topdown_L1

Methodology

Stage 1

6.2 Cycle_Accounting metrics for Neoverse V2

Cycle Accounting. This metric group contains a set of metrics that measure the percentage of processor cycles stalled in either frontend or backend of the processor.

Summary of metrics in Cycle_Accounting:

- Total metrics: 2

Table 6-2: Cycle_Accounting metrics summary

Metric	Name	Description
backend_stalled_cycles	Backend Stalled Cycles	This metric is the percentage of cycles that were stalled due to resource constraints in the...
frontend_stalled_cycles	Frontend Stalled Cycles	This metric is the percentage of cycles that were stalled due to resource constraints in the...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

backend_stalled_cycles, Backend Stalled Cycles, metric

This metric is the percentage of cycles that were stalled due to resource constraints in the backend unit of the processor.

Units

This unit is expressed as percent of cycles.

Formula

STALL_BACKEND / CPU_CYCLES * 100

Related telemetry artifacts

Events

CPU_CYCLES
STALL_BACKEND

Metric group

Cycle_Accounting

Methodology

Stage 2

frontend_stalled_cycles, Frontend Stalled Cycles, metric

This metric is the percentage of cycles that were stalled due to resource constraints in the frontend unit of the processor.

Units

This unit is expressed as percent of cycles.

Formula

STALL_FRONTEND / CPU_CYCLES * 100

Related telemetry artifacts

Events

CPU_CYCLES
STALL_FRONTEND

Metric group

Cycle_Accounting

Methodology

Stage 2

6.3 General metrics for Neoverse V2

General. This metric group contains general CPU metrics for performance analysis.

Summary of metrics in General:

- Total metrics: 1

Table 6-3: General metrics summary

Metric	Name	Description
ipc	Instructions Per Cycle	This metric measures the number of instructions retired per cycle.

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

ipc, Instructions Per Cycle, metric

This metric measures the number of instructions retired per cycle.

Units

This unit is expressed as per cycle.

Formula

INST_RETIRED / CPU_CYCLES

Related telemetry artifacts

Events

CPU_CYCLES
INST_RETIRED

Metric group

General

Methodology

Stage 2

6.4 MPKI metrics for Neoverse V2

Misses Per Kilo Instructions. This metric group contains metrics for different CPU resources that can be measured as misses per kilo instructions.

Summary of metrics in MPKI:

- Total metrics: 10

Table 6-4: MPKI metrics summary

Metric	Name	Description
branch_mpki	Branch MPKI	This metric measures the number of branch mispredictions per thousand instructions executed.
dtlb_mpki	DTLB MPKI	This metric measures the number of data TLB Walks per thousand instructions executed.
itlb_mpki	ITLB MPKI	This metric measures the number of instruction TLB Walks per thousand instructions executed.
l1d_cache_mpki	L1D Cache MPKI	This metric measures the number of level 1 data cache accesses missed per thousand instructions...
l1d_tlb_mpki	L1 Data TLB MPKI	This metric measures the number of level 1 data TLB accesses missed per thousand instructions...
l1i_cache_mpki	L1I Cache MPKI	This metric measures the number of level 1 instruction cache accesses missed per thousand...
l1i_tlb_mpki	L1 Instruction TLB MPKI	This metric measures the number of level 1 instruction TLB accesses missed per thousand...
l2_cache_mpki	L2 Cache MPKI	This metric measures the number of level 2 unified cache accesses missed per thousand...
l2_tlb_mpki	L2 Unified TLB MPKI	This metric measures the number of level 2 unified TLB accesses missed per thousand instructions...
ll_cache_read_mpki	LL Cache Read MPKI	This metric measures the number of last level cache read accesses missed per thousand...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

branch_mpki, Branch MPKI, metric

This metric measures the number of branch mispredictions per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

BR_MIS_PRED_RETIRED / INST_RETIRED * 1000

Related telemetry artifacts

Events

BR_MIS_PRED_RETIRED
INST_RETIRED

Metric group

MPKI

Other metric group: Branch_Effectiveness

Methodology

Stage 2

dtlb_mpki, DTLB MPKI, metric

This metric measures the number of data TLB Walks per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

DTLB_WALK / INST_RETIRED * 1000

Related telemetry artifacts

Events

DTLB_WALK

INST_RETIRED

Metric group

MPKI

Other metric group: DTLB_Effectiveness

Methodology

Stage 2

itlb_mpki, ITLB MPKI, metric

This metric measures the number of instruction TLB Walks per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

ITLB_WALK / INST_RETIRED * 1000

Related telemetry artifacts

Events

INST_RETIRED

ITLB_WALK

Metric group

MPKI

Other metric group: ITLB_Effectiveness

Methodology

Stage 2

l1d_cache_mpki, L1D Cache MPKI, metric

This metric measures the number of level 1 data cache accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L1D_CACHE_REFILL / INST_RETIRED * 1000

Related telemetry artifacts

Events

INST_RETIRED

L1D_CACHE_REFILL

Metric group

MPKI

Other metric group: L1D_Cache_Effectiveness

Methodology

Stage 2

l1d_tlb_mpki, L1 Data TLB MPKI, metric

This metric measures the number of level 1 data TLB accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L1D_TLB_REFILL / INST_RETIRED * 1000

Related telemetry artifacts

Events

INST_RETIRED

L1D_TLB_REFILL

Metric group

MPKI

Other metric group: DTLB_Effectiveness

Methodology

Stage 2

l1i_cache_mpki, L1I Cache MPKI, metric

This metric measures the number of level 1 instruction cache accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L1I_CACHE_REFILL / INST_RETIRED * 1000

Related telemetry artifacts

Events

INST_RETIRED

L1I_CACHE_REFILL

Metric group

MPKI

Other metric group: L1I_Cache_Effectiveness

Methodology

Stage 2

l1i_tlb_mpki, L1 Instruction TLB MPKI, metric

This metric measures the number of level 1 instruction TLB accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L1I_TLB_REFILL / INST_RETIRED * 1000

Related telemetry artifacts

Events

INST_RETIRED

L1I_TLB_REFILL

Metric group

MPKI

Other metric group: ITLB_Effectiveness

Methodology

Stage 2

l2_cache_mpki, L2 Cache MPKI, metric

This metric measures the number of level 2 unified cache accesses missed per thousand instructions executed. Note that cache accesses in this cache are either data memory access or instruction fetch as this is a unified cache.

Units

This unit is expressed as mpki.

Formula

L2D_CACHE_REFILL / INST_RETIRED * 1000

Related telemetry artifacts

Events

INST_RETIRED

L2D_CACHE_REFILL

Metric group

MPKI

Other metric group: L2_Cache_Effectiveness

Methodology

Stage 2

l2_tlb_mpki, L2 Unified TLB MPKI, metric

This metric measures the number of level 2 unified TLB accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L2D_TLB_REFILL / INST_RETIRED * 1000

Related telemetry artifacts

Events

INST_RETIRED

L2D_TLB_REFILL

Metric group

MPKI

Other metric group: ITLB_Effectiveness

Other metric group: DTLB_Effectiveness

Methodology

Stage 2

ll_cache_read_mpki, LL Cache Read MPKI, metric

This metric measures the number of last level cache read accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

LL_CACHE_MISS_RD / INST_RETIRED * 1000

Related telemetry artifacts

Events

INST_RETIRED

LL_CACHE_MISS_RD

Metric group

MPKI

Other metric group: LL_Cache_Effectiveness

Methodology

Stage 2

6.5 Miss_Ratio metrics for Neoverse V2

Miss Ratio. This metric group contains metrics to measure miss ratios of different processor resources.

Summary of metrics in Miss_Ratio:

- Total metrics: 10

Table 6-5: Miss_Ratio metrics summary

Metric	Name	Description
branch_misprediction_ratio	Branch Misprediction Ratio	This metric measures the ratio of branches mispredicted to the total number of branches...
dtlb_walk_ratio	DTLB Walk Ratio	This metric measures the ratio of data TLB Walks to the total number of data TLB accesses. This...
itlb_walk_ratio	ITLB Walk Ratio	This metric measures the ratio of instruction TLB Walks to the total number of instruction TLB...
l1d_cache_miss_ratio	L1D Cache Miss Ratio	This metric measures the ratio of level 1 data cache accesses missed to the total number of level...
l1d_tlb_miss_ratio	L1 Data TLB Miss Ratio	This metric measures the ratio of level 1 data TLB accesses missed to the total number of level 1...
l1i_cache_miss_ratio	L1I Cache Miss Ratio	This metric measures the ratio of level 1 instruction cache accesses missed to the total number...
l1i_tlb_miss_ratio	L1 Instruction TLB Miss Ratio	This metric measures the ratio of level 1 instruction TLB accesses missed to the total number of...
l2_cache_miss_ratio	L2 Cache Miss Ratio	This metric measures the ratio of level 2 cache accesses missed to the total number of level 2...
l2_tlb_miss_ratio	L2 Unified TLB Miss Ratio	This metric measures the ratio of level 2 unified TLB accesses missed to the total number of...
ll_cache_read_miss_ratio	LL Cache Read Miss Ratio	This metric measures the ratio of last level cache read accesses missed to the total number of...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

branch_misprediction_ratio, Branch Misprediction Ratio, metric

This metric measures the ratio of branches mispredicted to the total number of branches architecturally executed. This gives an indication of the effectiveness of the branch prediction unit.

Units

This unit is expressed as per branch.

Formula

BR_MIS_PRED_RETIRED / BR_RETIRED

Related telemetry artifacts

Events

BR_MIS_PRED_RETIRED
BR_RETIRED

Metric group

Miss_Ratio

Other metric group: Branch_Effectiveness

Methodology

Stage 2

dtlb_walk_ratio, DTLB Walk Ratio, metric

This metric measures the ratio of data TLB Walks to the total number of data TLB accesses. This gives an indication of the effectiveness of the data TLB accesses.

Units

This unit is expressed as per tlb access.

Formula

DTLB_WALK / L1D_TLB

Related telemetry artifacts

Events

DTLB_WALK
L1D_TLB

Metric group

Miss_Ratio

Other metric group: DTLB_Effectiveness

Methodology

Stage 2

itlb_walk_ratio, ITLB Walk Ratio, metric

This metric measures the ratio of instruction TLB Walks to the total number of instruction TLB accesses. This gives an indication of the effectiveness of the instruction TLB accesses.

Units

This unit is expressed as per tlb access.

Formula

ITLB_WALK / L1I_TLB

Related telemetry artifacts

Events

ITLB_WALK
L1I_TLB

Metric group

Miss_Ratio

Other metric group: ITLB_Effectiveness

Methodology

Stage 2

l1d_cache_miss_ratio, L1D Cache Miss Ratio, metric

This metric measures the ratio of level 1 data cache accesses missed to the total number of level 1 data cache accesses. This gives an indication of the effectiveness of the level 1 data cache.

Units

This unit is expressed as per cache access.

Formula

L1D_CACHE_REFILL / L1D_CACHE

Related telemetry artifacts

Events

L1D_CACHE

L1D_CACHE_REFILL

Metric group

Miss_Ratio

Other metric group: L1D_Cache_Effectiveness

Methodology

Stage 2

l1d_tlb_miss_ratio, L1 Data TLB Miss Ratio, metric

This metric measures the ratio of level 1 data TLB accesses missed to the total number of level 1 data TLB accesses. This gives an indication of the effectiveness of the level 1 data TLB.

Units

This unit is expressed as per tlb access.

Formula

L1D_TLB_REFILL / L1D_TLB

Related telemetry artifacts

Events

L1D_TLB

L1D_TLB_REFILL

Metric group

Miss_Ratio

Other metric group: DTLB_Effectiveness

Methodology

Stage 2

l1i_cache_miss_ratio, L1I Cache Miss Ratio, metric

This metric measures the ratio of level 1 instruction cache accesses missed to the total number of level 1 instruction cache accesses. This gives an indication of the effectiveness of the level 1 instruction cache.

Units

This unit is expressed as per cache access.

Formula

L1I_CACHE_REFILL / L1I_CACHE

Related telemetry artifacts

Events

L1I_CACHE
L1I_CACHE_REFILL

Metric group

Miss_Ratio

Other metric group: L1I_Cache_Effectiveness

Methodology

Stage 2

l1i_tlb_miss_ratio, L1 Instruction TLB Miss Ratio, metric

This metric measures the ratio of level 1 instruction TLB accesses missed to the total number of level 1 instruction TLB accesses. This gives an indication of the effectiveness of the level 1 instruction TLB.

Units

This unit is expressed as per tlb access.

Formula

L1I_TLB_REFILL / L1I_TLB

Related telemetry artifacts

Events

L1I_TLB
L1I_TLB_REFILL

Metric group

Miss_Ratio

Other metric group: ITLB_Effectiveness

Methodology

Stage 2

l2_cache_miss_ratio, L2 Cache Miss Ratio, metric

This metric measures the ratio of level 2 cache accesses missed to the total number of level 2 cache accesses. This gives an indication of the effectiveness of the level 2 cache, which is a unified

cache that stores both data and instruction. Note that cache accesses in this cache are either data memory access or instruction fetch as this is a unified cache.

Units

This unit is expressed as per cache access.

Formula

L2D_CACHE_REFILL / L2D_CACHE

Related telemetry artifacts

Events

L2D_CACHE
L2D_CACHE_REFILL

Metric group

Miss_Ratio
Other metric group: L2_Cache_Effectiveness

Methodology

Stage 2

l2_tlb_miss_ratio, L2 Unified TLB Miss Ratio, metric

This metric measures the ratio of level 2 unified TLB accesses missed to the total number of level 2 unified TLB accesses. This gives an indication of the effectiveness of the level 2 TLB.

Units

This unit is expressed as per tlb access.

Formula

L2D_TLB_REFILL / L2D_TLB

Related telemetry artifacts

Events

L2D_TLB
L2D_TLB_REFILL

Metric group

Miss_Ratio
Other metric group: ITLB_Effectiveness
Other metric group: DTLB_Effectiveness

Methodology

Stage 2

ll_cache_read_miss_ratio, LL Cache Read Miss Ratio, metric

This metric measures the ratio of last level cache read accesses missed to the total number of last level cache accesses. This gives an indication of the effectiveness of the last level cache for read traffic. Note that cache accesses in this cache are either data memory access or instruction fetch as this is a system level cache.

Units

This unit is expressed as per cache access.

Formula

LL_CACHE_MISS_RD / LL_CACHE_RD

Related telemetry artifacts

Events

LL_CACHE_MISS_RD
LL_CACHE_RD

Metric group

Miss_Ratio

Other metric group: LL_Cache_Effectiveness

Methodology

Stage 2

6.6 Branch_Effectiveness metrics for Neoverse V2

Branch Effectiveness. This metric group contains metrics to evaluate the effectiveness of branch instruction execution on this processor.

Summary of metrics in Branch_Effectiveness:

- Total metrics: 2

Table 6-6: Branch_Effectiveness metrics summary

Metric	Name	Description
branch_misprediction_ratio	Branch Misprediction Ratio	This metric measures the ratio of branches mispredicted to the total number of branches...
branch_mpki	Branch MPKI	This metric measures the number of branch mispredictions per thousand instructions executed.

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

branch_misprediction_ratio**, Branch Misprediction Ratio, metric

This metric measures the ratio of branches mispredicted to the total number of branches architecturally executed. This gives an indication of the effectiveness of the branch prediction unit.

Units

This unit is expressed as per branch.

Formula

BR_MIS_PRED_RETIRED / BR_RETIRED

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

BR_MIS_PRED_RETIRED
BR_RETIRED

Metric group

Branch_Effectiveness
Other metric group: Miss_Ratio

Methodology

Stage 2

branch_mpki**, Branch MPKI, metric

This metric measures the number of branch mispredictions per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

BR_MIS_PRED_RETIRED / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

BR_MIS_PRED_RETIRED
INST_RETIRED

Metric group

Branch_Effectiveness
Other metric group: MPKI

Methodology

Stage 2

6.7 ITLB_Effectiveness metrics for Neoverse V2

Instruction TLB Effectiveness. This metric group contains metrics to evaluate the effectiveness of instruction TLB on this processor.

Summary of metrics in ITLB_Effectiveness:

- Total metrics: 6

Table 6-7: ITLB_Effectiveness metrics summary

Metric	Name	Description
itlb_mpki	ITLB MPKI	This metric measures the number of instruction TLB Walks per thousand instructions executed.

Metric	Name	Description
itlb_walk_ratio	ITLB Walk Ratio	This metric measures the ratio of instruction TLB Walks to the total number of instruction TLB...
l1i_tlb_miss_ratio	L1 Instruction TLB Miss Ratio	This metric measures the ratio of level 1 instruction TLB accesses missed to the total number of...
l1i_tlb_mpki	L1 Instruction TLB MPKI	This metric measures the number of level 1 instruction TLB accesses missed per thousand...
l2_tlb_miss_ratio	L2 Unified TLB Miss Ratio	This metric measures the ratio of level 2 unified TLB accesses missed to the total number of...
l2_tlb_mpki	L2 Unified TLB MPKI	This metric measures the number of level 2 unified TLB accesses missed per thousand instructions...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

itlb_mpki**, ITLB MPKI, metric

This metric measures the number of instruction TLB Walks per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

ITLB_WALK / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

INST_RETIRED
ITLB_WALK

Metric group

ITLB_Effectiveness
Other metric group: MPKI

Methodology

Stage 2

itlb_walk_ratio**, ITLB Walk Ratio, metric

This metric measures the ratio of instruction TLB Walks to the total number of instruction TLB accesses. This gives an indication of the effectiveness of the instruction TLB accesses.

Units

This unit is expressed as per tlb access.

Formula

ITLB_WALK / L1I_TLB

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

ITLB_WALK
L1I_TLB

Metric group

ITLB_Effectiveness
Other metric group: Miss_Ratio

Methodology

Stage 2

l1i_tlb_miss_ratio**, L1 Instruction TLB Miss Ratio, metric

This metric measures the ratio of level 1 instruction TLB accesses missed to the total number of level 1 instruction TLB accesses. This gives an indication of the effectiveness of the level 1 instruction TLB.

Units

This unit is expressed as per tlb access.

Formula

L1I_TLB_REFILL / L1I_TLB

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

L1I_TLB
L1I_TLB_REFILL

Metric group

ITLB_Effectiveness
Other metric group: Miss_Ratio

Methodology

Stage 2

l1i_tlb_mpki**, L1 Instruction TLB MPKI, metric

This metric measures the number of level 1 instruction TLB accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L1I_TLB_REFILL / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

INST_RETIRED

L1I_TLB_REFILL

Metric group

ITLB_Effectiveness

Other metric group: MPKI

Methodology

Stage 2

l2_tlb_miss_ratio**, L2 Unified TLB Miss Ratio, metric

This metric measures the ratio of level 2 unified TLB accesses missed to the total number of level 2 unified TLB accesses. This gives an indication of the effectiveness of the level 2 TLB.

Units

This unit is expressed as per tlb access.

Formula

L2D_TLB_REFILL / L2D_TLB

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

L2D_TLB

L2D_TLB_REFILL

Metric group

ITLB_Effectiveness

Other metric group: DTLB_Effectiveness

Other metric group: Miss_Ratio

Methodology

Stage 2

l2_tlb_mpki**, L2 Unified TLB MPKI, metric

This metric measures the number of level 2 unified TLB accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L2D_TLB_REFILL / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

INST_RETIRED
L2D_TLB_REFILL

Metric group

ITLB_Effectiveness
Other metric group: DTLB_Effectiveness
Other metric group: MPKI

Methodology

Stage 2

6.8 DTLB_Effectiveness metrics for Neoverse V2

Data TLB Effectiveness. This metric group contains metrics to evaluate the effectiveness of data TLB on this processor.

Summary of metrics in DTLB_Effectiveness:

- Total metrics: 6

Table 6-8: DTLB_Effectiveness metrics summary

Metric	Name	Description
dtlb_mpki	DTLB MPKI	This metric measures the number of data TLB Walks per thousand instructions executed.
dtlb_walk_ratio	DTLB Walk Ratio	This metric measures the ratio of data TLB Walks to the total number of data TLB accesses. This...
l1d_tlb_miss_ratio	L1 Data TLB Miss Ratio	This metric measures the ratio of level 1 data TLB accesses missed to the total number of level 1...
l1d_tlb_mpki	L1 Data TLB MPKI	This metric measures the number of level 1 data TLB accesses missed per thousand instructions...
l2_tlb_miss_ratio	L2 Unified TLB Miss Ratio	This metric measures the ratio of level 2 unified TLB accesses missed to the total number of...
l2_tlb_mpki	L2 Unified TLB MPKI	This metric measures the number of level 2 unified TLB accesses missed per thousand instructions...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

dtlb_mpki**, DTLB MPKI, metric

This metric measures the number of data TLB Walks per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

DTLB_WALK / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

DTLB_WALK

INST_RETIRED

Metric group

DTLB_Effectiveness

Other metric group: MPKI

Methodology

Stage 2

dtlb_walk_ratio**, DTLB Walk Ratio, metric

This metric measures the ratio of data TLB Walks to the total number of data TLB accesses. This gives an indication of the effectiveness of the data TLB accesses.

Units

This unit is expressed as per tlb access.

Formula

DTLB_WALK / L1D_TLB

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

DTLB_WALK

L1D_TLB

Metric group

DTLB_Effectiveness

Other metric group: Miss_Ratio

Methodology

Stage 2

l1d_tlb_miss_ratio**, L1 Data TLB Miss Ratio, metric

This metric measures the ratio of level 1 data TLB accesses missed to the total number of level 1 data TLB accesses. This gives an indication of the effectiveness of the level 1 data TLB.

Units

This unit is expressed as per tlb access.

Formula

L1D_TLB_REFILL / L1D_TLB

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

L1D_TLB

L1D_TLB_REFILL

Metric group

DTLB_Effectiveness

Other metric group: Miss_Ratio

Methodology

Stage 2

l1d_tlb_mpki**, L1 Data TLB MPKI, metric

This metric measures the number of level 1 data TLB accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L1D_TLB_REFILL / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

INST_RETIRED

L1D_TLB_REFILL

Metric group

DTLB_Effectiveness

Other metric group: MPKI

Methodology

Stage 2

l2_tlb_miss_ratio***, L2 Unified TLB Miss Ratio, metric

This metric measures the ratio of level 2 unified TLB accesses missed to the total number of level 2 unified TLB accesses. This gives an indication of the effectiveness of the level 2 TLB.

Units

This unit is expressed as per tlb access.

Formula

L2D_TLB_REFILL / L2D_TLB

*** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

L2D_TLB
L2D_TLB_REFILL

Metric group

DTLB_Effectiveness
Other metric group: ITLB_Effectiveness
Other metric group: Miss_Ratio

Methodology

Stage 2

l2_tlb_mpki***, L2 Unified TLB MPKI, metric

This metric measures the number of level 2 unified TLB accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L2D_TLB_REFILL / INST_RETIRED * 1000

*** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

INST_RETIRED
L2D_TLB_REFILL

Metric group

DTLB_Effectiveness
Other metric group: ITLB_Effectiveness
Other metric group: MPKI

Methodology

Stage 2

6.9 L1I_Cache_Effectiveness metrics for Neoverse V2

L1 Instruction Cache Effectiveness. This metric group contains metrics to evaluate the effectiveness of L1 Instruction cache on this processor.

Summary of metrics in L1I_Cache_Effectiveness:

- Total metrics: 2

Table 6-9: L1I_Cache_Effectiveness metrics summary

Metric	Name	Description
l1i_cache_miss_ratio	L1I Cache Miss Ratio	This metric measures the ratio of level 1 instruction cache accesses missed to the total number...
l1i_cache_mpki	L1I Cache MPKI	This metric measures the number of level 1 instruction cache accesses missed per thousand...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

l1i_cache_miss_ratio**, L1I Cache Miss Ratio, metric

This metric measures the ratio of level 1 instruction cache accesses missed to the total number of level 1 instruction cache accesses. This gives an indication of the effectiveness of the level 1 instruction cache.

Units

This unit is expressed as per cache access.

Formula

L1I_CACHE_REFILL / L1I_CACHE

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

L1I_CACHE
L1I_CACHE_REFILL

Metric group

L1I_Cache_Effectiveness
Other metric group: Miss_Ratio

Methodology

Stage 2

l1i_cache_mpki**, L1I Cache MPKI, metric

This metric measures the number of level 1 instruction cache accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L1I_CACHE_REFILL / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

INST_RETIRED

L1I_CACHE_REFILL

Metric group

L1I_Cache_Effectiveness

Other metric group: MPKI

Methodology

Stage 2

6.10 L1D_Cache_Effectiveness metrics for Neoverse V2

L1 Data Cache Effectiveness. This metric group contains metrics to evaluate the effectiveness of L1 Data Cache on this processor.

Summary of metrics in L1D_Cache_Effectiveness:

- Total metrics: 2

Table 6-10: L1D_Cache_Effectiveness metrics summary

Metric	Name	Description
l1d_cache_miss_ratio	L1D Cache Miss Ratio	This metric measures the ratio of level 1 data cache accesses missed to the total number of level...
l1d_cache_mpki	L1D Cache MPKI	This metric measures the number of level 1 data cache accesses missed per thousand instructions...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

l1d_cache_miss_ratio**, L1D Cache Miss Ratio, metric

This metric measures the ratio of level 1 data cache accesses missed to the total number of level 1 data cache accesses. This gives an indication of the effectiveness of the level 1 data cache.

Units

This unit is expressed as per cache access.

Formula

L1D_CACHE_REFILL / L1D_CACHE

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

L1D_CACHE

L1D_CACHE_REFILL

Metric group

L1D_Cache_Effectiveness

Other metric group: Miss_Ratio

Methodology

Stage 2

l1d_cache_mpki**, L1D Cache MPKI, metric

This metric measures the number of level 1 data cache accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

L1D_CACHE_REFILL / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

INST_RETIRED
L1D_CACHE_REFILL

Metric group

L1D_Cache_Effectiveness
Other metric group: MPKI

Methodology

Stage 2

6.11 L2_Cache_Effectiveness metrics for Neoverse V2

L2 Unified Cache Effectiveness. This metric group contains metrics to evaluate the effectiveness of L2 Unified Cache on this processor.

Summary of metrics in L2_Cache_Effectiveness:

- Total metrics: 2

Table 6-11: L2_Cache_Effectiveness metrics summary

Metric	Name	Description
l2_cache_miss_ratio	L2 Cache Miss Ratio	This metric measures the ratio of level 2 cache accesses missed to the total number of level 2...
l2_cache_mpki	L2 Cache MPKI	This metric measures the number of level 2 unified cache accesses missed per thousand...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

l2_cache_miss_ratio**, L2 Cache Miss Ratio, metric

This metric measures the ratio of level 2 cache accesses missed to the total number of level 2 cache accesses. This gives an indication of the effectiveness of the level 2 cache, which is a unified

cache that stores both data and instruction. Note that cache accesses in this cache are either data memory access or instruction fetch as this is a unified cache.

Units

This unit is expressed as per cache access.

Formula

L2D_CACHE_REFILL / L2D_CACHE

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

L2D_CACHE
L2D_CACHE_REFILL

Metric group

L2_Cache_Effectiveness
Other metric group: Miss_Ratio

Methodology

Stage 2

l2_cache_mpki**, L2 Cache MPKI, metric

This metric measures the number of level 2 unified cache accesses missed per thousand instructions executed. Note that cache accesses in this cache are either data memory access or instruction fetch as this is a unified cache.

Units

This unit is expressed as mpki.

Formula

L2D_CACHE_REFILL / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

INST_RETIRED
L2D_CACHE_REFILL

Metric group

L2_Cache_Effectiveness
Other metric group: MPKI

Methodology

Stage 2

6.12 LL_Cache_Effectiveness metrics for Neoverse V2

Last Level Cache Effectiveness. This metric group contains metrics to evaluate the effectiveness of Last Level Cache on this processor.

Summary of metrics in LL_Cache_Effectiveness:

- Total metrics: 3

Table 6-12: LL_Cache_Effectiveness metrics summary

Metric	Name	Description
ll_cache_read_hit_ratio	LL Cache Read Hit Ratio	This metric measures the ratio of last level cache read accesses hit in the cache to the total...
ll_cache_read_miss_ratio	LL Cache Read Miss Ratio	This metric measures the ratio of last level cache read accesses missed to the total number of...
ll_cache_read_mpki	LL Cache Read MPKI	This metric measures the number of last level cache read accesses missed per thousand...

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

ll_cache_read_hit_ratio, LL Cache Read Hit Ratio, metric

This metric measures the ratio of last level cache read accesses hit in the cache to the total number of last level cache accesses. This gives an indication of the effectiveness of the last level cache for read traffic. Note that cache accesses in this cache are either data memory access or instruction fetch as this is a system level cache.

Units

This unit is expressed as per cache access.

Formula

(LL_CACHE_RD - LL_CACHE_MISS_RD) / LL_CACHE_RD

Related telemetry artifacts

Events

LL_CACHE_MISS_RD
LL_CACHE_RD

Metric group

LL_Cache_Effectiveness

Methodology

Stage 2

ll_cache_read_miss_ratio**, LL Cache Read Miss Ratio, metric

This metric measures the ratio of last level cache read accesses missed to the total number of last level cache accesses. This gives an indication of the effectiveness of the last level cache for read traffic. Note that cache accesses in this cache are either data memory access or instruction fetch as this is a system level cache.

Units

This unit is expressed as per cache access.

Formula

LL_CACHE_MISS_RD / LL_CACHE_RD

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

LL_CACHE_MISS_RD
LL_CACHE_RD

Metric group

LL_Cache_Effectiveness
Other metric group: Miss_Ratio

Methodology

Stage 2

ll_cache_read_mpki**, LL Cache Read MPKI, metric

This metric measures the number of last level cache read accesses missed per thousand instructions executed.

Units

This unit is expressed as mpki.

Formula

LL_CACHE_MISS_RD / INST_RETIRED * 1000

** This metric is used in multiple metric groups. See the following for more information.

Related telemetry artifacts

Events

INST_RETIRED
LL_CACHE_MISS_RD

Metric group

LL_Cache_Effectiveness
Other metric group: MPKI

Methodology

Stage 2

6.13 Operation_Mix metrics for Neoverse V2

Speculative Operation Mix. This metric group provides the distribution of micro-operation types executed for the program.

Summary of metrics in Operation_Mix:

- Total metrics: 7

Table 6-13: Operation_Mix metrics summary

Metric	Name	Description
branch_percentage	Branch Operations Percentage	This metric measures branch operations as a percentage of operations speculatively executed.
crypto_percentage	Crypto Operations Percentage	This metric measures crypto operations as a percentage of operations speculatively executed.
integer_dp_percentage	Integer Operations Percentage	This metric measures scalar integer operations as a percentage of operations speculatively executed.
load_percentage	Load Operations Percentage	This metric measures load operations as a percentage of operations speculatively executed.
scalar_fp_percentage	Floating Point Operations Percentage	This metric measures scalar floating point operations as a percentage of operations speculatively...
simd_percentage	Advanced SIMD Operations Percentage	This metric measures advanced SIMD operations as a percentage of total operations speculatively...
store_percentage	Store Operations Percentage	This metric measures store operations as a percentage of operations speculatively executed.

For a complete list of the metrics in Neoverse V2, see Metrics cheat sheet for Neoverse V2 and Metrics lookup table for Neoverse V2.

branch_percentage, Branch Operations Percentage, metric

This metric measures branch operations as a percentage of operations speculatively executed.

Units

This unit is expressed as percent of operations.

Formula

(BR_IMMED_SPEC + BR_INDIRECT_SPEC) / INST_SPEC * 100

Related telemetry artifacts

Events

BR_IMMED_SPEC
BR_INDIRECT_SPEC
INST_SPEC

Metric group

Operation_Mix

Methodology

Stage 2

crypto_percentage, Crypto Operations Percentage, metric

This metric measures crypto operations as a percentage of operations speculatively executed.

Units

This unit is expressed as percent of operations.

Formula

CRYPTO_SPEC / INST_SPEC * 100

Related telemetry artifacts

Events

CRYPTO_SPEC
INST_SPEC

Metric group

Operation_Mix

Methodology

Stage 2

integer_dp_percentage, Integer Operations Percentage, metric

This metric measures scalar integer operations as a percentage of operations speculatively executed.

Units

This unit is expressed as percent of operations.

Formula

DP_SPEC / INST_SPEC * 100

Related telemetry artifacts

Events

DP_SPEC
INST_SPEC

Metric group

Operation_Mix

Methodology

Stage 2

load_percentage, Load Operations Percentage, metric

This metric measures load operations as a percentage of operations speculatively executed.

Units

This unit is expressed as percent of operations.

Formula

LD_SPEC / INST_SPEC * 100

Related telemetry artifacts

Events

INST_SPEC
LD_SPEC

Metric group

Operation_Mix

Methodology

Stage 2

scalar_fp_percentage, Floating Point Operations Percentage, metric

This metric measures scalar floating point operations as a percentage of operations speculatively executed.

Units

This unit is expressed as percent of operations.

Formula

VFP_SPEC / INST_SPEC * 100

Related telemetry artifacts

Events

INST_SPEC
VFP_SPEC

Metric group

Operation_Mix

Methodology

Stage 2

simd_percentage, Advanced SIMD Operations Percentage, metric

This metric measures advanced SIMD operations as a percentage of total operations speculatively executed.

Units

This unit is expressed as percent of operations.

Formula

ASE_SPEC / INST_SPEC * 100

Related telemetry artifacts

Events

ASE_SPEC
INST_SPEC

Metric group

Operation_Mix

Methodology

Stage 2

store_percentage, Store Operations Percentage, metric

This metric measures store operations as a percentage of operations speculatively executed.

Units

This unit is expressed as percent of operations.

Formula

ST_SPEC / INST_SPEC * 100

Related telemetry artifacts

Events

INST_SPEC
ST_SPEC

Metric group

Operation_Mix

Methodology

Stage 2

7. PMU events by functional group in Neoverse V2

The Performance Monitoring Unit (PMU) collects events through an event interface from other units in the design. These events are used as triggers for event counters. Not all of the possible events are used in the Methodology, however, they are all listed for completeness.

Neoverse V2 provides the following types of PMU events:

- Total implemented Common events: 155
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

PMU events for Neoverse V2 are grouped into the following functional groups:

- Bus, BUS (4 events)
- Chain, CHAIN (1 events)
- Exception, EXCEPTION (15 events)
- L1D_Cache, L1D CACHE (13 events)
- L1I_Cache, L1I CACHE (3 events)
- L2_Cache, L2 CACHE (12 events)
- L3_Cache, L3 CACHE (5 events)
- LL_Cache, LL CACHE (2 events)
- Memory, MEMORY (11 events)
- Retired, RETIRED (7 events)
- SPE, SPE (4 events)
- Spec_Operation, SPEC OPERATION (27 events)
- FP_Operation, FP OPERATION (5 events)
- Stall, STALL (7 events)
- General, GENERAL (2 events)
- TLB, TLB (16 events)
- SVE, SVE (12 events)
- TRACE, TRACE (9 events)

7.1 Bus (BUS) events for Neoverse V2

Bus transaction related events.

Summary of events in Bus:

- Total implemented Common events: 4
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-1: Bus events summary

Code	Mnemonic	Name	Description
0x0019	BUS_ACCESS	Bus access	Counts memory transactions issued by the CPU to the external bus, including snoop requests and...
0x001D	BUS_CYCLES	Bus cycle	Counts bus cycles in the CPU. Bus cycles represent a clock cycle in which a transaction could be...
0x0060	BUS_ACCESS_RD	Bus access, read	Counts memory read transactions seen on the external bus. Each beat of data is counted individually.
0x0061	BUS_ACCESS_WR	Bus access, write	Counts memory write transactions seen on the external bus. Each beat of data is counted...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0019 BUS_ACCESS, Bus access, event

Counts memory transactions issued by the CPU to the external bus, including snoop requests and snoop responses. Each beat of data is counted individually.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Bus

0x001D BUS_CYCLES, Bus cycle, event

Counts bus cycles in the CPU. Bus cycles represent a clock cycle in which a transaction could be sent or received on the interface from the CPU to the external bus. Since that interface is driven at the same clock speed as the CPU, this event is a duplicate of CPU_CYCLES.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Bus

0x0060 BUS_ACCESS_RD, Bus access, read, event

Counts memory read transactions seen on the external bus. Each beat of data is counted individually.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Bus

0x0061 BUS_ACCESS_WR, Bus access, write, event

Counts memory write transactions seen on the external bus. Each beat of data is counted individually.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Bus

7.2 Chain (CHAIN) events for Neoverse V2

Chain related events.

Summary of events in Chain:

- Total implemented Common events: 1
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-2: Chain events summary

Code	Mnemonic	Name	Description
0x001E	CHAIN	Chain a pair of event counters	Counts whenever the even numbered PMU counter registers overflow. This event is used when the...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x001E CHAIN, Chain a pair of event counters, event

Counts whenever the even numbered PMU counter registers overflow. This event is used when the even/odd pairs of registers are used as a single counter.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Chain

7.3 Exception (EXCEPTION) events for Neoverse V2

Exception related events.

Summary of events in Exception:

- Total implemented Common events: 15
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-3: Exception events summary

Code	Mnemonic	Name	Description
0x0009	EXC_TAKEN	Exception taken	Counts any taken architecturally visible exceptions such as IRQ, FIQ, SError, and other...
0x000A	EXC_RETURN	Instruction architecturally executed, Condition code check pass, exception return	Counts any architecturally executed exception return instructions. Eg: AArch64: ERET
0x0081	EXC_UNDEF	Exception taken, other synchronous	Counts the number of synchronous exceptions which are taken locally that are due to attempting to...
0x0082	EXC_SVC	Exception taken, Supervisor Call	Counts SVC exceptions taken locally.
0x0083	EXC_PABORT	Exception taken, Instruction Abort	Counts synchronous exceptions that are taken locally and caused by Instruction Aborts.
0x0084	EXC_DABORT	Exception taken, Data Abort or SError	Counts exceptions that are taken locally and are caused by data aborts or SErrors. Conditions...
0x0086	EXC_IRQ	Exception taken, IRQ	Counts IRQ exceptions including the virtual IRQs that are taken locally.
0x0087	EXC_FIQ	Exception taken, FIQ	Counts FIQ exceptions including the virtual FIQs that are taken locally.
0x0088	EXC_SMC	Exception taken, Secure Monitor Call	Counts SMC exceptions take to EL3.
0x008A	EXC_HVC	Exception taken, Hypervisor Call	Counts HVC exceptions taken to EL2.
0x008B	EXC_TRAP_PABORT	Exception taken, Instruction Abort not Taken locally	Counts exceptions which are traps not taken locally and are caused by Instruction Aborts. For...
0x008C	EXC_TRAP_DABORT	Exception taken, Data Abort or SError not Taken locally	Counts exceptions which are traps not taken locally and are caused by Data Aborts or SError...
0x008D	EXC_TRAP_OTHER	Exception taken, other traps not Taken locally	Counts the number of synchronous trap exceptions which are not taken locally and are not SVC,...
0x008E	EXC_TRAP_IRQ	Exception taken, IRQ not Taken locally	Counts IRQ exceptions including the virtual IRQs that are not taken locally.

Code	Mnemonic	Name	Description
0x008F	EXC_TRAP_FIQ	Exception taken, FIQ not Taken locally	Counts FIQs which are not taken locally but taken from EL0, EL1, or EL2 to EL3 (which would be...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0009 EXC_TAKEN, Exception taken, event

Counts any taken architecturally visible exceptions such as IRQ, FIQ, SError, and other synchronous exceptions. Exceptions are counted whether or not they are taken locally.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x000A EXC_RETURN, Instruction architecturally executed, Condition code check pass, exception return, event

Counts any architecturally executed exception return instructions. Eg: AArch64: ERET

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x0081 EXC_UNDEF, Exception taken, other synchronous, event

Counts the number of synchronous exceptions which are taken locally that are due to attempting to execute an instruction that is **UNDEFINED**. Attempting to execute instruction bit patterns that have not been allocated. Attempting to execute instructions when they are disabled. Attempting to execute instructions at an inappropriate Exception level. Attempting to execute an instruction when the value of PSTATE.IL is 1.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x0082 EXC_SVC, Exception taken, Supervisor Call, event

Counts SVC exceptions taken locally.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x0083 EXC_PABORT, Exception taken, Instruction Abort, event

Counts synchronous exceptions that are taken locally and caused by Instruction Aborts.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x0084 EXC_DABORT, Exception taken, Data Abort or SError, event

Counts exceptions that are taken locally and are caused by data aborts or SErrors. Conditions that could cause those exceptions are attempting to read or write memory where the MMU generates a fault, attempting to read or write memory with a misaligned address, interrupts from the nSEI inputs and internally generated SErrors.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x0086 EXC_IRQ, Exception taken, IRQ, event

Counts IRQ exceptions including the virtual IRQs that are taken locally.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x0087 EXC_FIQ, Exception taken, FIQ, event

Counts FIQ exceptions including the virtual FIQs that are taken locally.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x0088 EXC_SMC, Exception taken, Secure Monitor Call, event

Counts SMC exceptions take to EL3.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x008A EXC_HVC, Exception taken, Hypervisor Call, event

Counts HVC exceptions taken to EL2.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x008B EXC_TRAP_PABORT, Exception taken, Instruction Abort not Taken locally, event

Counts exceptions which are traps not taken locally and are caused by Instruction Aborts. For example, attempting to execute an instruction with a misaligned PC.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x008C EXC_TRAP_DABORT, Exception taken, Data Abort or SError not Taken locally, event

Counts exceptions which are traps not taken locally and are caused by Data Aborts or SError interrupts. Conditions that could cause those exceptions are:

1. Attempting to read or write memory where the MMU generates a fault,
2. Attempting to read or write memory with a misaligned address,
3. Interrupts from the SEI input.
4. internally generated SErrors.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x008D EXC_TRAP_OTHER, Exception taken, other traps not Taken locally, event

Counts the number of synchronous trap exceptions which are not taken locally and are not SVC, SMC, HVC, data aborts, Instruction Aborts, or interrupts.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x008E EXC_TRAP_IRQ, Exception taken, IRQ not Taken locally, event

Counts IRQ exceptions including the virtual IRQs that are not taken locally.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

0x008F EXC_TRAP_FIQ, Exception taken, FIQ not Taken locally, event

Counts FIQs which are not taken locally but taken from EL0, EL1, or EL2 to EL3 (which would be the normal behavior for FIQs when not executing in EL3).

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Exception

7.4 L1D_Cache (L1D CACHE) events for Neoverse V2

L1 data cache related events.

Summary of events in L1D_Cache:

- Total implemented Common events: 13
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-4: L1D_Cache events summary

Code	Mnemonic	Name	Description
0x0003	L1D_CACHE_REFILL	Level 1 data cache refill	Counts level 1 data cache refills caused by speculatively executed load or store operations that...
0x0004	L1D_CACHE	Level 1 data cache access	Counts level 1 data cache accesses from any load/store operations. Atomic operations that resolve...
0x0015	L1D_CACHE_WB	Level 1 data cache write-back	Counts write-backs of dirty data from the L1 data cache to the L2 cache. This occurs when either...

Code	Mnemonic	Name	Description
0x0039	L1D_CACHE_LMISS_RD	Level 1 data cache long-latency read miss	Counts cache line refills into the level 1 data cache from any memory read operations, that...
0x0040	L1D_CACHE_RD	Level 1 data cache access, read	Counts level 1 data cache accesses from any load operation. Atomic load operations that resolve...
0x0041	L1D_CACHE_WR	Level 1 data cache access, write	Counts level 1 data cache accesses generated by store operations. This event also counts accesses...
0x0042	L1D_CACHE_REFILL_RD	Level 1 data cache refill, read	Counts level 1 data cache refills caused by speculatively executed load instructions where the...
0x0043	L1D_CACHE_REFILL_WR	Level 1 data cache refill, write	Counts level 1 data cache refills caused by speculatively executed store instructions where the...
0x0044	L1D_CACHE_REFILL_INNER	Level 1 data cache refill, inner	Counts level 1 data cache refills where the cache line data came from caches inside the immediate...
0x0045	L1D_CACHE_REFILL_OUTER	Level 1 data cache refill, outer	Counts level 1 data cache refills for which the cache line data came from outside the immediate...
0x0046	L1D_CACHE_WB_VICTIM	Level 1 data cache write-back, victim	Counts dirty cache line evictions from the level 1 data cache caused by a new cache line...
0x0047	L1D_CACHE_WB_CLEAN	Level 1 data cache write-back, cleaning and coherency	Counts write-backs from the level 1 data cache that are a result of a coherency operation made by...
0x0048	L1D_CACHE_INVAL	Level 1 data cache invalidate	Counts each explicit invalidation of a cache line in the level 1 data cache caused by: • Cache...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0003 L1D_CACHE_REFILL, Level 1 data cache refill, event

Counts level 1 data cache refills caused by speculatively executed load or store operations that missed in the level 1 data cache. This event only counts one event per cache line. This event does not count cache line allocations from preload instructions or from hardware cache prefetching.

Related telemetry artifacts

Metrics

- l1d_cache_mpki in L1D_Cache_Effectiveness
- l1d_cache_mpki in MPKI
- l1d_cache_miss_ratio in L1D_Cache_Effectiveness
- l1d_cache_miss_ratio in Miss_Ratio

Metric groups

- L1D_Cache_Effectiveness
- MPKI
- Miss_Ratio

Functional groups

- L1D_Cache

0x0004 L1D_CACHE, Level 1 data cache access, event

Counts level 1 data cache accesses from any load/store operations. Atomic operations that resolve in the CPUs caches (near atomic operations) counts as both a write access and read access. Each access to a cache line is counted including the multiple accesses caused by single instructions such as LDM or STM. Each access to other level 1 data or unified memory structures, for example refill buffers, write buffers, and write-back buffers, are also counted.

Related telemetry artifacts

Metrics

l1d_cache_miss_ratio in L1D_Cache_Effectiveness

l1d_cache_miss_ratio in Miss_Ratio

Metric groups

L1D_Cache_Effectiveness

Miss_Ratio

Functional groups

L1D_Cache

0x0015 L1D_CACHE_WB, Level 1 data cache write-back, event

Counts write-backs of dirty data from the L1 data cache to the L2 cache. This occurs when either a dirty cache line is evicted from L1 data cache and allocated in the L2 cache or dirty data is written to the L2 and possibly to the next level of cache. This event counts both victim cache line evictions and cache write-backs from snoops or cache maintenance operations. The following cache operations are not counted:

1. Invalidations which do not result in data being transferred out of the L1 (such as evictions of clean data),
2. Full line writes which write to L2 without writing L1, such as write streaming mode.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0039 L1D_CACHE_LMISS_RD, Level 1 data cache long-latency read miss, event

Counts cache line refills into the level 1 data cache from any memory read operations, that incurred additional latency.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0040 L1D_CACHE_RD, Level 1 data cache access, read, event

Counts level 1 data cache accesses from any load operation. Atomic load operations that resolve in the CPUs caches counts as both a write access and read access.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0041 L1D_CACHE_WR, Level 1 data cache access, write, event

Counts level 1 data cache accesses generated by store operations. This event also counts accesses caused by a DC ZVA (data cache zero, specified by virtual address) instruction. Near atomic operations that resolve in the CPUs caches count as a write access and read access.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0042 L1D_CACHE_REFILL_RD, Level 1 data cache refill, read, event

Counts level 1 data cache refills caused by speculatively executed load instructions where the memory read operation misses in the level 1 data cache. This event only counts one event per cache line.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0043 L1D_CACHE_REFILL_WR, Level 1 data cache refill, write, event

Counts level 1 data cache refills caused by speculatively executed store instructions where the memory write operation misses in the level 1 data cache. This event only counts one event per cache line.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0044 L1D_CACHE_REFILL_INNER, Level 1 data cache refill, inner, event

Counts level 1 data cache refills where the cache line data came from caches inside the immediate cluster of the core.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0045 L1D_CACHE_REFILL_OUTER, Level 1 data cache refill, outer, event

Counts level 1 data cache refills for which the cache line data came from outside the immediate cluster of the core, like an SLC in the system interconnect or DRAM.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0046 L1D_CACHE_WB_VICTIM, Level 1 data cache write-back, victim, event

Counts dirty cache line evictions from the level 1 data cache caused by a new cache line allocation. This event does not count evictions caused by cache maintenance operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0047 L1D_CACHE_WB_CLEAN, Level 1 data cache write-back, cleaning and coherency, event

Counts write-backs from the level 1 data cache that are a result of a coherency operation made by another CPU. Event count includes cache maintenance operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

0x0048 L1D_CACHE_INVAL, Level 1 data cache invalidate, event

Counts each explicit invalidation of a cache line in the level 1 data cache caused by:

- Cache Maintenance Operations (CMO) that operate by a virtual address.
- Broadcast cache coherency operations from another CPU in the system.

This event does not count for the following conditions:

1. A cache refill invalidates a cache line.

2. A CMO which is executed on that CPU and invalidates a cache line specified by set/way.

Note that CMOs that operate by set/way cannot be broadcast from one CPU to another.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1D_Cache

7.5 L1I_Cache (L1I CACHE) events for Neoverse V2

L1 instruction cache related events.

Summary of events in L1I_Cache:

- Total implemented Common events: 3
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-5: L1I_Cache events summary

Code	Mnemonic	Name	Description
0x0001	L1I_CACHE_REFILL	Level 1 instruction cache refill	Counts cache line refills in the level 1 instruction cache caused by a missed instruction fetch....
0x0014	L1I_CACHE	Level 1 instruction cache access	Counts instruction fetches which access the level 1 instruction cache. Instruction cache accesses...
0x4006	L1I_CACHE_LMISS	Level 1 instruction cache long-latency miss	Counts cache line refills into the level 1 instruction cache, that incurred additional latency.

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0001 L1I_CACHE_REFILL, Level 1 instruction cache refill, event

Counts cache line refills in the level 1 instruction cache caused by a missed instruction fetch. Instruction fetches may include accessing multiple instructions, but the single cache line allocation is counted once.

Related telemetry artifacts

Metrics

l1i_cache_mpki in L1I_Cache_Effectiveness
l1i_cache_mpki in MPKI
l1i_cache_miss_ratio in L1I_Cache_Effectiveness
l1i_cache_miss_ratio in Miss_Ratio

Metric groups

L1I_Cache_Effectiveness
MPKI
Miss_Ratio

Functional groups

L1I_Cache

`0x0014` L1I_CACHE, Level 1 instruction cache access, event

Counts instruction fetches which access the level 1 instruction cache. Instruction cache accesses caused by cache maintenance operations are not counted.

Related telemetry artifacts

Metrics

l1i_cache_miss_ratio in L1I_Cache_Effectiveness
l1i_cache_miss_ratio in Miss_Ratio

Metric groups

L1I_Cache_Effectiveness
Miss_Ratio

Functional groups

L1I_Cache

`0x4006` L1I_CACHE_LMISS, Level 1 instruction cache long-latency miss, event

Counts cache line refills into the level 1 instruction cache, that incurred additional latency.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L1I_Cache

7.6 L2_Cache (L2 CACHE) events for Neoverse V2

L2 unified cache related events.

Summary of events in L2_Cache:

- Total implemented Common events: 12
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-6: L2_Cache events summary

Code	Mnemonic	Name	Description
0x0016	L2D_CACHE	Level 2 data cache access	Counts level 2 cache accesses. Level 2 cache is a unified cache for data and instruction...
0x0017	L2D_CACHE_REFILL	Level 2 data cache refill	Counts cache line refills into the level 2 cache. Level 2 cache is a unified cache for data and...
0x0018	L2D_CACHE_WB	Level 2 data cache write-back	Counts write-backs of data from the L2 cache to outside the CPU. This includes snoops to the L2...
0x0020	L2D_CACHE_ALLOCATE	Level 2 data cache allocation without refill	Counts level 2 cache line allocates that do not fetch data from outside the level 2 data or...
0x0050	L2D_CACHE_RD	Level 2 data cache access, read	Counts level 2 cache accesses due to memory read operations. Level 2 cache is a unified cache for...
0x0051	L2D_CACHE_WR	Level 2 data cache access, write	Counts level 2 cache accesses due to memory write operations. Level 2 cache is a unified cache...
0x0052	L2D_CACHE_REFILL_RD	Level 2 data cache refill, read	Counts refills for memory accesses due to memory read operation counted by L2D_CACHE_RD. Level 2...
0x0053	L2D_CACHE_REFILL_WR	Level 2 data cache refill, write	Counts refills for memory accesses due to memory write operation counted by L2D_CACHE_WR. Level 2...
0x0056	L2D_CACHE_WB_VICTIM	Level 2 data cache write-back, victim	Counts evictions from the level 2 cache because of a line being allocated into the L2 cache.
0x0057	L2D_CACHE_WB_CLEAN	Level 2 data cache write-back, cleaning and coherency	Counts write-backs from the level 2 cache that are a result of either: 1. Cache maintenance...
0x0058	L2D_CACHE_INVAL	Level 2 data cache invalidate	Counts each explicit invalidation of a cache line in the level 2 cache by cache maintenance...
0x4009	L2D_CACHE_LMISS_RD	Level 2 data cache long-latency read miss	Counts cache line refills into the level 2 unified cache from any memory read operations that...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0016 L2D_CACHE, Level 2 data cache access, event

Counts level 2 cache accesses. Level 2 cache is a unified cache for data and instruction accesses. Accesses are for misses in the first level caches or translation resolutions due to accesses. This event also counts write back of dirty data from level 1 data cache to the L2 cache.

Related telemetry artifacts

Metrics

l2_cache_miss_ratio in L2_Cache_Effectiveness

l2_cache_miss_ratio in Miss_Ratio

Metric groups

L2_Cache_Effectiveness

Miss_Ratio

Functional groups

L2_Cache

0x0017 L2D_CACHE_REFILL, Level 2 data cache refill, event

Counts cache line refills into the level 2 cache. Level 2 cache is a unified cache for data and instruction accesses. Accesses are for misses in the level 1 caches or translation resolutions due to accesses.

Related telemetry artifacts

Metrics

- l2_cache_mpki in L2_Cache_Effectiveness
- l2_cache_mpki in MPKI
- l2_cache_miss_ratio in L2_Cache_Effectiveness
- l2_cache_miss_ratio in Miss_Ratio

Metric groups

- Miss_Ratio
- MPKI
- L2_Cache_Effectiveness

Functional groups

- L2_Cache

0x0018 L2D_CACHE_WB, Level 2 data cache write-back, event

Counts write-backs of data from the L2 cache to outside the CPU. This includes snoops to the L2 (from other CPUs) which return data even if the snoops cause an invalidation. L2 cache line invalidations which do not write data outside the CPU and snoops which return data from an L1 cache are not counted. Data would not be written outside the cache when invalidating a clean cache line.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

- L2_Cache

0x0020 L2D_CACHE_ALLOCATE, Level 2 data cache allocation without refill, event

Counts level 2 cache line allocates that do not fetch data from outside the level 2 data or unified cache.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

- L2_Cache

0x0050 L2D_CACHE_RD, Level 2 data cache access, read, event

Counts level 2 cache accesses due to memory read operations. Level 2 cache is a unified cache for data and instruction accesses, accesses are for misses in the level 1 caches or translation resolutions due to accesses.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L2_Cache

0x0051 L2D_CACHE_WR, Level 2 data cache access, write, event

Counts level 2 cache accesses due to memory write operations. Level 2 cache is a unified cache for data and instruction accesses, accesses are for misses in the level 1 caches or translation resolutions due to accesses.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L2_Cache

0x0052 L2D_CACHE_REFILL_RD, Level 2 data cache refill, read, event

Counts refills for memory accesses due to memory read operation counted by L2D_CACHE_RD. Level 2 cache is a unified cache for data and instruction accesses, accesses are for misses in the level 1 caches or translation resolutions due to accesses.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L2_Cache

0x0053 L2D_CACHE_REFILL_WR, Level 2 data cache refill, write, event

Counts refills for memory accesses due to memory write operation counted by L2D_CACHE_WR. Level 2 cache is a unified cache for data and instruction accesses, accesses are for misses in the level 1 caches or translation resolutions due to accesses.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L2_Cache

0x0056 L2D_CACHE_WB_VICTIM, Level 2 data cache write-back, victim, event

Counts evictions from the level 2 cache because of a line being allocated into the L2 cache.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L2_Cache

0x0057 L2D_CACHE_WB_CLEAN, Level 2 data cache write-back, cleaning and coherency, event

Counts write-backs from the level 2 cache that are a result of either:

1. Cache maintenance operations,
2. Snoop responses or,
3. Direct cache transfers to another CPU due to a forwarding snoop request.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L2_Cache

0x0058 L2D_CACHE_INVAL, Level 2 data cache invalidate, event

Counts each explicit invalidation of a cache line in the level 2 cache by cache maintenance operations that operate by a virtual address, or by external coherency operations. This event does not count if either:

1. A cache refill invalidates a cache line or,
2. A Cache Maintenance Operation (CMO), which invalidates a cache line specified by set/way, is executed on that CPU.

CMOs that operate by set/way cannot be broadcast from one CPU to another.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L2_Cache

0x4009 L2D_CACHE_LMISS_RD, Level 2 data cache long-latency read miss, event

Counts cache line refills into the level 2 unified cache from any memory read operations that incurred additional latency.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L2_Cache

7.7 L3_Cache (L3 CACHE) events for Neoverse V2

L3 unified cache related events.

Summary of events in L3_Cache:

- Total implemented Common events: 5
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-7: L3_Cache events summary

Code	Mnemonic	Name	Description
0x0029	L3D_CACHE_ALLOCATE	Level 3 data cache allocation without refill	Counts level 3 cache line allocates that do not fetch data from outside the level 3 data or...
0x002A	L3D_CACHE_REFILL	Level 3 data cache refill	Counts level 3 accesses that receive data from outside the L3 cache.
0x002B	L3D_CACHE	Level 3 data cache access	Counts level 3 cache accesses. Level 3 cache is a unified cache for data and instruction...
0x00A0	L3D_CACHE_RD	Level 3 data cache access, read	Counts level 3 cache accesses caused by any memory read operation. Level 3 cache is a unified...
0x400B	L3D_CACHE_LMISS_RD	Level 3 data cache long-latency read miss	Counts any cache line refill into the level 3 cache from memory read operations that incurred...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0029 L3D_CACHE_ALLOCATE, Level 3 data cache allocation without refill, event

Counts level 3 cache line allocates that do not fetch data from outside the level 3 data or unified cache. For example, allocates due to streaming stores.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L3_Cache

0x002A L3D_CACHE_REFILL, Level 3 data cache refill, event

Counts level 3 accesses that receive data from outside the L3 cache.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L3_Cache

0x002B L3D_CACHE, Level 3 data cache access, event

Counts level 3 cache accesses. Level 3 cache is a unified cache for data and instruction accesses. Accesses are for misses in the lower level caches or translation resolutions due to accesses.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L3_Cache

0x00A0 L3D_CACHE_RD, Level 3 data cache access, read, event

Counts level 3 cache accesses caused by any memory read operation. Level 3 cache is a unified cache for data and instruction accesses. Accesses are for misses in the lower level caches or translation resolutions due to accesses.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L3_Cache

0x400B L3D_CACHE_LMISS_RD, Level 3 data cache long-latency read miss, event

Counts any cache line refill into the level 3 cache from memory read operations that incurred additional latency.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

L3_Cache

7.8 LL_Cache (LL CACHE) events for Neoverse V2

Last Level Cache related events.

Summary of events in LL_Cache:

- Total implemented Common events: 2
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-8: LL_Cache events summary

Code	Mnemonic	Name	Description
0x0036	LL_CACHE_RD	Last level cache access, read	Counts read transactions that were returned from outside the core cluster. This event counts for...
0x0037	LL_CACHE_MISS_RD	Last level cache miss, read	Counts read transactions that were returned from outside the core cluster but missed in the...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0036 LL_CACHE_RD, Last level cache access, read, event

Counts read transactions that were returned from outside the core cluster. This event counts for external last level cache when the system register CPUECTLR.EXTLLC bit is set, otherwise it counts for the L3 cache. This event counts read transactions returned from outside the core if those transactions are either hit in the system level cache or missed in the SLC and are returned from any other external sources.

Related telemetry artifacts

Metrics

ll_cache_read_miss_ratio in LL_Cache_Effectiveness
ll_cache_read_miss_ratio in Miss_Ratio
ll_cache_read_hit_ratio

Metric groups

LL_Cache_Effectiveness
Miss_Ratio

Functional groups

LL_Cache

0x0037 LL_CACHE_MISS_RD, Last level cache miss, read, event

Counts read transactions that were returned from outside the core cluster but missed in the system level cache. This event counts for external last level cache when the system register CPUECTLR.EXTLLC bit is set, otherwise it counts for L3 cache. This event counts read transactions returned from outside the core if those transactions are missed in the System level Cache. The data source of the transaction is indicated by a field in the CHI transaction returning to the CPU. This event does not count reads caused by cache maintenance operations.

Related telemetry artifacts

Metrics

ll_cache_read_mpki in LL_Cache_Effectiveness
ll_cache_read_mpki in MPKI
ll_cache_read_miss_ratio in LL_Cache_Effectiveness
ll_cache_read_miss_ratio in Miss_Ratio
ll_cache_read_hit_ratio

Metric groups

LL_Cache_Effectiveness
MPKI
Miss_Ratio

Functional groups

LL_Cache

7.9 Memory (MEMORY) events for Neoverse V2

Memory system related events.

Summary of events in Memory:

- Total implemented Common events: 11
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-9: Memory events summary

Code	Mnemonic	Name	Description
0x0013	MEM_ACCESS	Data memory access	Counts memory accesses issued by the CPU load store unit, where those accesses are issued due to...
0x001A	MEMORY_ERROR	Local memory error	Counts any detected correctable or uncorrectable physical memory errors (ECC or parity) in...
0x0031	REMOTE_ACCESS	Access to another socket in a multi-socket system	Counts accesses to another chip, which is implemented as a different CMN mesh in the system. If...
0x0066	MEM_ACCESS_RD	Data memory access, read	Counts memory accesses issued by the CPU due to load operations. The event counts any memory load...
0x0067	MEM_ACCESS_WR	Data memory access, write	Counts memory accesses issued by the CPU due to store operations. The event counts any memory...
0x4020	LDST_ALIGN_LAT	Access with additional latency from alignment	Counts the number of memory read and write accesses in a cycle that incurred additional latency,...
0x4021	LD_ALIGN_LAT	Load with additional latency from alignment	Counts the number of memory read accesses in a cycle that incurred additional latency, due to the...
0x4022	ST_ALIGN_LAT	Store with additional latency from alignment	Counts the number of memory write access in a cycle that incurred additional latency, due to the...
0x4024	MEM_ACCESS_CHECKED	Checked data memory access	Counts the number of memory read and write accesses in a cycle that are tag checked by the Memory...
0x4025	MEM_ACCESS_CHECKED_RD	Checked data memory access, read	Counts the number of memory read accesses in a cycle that are tag checked by the Memory Tagging...
0x4026	MEM_ACCESS_CHECKED_WR	Checked data memory access, write	Counts the number of memory write accesses in a cycle that is tag checked by the Memory Tagging...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0013 MEM_ACCESS, Data memory access, event

Counts memory accesses issued by the CPU load store unit, where those accesses are issued due to load or store operations. This event counts memory accesses no matter whether the data is received from any level of cache hierarchy or external memory. If memory accesses are broken up into smaller transactions than what were specified in the load or store instructions, then the event counts those smaller memory transactions.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x001A MEMORY_ERROR, Local memory error, event

Counts any detected correctable or uncorrectable physical memory errors (ECC or parity) in protected CPUs RAMs. On the core, this event counts errors in the caches (including data and tag rams). Any detected memory error (from either a speculative and abandoned access, or an architecturally executed access) is counted. Note that errors are only detected when the actual protected memory is accessed by an operation.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x0031 REMOTE_ACCESS, Access to another socket in a multi-socket system, event

Counts accesses to another chip, which is implemented as a different CMN mesh in the system. If the CHI bus response back to the core indicates that the data source is from another chip (mesh), then the counter is updated. If no data is returned, even if the system snoops another chip/mesh, then the counter is not updated.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x0066 MEM_ACCESS_RD, Data memory access, read, event

Counts memory accesses issued by the CPU due to load operations. The event counts any memory load access, no matter whether the data is received from any level of cache hierarchy or external memory. The event also counts atomic load operations. If memory accesses are broken up by the load/store unit into smaller transactions that are issued by the bus interface, then the event counts those smaller transactions.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x0067 MEM_ACCESS_WR, Data memory access, write, event

Counts memory accesses issued by the CPU due to store operations. The event counts any memory store access, no matter whether the data is located in any level of cache or external memory. The event also counts atomic load and store operations. If memory accesses are broken up by the load/store unit into smaller transactions that are issued by the bus interface, then the event counts those smaller transactions.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x4020 LDST_ALIGN_LAT, Access with additional latency from alignment, event

Counts the number of memory read and write accesses in a cycle that incurred additional latency, due to the alignment of the address and the size of data being accessed, which results in store crossing a single cache line.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x4021 LD_ALIGN_LAT, Load with additional latency from alignment, event

Counts the number of memory read accesses in a cycle that incurred additional latency, due to the alignment of the address and size of data being accessed, which results in load crossing a single cache line.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x4022 ST_ALIGN_LAT, Store with additional latency from alignment, event

Counts the number of memory write access in a cycle that incurred additional latency, due to the alignment of the address and size of data being accessed incurred additional latency.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x4024 MEM_ACCESS_CHECKED, Checked data memory access, event

Counts the number of memory read and write accesses in a cycle that are tag checked by the Memory Tagging Extension (MTE).

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x4025 MEM_ACCESS_CHECKED_RD, Checked data memory access, read, event

Counts the number of memory read accesses in a cycle that are tag checked by the Memory Tagging Extension (MTE).

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

0x4026 MEM_ACCESS_CHECKED_WR, Checked data memory access, write, event

Counts the number of memory write accesses in a cycle that is tag checked by the Memory Tagging Extension (MTE).

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Memory

7.10 Retired (RETIRED) events for Neoverse V2

Retired instruction and operation events.

Summary of events in Retired:

- Total implemented Common events: 7
- Total Implemented Product ImpDef events: 0

- PMU Only events : 0
- ETE Only events : 0

Table 7-10: Retired events summary

Code	Mnemonic	Name	Description
0x0000	SW_INCR	Instruction architecturally executed, Condition code check pass, software increment	Counts software writes to the PMSWINC_EL0 (software PMU increment) register. The PMSWINC_EL0...
0x0008	INST_RETIRED	Instruction architecturally executed	Counts instructions that have been architecturally executed.
0x000B	CID_WRITE_RETIRED	Instruction architecturally executed, Condition code check pass, write to CONTEXTIDR	Counts architecturally executed writes to the CONTEXTIDR register, which usually contain the...
0x001C	TTBR_WRITE_RETIRED	Instruction architecturally executed, Condition code check pass, write to TTBR	Counts architectural writes to TTBR0/1_EL1. If virtualization host extensions are enabled (by...
0x0021	BR_RETIRED	Instruction architecturally executed, branch	Counts architecturally executed branches, whether the branch is taken or not. Instructions that...
0x0022	BR_MIS_PRED_RETIRED	Branch instruction architecturally executed, mispredicted	Counts branches counted by BR_RETIRED which were mispredicted and caused a pipeline flush.
0x003A	OP_RETIRED	Micro-operation architecturally executed	Counts micro-operations that are architecturally executed. This is a count of number of...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0000 SW_INCR, Instruction architecturally executed, Condition code check pass, software increment, event

Counts software writes to the PMSWINC_EL0 (software PMU increment) register. The PMSWINC_EL0 register is a manually updated counter for use by application software.

This event could be used to measure any user program event, such as accesses to a particular data structure (by writing to the PMSWINC_EL0 register each time the data structure is accessed).

To use the PMSWINC_EL0 register and event, developers must insert instructions that write to the PMSWINC_EL0 register into the source code.

Since the SW_INCR event records writes to the PMSWINC_EL0 register, there is no need to do a read/increment/write sequence to the PMSWINC_EL0 register.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Retired

0x0008 INST_RETIRED, Instruction architecturally executed, event

Counts instructions that have been architecturally executed.

Related telemetry artifacts

Metrics

- ipc
- branch_mpki in Branch_Effectiveness
- branch_mpki in MPKI
- itlb_mpki in ITLB_Effectiveness
- itlb_mpki in MPKI
- l1i_tlb_mpki in ITLB_Effectiveness
- l1i_tlb_mpki in MPKI
- dtlb_mpki in DTLB_Effectiveness
- dtlb_mpki in MPKI
- l1d_tlb_mpki in DTLB_Effectiveness
- l1d_tlb_mpki in MPKI
- l2_tlb_mpki in ITLB_Effectiveness
- l2_tlb_mpki in DTLB_Effectiveness
- l2_tlb_mpki in MPKI
- l1i_cache_mpki in L1I_Cache_Effectiveness
- l1i_cache_mpki in MPKI
- l1d_cache_mpki in L1D_Cache_Effectiveness
- l1d_cache_mpki in MPKI
- l2_cache_mpki in L2_Cache_Effectiveness
- l2_cache_mpki in MPKI
- ll_cache_read_mpki in LL_Cache_Effectiveness
- ll_cache_read_mpki in MPKI

Metric groups

- LL_Cache_Effectiveness
- MPKI
- L1D_Cache_Effectiveness
- L2_Cache_Effectiveness
- Branch_Effectiveness
- General
- ITLB_Effectiveness
- L1I_Cache_Effectiveness
- DTLB_Effectiveness

Functional groups

- Retired

0x000B CID_WRITE_RETIRED, Instruction architecturally executed, Condition code check pass, write to CONTEXTIDR, event

Counts architecturally executed writes to the CONTEXTIDR register, which usually contain the kernel PID and can be output with hardware trace.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Retired

0x001C TTBR_WRITE_RETIRED, Instruction architecturally executed, Condition code check pass, write to TTBR, event

Counts architectural writes to TTBR0/1_EL1. If virtualization host extensions are enabled (by setting the HCR_EL2.E2H bit to 1), then accesses to TTBR0/1_EL1 that are redirected to TTBR0/1_EL2, or accesses to TTBR0/1_EL12, are counted. TTBRn registers are typically updated when the kernel is swapping user-space threads or applications.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Retired

0x0021 BR_RETIRED, Instruction architecturally executed, branch, event

Counts architecturally executed branches, whether the branch is taken or not. Instructions that explicitly write to the PC are also counted.

Related telemetry artifacts

Metrics

branch_misprediction_ratio in Branch_Effectiveness
branch_misprediction_ratio in Miss_Ratio

Metric groups

Branch_Effectiveness
Miss_Ratio

Functional groups

Retired

0x0022 BR_MIS_PRED_RETIRED, Branch instruction architecturally executed, mispredicted, event

Counts branches counted by BR_RETIRED which were mispredicted and caused a pipeline flush.

Related telemetry artifacts

Metrics

branch_mpki in Branch_Effectiveness
branch_mpki in MPKI
branch_misprediction_ratio in Branch_Effectiveness
branch_misprediction_ratio in Miss_Ratio

Metric groups

Branch_Effectiveness
MPKI
Miss_Ratio

Functional groups

Retired

0x003A OP_RETIRED, Micro-operation architecturally executed, event

Counts micro-operations that are architecturally executed. This is a count of number of micro-operations retired from the commit queue in a single cycle.

Related telemetry artifacts

Metrics

retiring
bad_speculation

Metric groups

Topdown_L1

Functional groups

Retired

7.11 SPE (SPE) events for Neoverse V2

SPE related events.

Summary of events in SPE:

- Total implemented Common events: 4
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-11: SPE events summary

Code	Mnemonic	Name	Description
0x4000	SAMPLE_POP	Statistical Profiling sample population	Counts statistical profiling sample population, the count of all operations that could be sampled...
0x4001	SAMPLE_FEED	Statistical Profiling sample taken	Counts statistical profiling samples taken for sampling.
0x4002	SAMPLE_FILTRATE	Statistical Profiling sample taken and not removed by filtering	Counts statistical profiling samples taken which are not removed by filtering.
0x4003	SAMPLE_COLLISION	Statistical Profiling sample collided with previous sample	Counts statistical profiling samples that have collided with a previous sample and so therefore...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x4000 SAMPLE_POP, Statistical Profiling sample population, event

Counts statistical profiling sample population, the count of all operations that could be sampled but may or may not be chosen for sampling.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SPE

0x4001 SAMPLE_FEED, Statistical Profiling sample taken, event

Counts statistical profiling samples taken for sampling.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SPE

0x4002 SAMPLE_FILTRATE, Statistical Profiling sample taken and not removed by filtering, event

Counts statistical profiling samples taken which are not removed by filtering.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SPE

0x4003 SAMPLE_COLLISION, Statistical Profiling sample collided with previous sample, event

Counts statistical profiling samples that have collided with a previous sample and so therefore not taken.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SPE

7.12 Spec_Operation (SPEC OPERATION) events for Neoverse V2

Speculatively executed operations related events.

Summary of events in Spec_Operation:

- Total implemented Common events: 27
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-12: Spec_Operation events summary

Code	Mnemonic	Name	Description
0x0010	BR_MIS_PRED	Branch instruction speculatively executed, mispredicted or not predicted	Counts branches which are speculatively executed and mispredicted.
0x0012	BR_PRED	Predictable branch instruction speculatively executed	Counts branches speculatively executed and were predicted right.
0x001B	INST_SPEC	Operation speculatively executed	Counts operations that have been speculatively executed.
0x003B	OP_SPEC	Micro-operation speculatively executed	Counts micro-operations speculatively executed. This is the count of the number of...
0x0068	UNALIGNED_LD_SPEC	Unaligned access, read	Counts unaligned memory read operations issued by the CPU. This event counts unaligned accesses...
0x0069	UNALIGNED_ST_SPEC	Unaligned access, write	Counts unaligned memory write operations issued by the CPU. This event counts unaligned accesses...
0x006A	UNALIGNED_LDST_SPEC	Unaligned access	Counts unaligned memory operations issued by the CPU. This event counts unaligned accesses (as...
0x006C	LDREX_SPEC	Exclusive operation speculatively executed, Load-Exclusive	Counts Load-Exclusive operations that have been speculatively executed. Eg: LDREX, LDX
0x006D	STREX_PASS_SPEC	Exclusive operation speculatively executed, Store-Exclusive pass	Counts store-exclusive operations that have been speculatively executed and have successfully...
0x006E	STREX_FAIL_SPEC	Exclusive operation speculatively executed, Store-Exclusive fail	Counts store-exclusive operations that have been speculatively executed and have not successfully...
0x006F	STREX_SPEC	Exclusive operation speculatively executed, Store-Exclusive	Counts store-exclusive operations that have been speculatively executed.
0x0070	LD_SPEC	Operation speculatively executed, load	Counts speculatively executed load operations including Single Instruction Multiple Data (SIMD)...
0x0071	ST_SPEC	Operation speculatively executed, store	Counts speculatively executed store operations including Single Instruction Multiple Data (SIMD)...
0x0073	DP_SPEC	Operation speculatively executed, integer data processing	Counts speculatively executed logical or arithmetic instructions such as MOV/MVN operations.
0x0074	ASE_SPEC	Operation speculatively executed, Advanced SIMD	Counts speculatively executed Advanced SIMD operations excluding load, store and move...
0x0075	VFP_SPEC	Operation speculatively executed, scalar floating-point	Counts speculatively executed floating point operations. This event does not count operations...

Code	Mnemonic	Name	Description
0x0076	PC_WRITE_SPEC	Operation speculatively executed, Software change of the PC	Counts speculatively executed operations which cause software changes of the PC. Those operations...
0x0077	CRYPTO_SPEC	Operation speculatively executed, Cryptographic instruction	Counts speculatively executed cryptographic operations except for PMULL and VMULL operations.
0x0078	BR_IMMED_SPEC	Branch speculatively executed, immediate branch	Counts immediate branch operations which are speculatively executed.
0x0079	BR_RETURN_SPEC	Branch speculatively executed, procedure return	Counts procedure return operations (RET) which are speculatively executed.
0x007A	BR_INDIRECT_SPEC	Branch speculatively executed, indirect branch	Counts indirect branch operations including procedure returns, which are speculatively executed....
0x007C	ISB_SPEC	Barrier speculatively executed, ISB	Counts ISB operations that are executed.
0x007D	DSB_SPEC	Barrier speculatively executed, DSB	Counts DSB operations that are speculatively issued to Load/Store unit in the CPU.
0x007E	DMB_SPEC	Barrier speculatively executed, DMB	Counts DMB operations that are speculatively issued to the Load/Store unit in the CPU. This event...
0x0090	RC_LD_SPEC	Release consistency operation speculatively executed, Load-Acquire	Counts any load acquire operations that are speculatively executed. Eg: LDAR, LDARH, LDARB
0x0091	RC_ST_SPEC	Release consistency operation speculatively executed, Store-Release	Counts any store release operations that are speculatively executed. Eg: STLR, STLRH, STLRB'
0x8005	ASE_INST_SPEC	Operation speculatively executed, Advanced SIMD	Counts speculatively executed Advanced SIMD operations.

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0010 BR_MIS_PRED, Branch instruction speculatively executed, mispredicted or not predicted, event

Counts branches which are speculatively executed and mispredicted.

Related telemetry artifacts

Metrics

frontend_bound
backend_bound
bad_speculation

Metric groups

Topdown_L1

Functional groups

Spec_Operation

0x0012 BR_PRED, Predictable branch instruction speculatively executed, event

Counts branches speculatively executed and were predicted right.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x001B INST_SPEC, Operation speculatively executed, event

Counts operations that have been speculatively executed.

Related telemetry artifacts

Metrics

load_percentage
store_percentage
integer_dp_percentage
simd_percentage
scalar_fp_percentage
branch_percentage
crypto_percentage

Metric groups

Operation_Mix

Functional groups

Spec_Operation

0x003B OP_SPEC, Micro-operation speculatively executed, event

Counts micro-operations speculatively executed. This is the count of the number of micro-operations dispatched in a cycle.

Related telemetry artifacts

Metrics

retiring
bad_speculation

Metric groups

Topdown_L1

Functional groups

Spec_Operation

0x0068 UNALIGNED_LD_SPEC, Unaligned access, read, event

Counts unaligned memory read operations issued by the CPU. This event counts unaligned accesses (as defined by the actual instruction), even if they are subsequently issued as multiple aligned accesses. The event does not count preload operations (PLD, PLI).

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x0069 UNALIGNED_ST_SPEC, Unaligned access, write, event

Counts unaligned memory write operations issued by the CPU. This event counts unaligned accesses (as defined by the actual instruction), even if they are subsequently issued as multiple aligned accesses.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x006A UNALIGNED_LDST_SPEC, Unaligned access, event

Counts unaligned memory operations issued by the CPU. This event counts unaligned accesses (as defined by the actual instruction), even if they are subsequently issued as multiple aligned accesses.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x006C LDREX_SPEC, Exclusive operation speculatively executed, Load-Exclusive, event

Counts Load-Exclusive operations that have been speculatively executed. Eg: LDREX, LDX

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x006D STREX_PASS_SPEC, Exclusive operation speculatively executed, Store-Exclusive pass, event

Counts store-exclusive operations that have been speculatively executed and have successfully completed the store operation.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x006E STREX_FAIL_SPEC, Exclusive operation speculatively executed, Store-Exclusive fail, event

Counts store-exclusive operations that have been speculatively executed and have not successfully completed the store operation.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x006F STREX_SPEC, Exclusive operation speculatively executed, Store-Exclusive, event

Counts store-exclusive operations that have been speculatively executed.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x0070 LD_SPEC, Operation speculatively executed, load, event

Counts speculatively executed load operations including Single Instruction Multiple Data (SIMD) load operations.

Related telemetry artifacts

Metrics

load_percentage

Metric groups

Operation_Mix

Functional groups

Spec_Operation

0x0071 ST_SPEC, Operation speculatively executed, store, event

Counts speculatively executed store operations including Single Instruction Multiple Data (SIMD) store operations.

Related telemetry artifacts

Metrics

store_percentage

Metric groups

Operation_Mix

Functional groups

Spec_Operation

0x0073 DP_SPEC, Operation speculatively executed, integer data processing, event

Counts speculatively executed logical or arithmetic instructions such as MOV/MVN operations.

Related telemetry artifacts

Metrics

integer_dp_percentage

Metric groups

Operation_Mix

Functional groups

Spec_Operation

0x0074 ASE_SPEC, Operation speculatively executed, Advanced SIMD, event

Counts speculatively executed Advanced SIMD operations excluding load, store and move micro-operations that move data to or from SIMD (vector) registers.

Related telemetry artifacts

Metrics

simd_percentage

Metric groups

Operation_Mix

Functional groups

Spec_Operation

0x0075 VFP_SPEC, Operation speculatively executed, scalar floating-point, event

Counts speculatively executed floating point operations. This event does not count operations that move data to or from floating point (vector) registers.

Related telemetry artifacts

Metrics

scalar_fp_percentage

Metric groups

Operation_Mix

Functional groups

Spec_Operation

0x0076 PC_WRITE_SPEC, Operation speculatively executed, Software change of the PC, event

Counts speculatively executed operations which cause software changes of the PC. Those operations include all taken branch operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x0077 CRYPTO_SPEC, Operation speculatively executed, Cryptographic instruction, event

Counts speculatively executed cryptographic operations except for PMULL and VMULL operations.

Related telemetry artifacts

Metrics

crypto_percentage

Metric groups

Operation_Mix

Functional groups

Spec_Operation

0x0078 BR_IMMED_SPEC, Branch speculatively executed, immediate branch, event

Counts immediate branch operations which are speculatively executed.

Related telemetry artifacts

Metrics

branch_percentage

Metric groups

Operation_Mix

Functional groups

Spec_Operation

0x0079 BR_RETURN_SPEC, Branch speculatively executed, procedure return, event

Counts procedure return operations (RET) which are speculatively executed.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x007A BR_INDIRECT_SPEC, Branch speculatively executed, indirect branch, event

Counts indirect branch operations including procedure returns, which are speculatively executed. This includes operations that force a software change of the PC, other than exception-generating operations. Eg: BR Xn, RET

Related telemetry artifacts

Metrics

branch_percentage

Metric groups

Operation_Mix

Functional groups

Spec_Operation

0x007C ISB_SPEC, Barrier speculatively executed, ISB, event

Counts ISB operations that are executed.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x007D DSB_SPEC, Barrier speculatively executed, DSB, event

Counts DSB operations that are speculatively issued to Load/Store unit in the CPU.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x007E DMB_SPEC, Barrier speculatively executed, DMB, event

Counts DMB operations that are speculatively issued to the Load/Store unit in the CPU. This event does not count implied barriers from load acquire/store release operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x0090 RC_LD_SPEC, Release consistency operation speculatively executed, Load-Acquire, event

Counts any load acquire operations that are speculatively executed. Eg: LDAR, LDARH, LDARB

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x0091 RC_ST_SPEC, Release consistency operation speculatively executed, Store-Release, event

Counts any store release operations that are speculatively executed. Eg: STLR, STLRH, STLRB’

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

0x8005 ASE_INST_SPEC, Operation speculatively executed, Advanced SIMD, event

Counts speculatively executed Advanced SIMD operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Spec_Operation

7.13 FP_Operation (FP OPERATION) events for Neoverse V2

Speculatively executed floating-point events.

Summary of events in FP_Operation:

- Total implemented Common events: 5
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-13: FP_Operation events summary

Code	Mnemonic	Name	Description
0x8014	FP_HP_SPEC	Floating-point operation speculatively executed, half precision	Counts speculatively executed half precision floating point operations.
0x8018	FP_SP_SPEC	Floating-point operation speculatively executed, single precision	Counts speculatively executed single precision floating point operations.
0x801C	FP_DP_SPEC	Floating-point operation speculatively executed, double precision	Counts speculatively executed double precision floating point operations.
0x80C0	FP_SCALE_OPS_SPEC	Scalable floating-point element ALU operations speculatively executed	Counts speculatively executed scalable single precision floating point operations.
0x80C1	FP_FIXED_OPS_SPEC	Non-scalable floating-point element ALU operations speculatively executed	Counts speculatively executed non-scalable single precision floating point operations.

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x8014 FP_HP_SPEC, Floating-point operation speculatively executed, half precision, event

Counts speculatively executed half precision floating point operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

FP_Operation

0x8018 FP_SP_SPEC, Floating-point operation speculatively executed, single precision, event

Counts speculatively executed single precision floating point operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

FP_Operation

0x801C FP_DP_SPEC, Floating-point operation speculatively executed, double precision, event

Counts speculatively executed double precision floating point operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

FP_Operation

0x80C0 FP_SCALE_OPS_SPEC, Scalable floating-point element ALU operations speculatively executed, event

Counts speculatively executed scalable single precision floating point operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

FP_Operation

0x80C1 FP_FIXED_OPS_SPEC, Non-scalable floating-point element ALU operations speculatively executed, event

Counts speculatively executed non-scalable single precision floating point operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

FP_Operation

7.14 Stall (STALL) events for Neoverse V2

Stall related events.

Summary of events in Stall:

- Total implemented Common events: 7
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-14: Stall events summary

Code	Mnemonic	Name	Description
0x0023	STALL_FRONTEND	No operation sent for execution due to the frontend	Counts cycles when frontend could not send any micro-operations to the rename stage because of...
0x0024	STALL_BACKEND	No operation sent for execution due to the backend	Counts cycles whenever the rename unit is unable to send any micro-operations to the backend of...
0x003C	STALL	No operation sent for execution	Counts cycles when no operations are sent to the rename unit from the frontend or from the rename...
0x003D	STALL_SLOT_BACKEND	No operation sent for execution on a Slot due to the backend	Counts slots per cycle in which no operations are sent from the rename unit to the backend due to...
0x003E	STALL_SLOT_FRONTEND	No operation sent for execution on a Slot due to the frontend	Counts slots per cycle in which no operations are sent to the rename unit from the frontend due...
0x003F	STALL_SLOT	No operation sent for execution on a Slot	Counts slots per cycle in which no operations are sent to the rename unit from the frontend or...
0x4005	STALL_BACKEND_MEM	Memory stall cycles	Counts cycles when the backend is stalled because there is a pending demand load request in...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0023 STALL_FRONTEND, No operation sent for execution due to the frontend, event

Counts cycles when frontend could not send any micro-operations to the rename stage because of frontend resource stalls caused by fetch memory latency or branch prediction flow stalls. All the frontend slots were empty during the cycle when this event counts.

Related telemetry artifacts

Metrics

frontend_stalled_cycles

Metric groups

Cycle_Accounting

Functional groups

Stall

0x0024 STALL_BACKEND, No operation sent for execution due to the backend, event

Counts cycles whenever the rename unit is unable to send any micro-operations to the backend of the pipeline because of backend resource constraints. Backend resource constraints can include issue stage fullness, execution stage fullness, or other internal pipeline resource fullness. All the backend slots were empty during the cycle when this event counts.

Related telemetry artifacts

Metrics

backend_stalled_cycles

Metric groups

Cycle_Accounting

Functional groups

Stall

0x003C STALL, No operation sent for execution, event

Counts cycles when no operations are sent to the rename unit from the frontend or from the rename unit to the backend for any reason (either frontend or backend stall).

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Stall

0x003D STALL_SLOT_BACKEND, No operation sent for execution on a Slot due to the backend, event

Counts slots per cycle in which no operations are sent from the rename unit to the backend due to backend resource constraints.

Related telemetry artifacts

Metrics

backend_bound

Metric groups

Topdown_L1

Functional groups

Stall

0x003E STALL_SLOT_FRONTEND, No operation sent for execution on a Slot due to the frontend, event

Counts slots per cycle in which no operations are sent to the rename unit from the frontend due to frontend resource constraints.

Related telemetry artifacts

Metrics

frontend_bound

Metric groups

Topdown_L1

Functional groups

Stall

0x003F STALL_SLOT, No operation sent for execution on a Slot, event

Counts slots per cycle in which no operations are sent to the rename unit from the frontend or from the rename unit to the backend for any reason (either frontend or backend stall).

Related telemetry artifacts

Metrics

retiring

bad_speculation

Metric groups

Topdown_L1

Functional groups

Stall

0x4005 STALL_BACKEND_MEM, Memory stall cycles, event

Counts cycles when the backend is stalled because there is a pending demand load request in progress in the last level core cache.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

Stall

7.15 General (GENERAL) events for Neoverse V2

General CPU related events.

Summary of events in General:

- Total implemented Common events: 2
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-15: General events summary

Code	Mnemonic	Name	Description
0x0011	CPU_CYCLES	Cycle	Counts CPU clock cycles (not timer cycles). The clock measured by this event is defined as the...

Code	Mnemonic	Name	Description
0x4004	CNT_CYCLES	Constant frequency cycles	Counts constant frequency cycles

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0011 CPU_CYCLES, Cycle, event

Counts CPU clock cycles (not timer cycles). The clock measured by this event is defined as the physical clock driving the CPU logic.

Related telemetry artifacts

Metrics

frontend_stalled_cycles
backend_stalled_cycles
frontend_bound
backend_bound
retiring
bad_speculation
ipc

Metric groups

Topdown_L1
Cycle_Accounting
General

Functional groups

General

0x4004 CNT_CYCLES, Constant frequency cycles, event

Counts constant frequency cycles

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

General

7.16 TLB (TLB) events for Neoverse V2

TLB and MMU related events.

Summary of events in TLB:

- Total implemented Common events: 16
- Total Implemented Product ImpDef events: 0

- PMU Only events : 0
- ETE Only events : 0

Table 7-16: TLB events summary

Code	Mnemonic	Name	Description
0x0002	L1I_TLB_REFILL	Level 1 instruction TLB refill	Counts level 1 instruction TLB refills from any Instruction fetch. If there are multiple misses...
0x0005	L1D_TLB_REFILL	Level 1 data TLB refill	Counts level 1 data TLB accesses that resulted in TLB refills. If there are multiple misses in...
0x0025	L1D_TLB	Level 1 data TLB access	Counts level 1 data TLB accesses caused by any memory load or store operation. Note that load or...
0x0026	L1I_TLB	Level 1 instruction TLB access	Counts level 1 instruction TLB accesses, whether the access hits or misses in the TLB. This event...
0x002D	L2D_TLB_REFILL	Level 2 data TLB refill	Counts level 2 TLB refills caused by memory operations from both data and instruction fetch,...
0x002F	L2D_TLB	Level 2 data TLB access	Counts level 2 TLB accesses except those caused by TLB maintenance operations.
0x0034	DTLB_WALK	Data TLB access with at least one translation table walk	Counts number of demand data translation table walks caused by a miss in the L2 TLB and...
0x0035	ITLB_WALK	Instruction TLB access with at least one translation table walk	Counts number of instruction translation table walks caused by a miss in the L2 TLB and...
0x004C	L1D_TLB_REFILL_RD	Level 1 data TLB refill, read	Counts level 1 data TLB refills caused by memory read operations. If there are multiple misses in...
0x004D	L1D_TLB_REFILL_WR	Level 1 data TLB refill, write	Counts level 1 data TLB refills caused by data side memory write operations. If there are...
0x004E	L1D_TLB_RD	Level 1 data TLB access, read	Counts level 1 data TLB accesses caused by memory read operations. This event counts whether the...
0x004F	L1D_TLB_WR	Level 1 data TLB access, write	Counts any L1 data side TLB accesses caused by memory write operations. This event counts whether...
0x005C	L2D_TLB_REFILL_RD	Level 2 data TLB refill, read	Counts level 2 TLB refills caused by memory read operations from both data and instruction fetch...
0x005D	L2D_TLB_REFILL_WR	Level 2 data TLB refill, write	Counts level 2 TLB refills caused by memory write operations from both data and instruction fetch...
0x005E	L2D_TLB_RD	Level 2 data TLB access, read	Counts level 2 TLB accesses caused by memory read operations from both data and instruction fetch...
0x005F	L2D_TLB_WR	Level 2 data TLB access, write	Counts level 2 TLB accesses caused by memory write operations from both data and instruction...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x0002 L1I_TLB_REFILL, Level 1 instruction TLB refill, event

Counts level 1 instruction TLB refills from any Instruction fetch. If there are multiple misses in the TLB that are resolved by the refill, then this event only counts once. This event will not count if the translation table walk results in a fault (such as a translation or access fault), since there is no new translation created for the TLB.

Related telemetry artifacts

Metrics

- l1i_tlb_mpki in ITLB_Effectiveness
- l1i_tlb_mpki in MPKI
- l1i_tlb_miss_ratio in ITLB_Effectiveness
- l1i_tlb_miss_ratio in Miss_Ratio

Metric groups

- Miss_Ratio
- MPKI
- ITLB_Effectiveness

Functional groups

- TLB

0x0005 L1D_TLB_REFILL, Level 1 data TLB refill, event

Counts level 1 data TLB accesses that resulted in TLB refills. If there are multiple misses in the TLB that are resolved by the refill, then this event only counts once. This event counts for refills caused by preload instructions or hardware prefetch accesses. This event counts regardless of whether the miss hits in L2 or results in a translation table walk. This event will not count if the translation table walk results in a fault (such as a translation or access fault), since there is no new translation created for the TLB. This event will not count on an access from an AT(address translation) instruction.

Related telemetry artifacts

Metrics

- l1d_tlb_mpki in DTLB_Effectiveness
- l1d_tlb_mpki in MPKI
- l1d_tlb_miss_ratio in DTLB_Effectiveness
- l1d_tlb_miss_ratio in Miss_Ratio

Metric groups

- DTLB_Effectiveness
- MPKI
- Miss_Ratio

Functional groups

- TLB

0x0025 L1D_TLB, Level 1 data TLB access, event

Counts level 1 data TLB accesses caused by any memory load or store operation. Note that load or store instructions can be broken up into multiple memory operations. This event does not count TLB maintenance operations.

Related telemetry artifacts

Metrics

- dtlb_walk_ratio in DTLB_Effectiveness

dtlb_walk_ratio in Miss_Ratio
l1d_tlb_miss_ratio in DTLB_Effectiveness
l1d_tlb_miss_ratio in Miss_Ratio

Metric groups

DTLB_Effectiveness
Miss_Ratio

Functional groups

TLB

0x0026 L1I_TLB, Level 1 instruction TLB access, event

Counts level 1 instruction TLB accesses, whether the access hits or misses in the TLB. This event counts both demand accesses and prefetch or preload generated accesses.

Related telemetry artifacts

Metrics

itlb_walk_ratio in ITLB_Effectiveness
itlb_walk_ratio in Miss_Ratio
l1i_tlb_miss_ratio in ITLB_Effectiveness
l1i_tlb_miss_ratio in Miss_Ratio

Metric groups

ITLB_Effectiveness
Miss_Ratio

Functional groups

TLB

0x002D L2D_TLB_REFILL, Level 2 data TLB refill, event

Counts level 2 TLB refills caused by memory operations from both data and instruction fetch, except for those caused by TLB maintenance operations and hardware prefetches.

Related telemetry artifacts

Metrics

l2_tlb_mpki in ITLB_Effectiveness
l2_tlb_mpki in DTLB_Effectiveness
l2_tlb_mpki in MPKI
l2_tlb_miss_ratio in ITLB_Effectiveness
l2_tlb_miss_ratio in DTLB_Effectiveness
l2_tlb_miss_ratio in Miss_Ratio

Metric groups

MPKI
ITLB_Effectiveness
DTLB_Effectiveness
Miss_Ratio

Functional groups

TLB

0x002F L2D_TLB, Level 2 data TLB access, event

Counts level 2 TLB accesses except those caused by TLB maintenance operations.

Related telemetry artifacts

Metrics

l2_tlb_miss_ratio in ITLB_Effectiveness
l2_tlb_miss_ratio in DTLB_Effectiveness
l2_tlb_miss_ratio in Miss_Ratio

Metric groups

Miss_Ratio
DTLB_Effectiveness
ITLB_Effectiveness

Functional groups

TLB

0x0034 DTLB_WALK, Data TLB access with at least one translation table walk, event

Counts number of demand data translation table walks caused by a miss in the L2 TLB and performing at least one memory access. Translation table walks are counted even if the translation ended up taking a translation fault for reasons different than EPD, E0PD and NFD. Note that partial translations that cause a translation table walk are also counted. Also note that this event counts walks triggered by software preloads, but not walks triggered by hardware prefetchers, and that this event does not count walks triggered by TLB maintenance operations.

Related telemetry artifacts

Metrics

dtlb_mpki in DTLB_Effectiveness
dtlb_mpki in MPKI
dtlb_walk_ratio in DTLB_Effectiveness
dtlb_walk_ratio in Miss_Ratio

Metric groups

DTLB_Effectiveness
MPKI
Miss_Ratio

Functional groups

TLB

0x0035 ITLB_WALK, Instruction TLB access with at least one translation table walk, event

Counts number of instruction translation table walks caused by a miss in the L2 TLB and performing at least one memory access. Translation table walks are counted even if the translation

ended up taking a translation fault for reasons different than EPD, E0PD and NFD. Note that partial translations that cause a translation table walk are also counted. Also note that this event does not count walks triggered by TLB maintenance operations.

Related telemetry artifacts

Metrics

itlb_mpki in ITLB_Effectiveness
itlb_mpki in MPKI
itlb_walk_ratio in ITLB_Effectiveness
itlb_walk_ratio in Miss_Ratio

Metric groups

Miss_Ratio
MPKI
ITLB_Effectiveness

Functional groups

TLB

0x004C L1D_TLB_REFILL_RD, Level 1 data TLB refill, read, event

Counts level 1 data TLB refills caused by memory read operations. If there are multiple misses in the TLB that are resolved by the refill, then this event only counts once. This event counts for refills caused by preload instructions or hardware prefetch accesses. This event counts regardless of whether the miss hits in L2 or results in a translation table walk. This event will not count if the translation table walk results in a fault (such as a translation or access fault), since there is no new translation created for the TLB. This event will not count on an access from an Address Translation (AT) instruction.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TLB

0x004D L1D_TLB_REFILL_WR, Level 1 data TLB refill, write, event

Counts level 1 data TLB refills caused by data side memory write operations. If there are multiple misses in the TLB that are resolved by the refill, then this event only counts once. This event counts for refills caused by preload instructions or hardware prefetch accesses. This event counts regardless of whether the miss hits in L2 or results in a translation table walk. This event will not count if the table walk results in a fault (such as a translation or access fault), since there is no new translation created for the TLB. This event will not count with an access from an Address Translation (AT) instruction.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TLB

0x004E L1D_TLB_RD, Level 1 data TLB access, read, event

Counts level 1 data TLB accesses caused by memory read operations. This event counts whether the access hits or misses in the TLB. This event does not count TLB maintenance operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TLB

0x004F L1D_TLB_WR, Level 1 data TLB access, write, event

Counts any L1 data side TLB accesses caused by memory write operations. This event counts whether the access hits or misses in the TLB. This event does not count TLB maintenance operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TLB

0x005C L2D_TLB_REFILL_RD, Level 2 data TLB refill, read, event

Counts level 2 TLB refills caused by memory read operations from both data and instruction fetch except for those caused by TLB maintenance operations or hardware prefetches.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TLB

0x005D L2D_TLB_REFILL_WR, Level 2 data TLB refill, write, event

Counts level 2 TLB refills caused by memory write operations from both data and instruction fetch except for those caused by TLB maintenance operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TLB

0x005E L2D_TLB_RD, Level 2 data TLB access, read, event

Counts level 2 TLB accesses caused by memory read operations from both data and instruction fetch except for those caused by TLB maintenance operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TLB

0x005F L2D_TLB_WR, Level 2 data TLB access, write, event

Counts level 2 TLB accesses caused by memory write operations from both data and instruction fetch except for those caused by TLB maintenance operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TLB

7.17 SVE (SVE) events for Neoverse V2

SVE related events.

Summary of events in SVE:

- Total implemented Common events: 12
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-17: SVE events summary

Code	Mnemonic	Name	Description
0x8006	SVE_INST_SPEC	Operation speculatively executed, SVE, including load and store	Counts speculatively executed operations that are SVE operations.
0x8074	SVE_PRED_SPEC	Operation speculatively executed, SVE predicated	Counts speculatively executed predicated SVE operations.
0x8075	SVE_PRED_EMPTY_SPEC	Operation speculatively executed, SVE predicated with no active predicates	Counts speculatively executed predicated SVE operations with no active predicate elements.
0x8076	SVE_PRED_FULL_SPEC	Operation speculatively executed, SVE predicated with all active predicates	Counts speculatively executed predicated SVE operations with all predicate elements active.
0x8077	SVE_PRED_PARTIAL_SPEC	Operation speculatively executed, SVE predicated with partially active predicates	Counts speculatively executed predicated SVE operations with at least one but not all active...

Code	Mnemonic	Name	Description
0x8079	SVE_PRED_NOT_FULL_SPEC	SVE predicated operations speculatively executed with no active or partially active predicates	Counts speculatively executed predicated SVE operations with at least one non active predicate...
0x80BC	SVE_LDFF_SPEC	Operation speculatively executed, SVE first-fault load	Counts speculatively executed SVE first fault or non-fault load operations.
0x80BD	SVE_LDFF_FAULT_SPEC	Operation speculatively executed, SVE first-fault load which set FFR bit to 0b0	Counts speculatively executed SVE first fault or non-fault load operations that clear at least...
0x80E3	ASE_SVE_INT8_SPEC	Integer operation speculatively executed, Advanced SIMD or SVE 8-bit	Counts speculatively executed Advanced SIMD or SVE integer operations with the largest data type...
0x80E7	ASE_SVE_INT16_SPEC	Integer operation speculatively executed, Advanced SIMD or SVE 16-bit	Counts speculatively executed Advanced SIMD or SVE integer operations with the largest data type...
0x80EB	ASE_SVE_INT32_SPEC	Integer operation speculatively executed, Advanced SIMD or SVE 32-bit	Counts speculatively executed Advanced SIMD or SVE integer operations with the largest data type...
0x80EF	ASE_SVE_INT64_SPEC	Integer operation speculatively executed, Advanced SIMD or SVE 64-bit	Counts speculatively executed Advanced SIMD or SVE integer operations with the largest data type...

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x8006 SVE_INST_SPEC, Operation speculatively executed, SVE, including load and store, event

Counts speculatively executed operations that are SVE operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x8074 SVE_PRED_SPEC, Operation speculatively executed, SVE predicated, event

Counts speculatively executed predicated SVE operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x8075 SVE_PRED_EMPTY_SPEC, Operation speculatively executed, SVE predicated with no active predicates, event

Counts speculatively executed predicated SVE operations with no active predicate elements.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x8076 SVE_PRED_FULL_SPEC, Operation speculatively executed, SVE predicated with all active predicates, event

Counts speculatively executed predicated SVE operations with all predicate elements active.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x8077 SVE_PRED_PARTIAL_SPEC, Operation speculatively executed, SVE predicated with partially active predicates, event

Counts speculatively executed predicated SVE operations with at least one but not all active predicate elements.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x8079 SVE_PRED_NOT_FULL_SPEC, SVE predicated operations speculatively executed with no active or partially active predicates, event

Counts speculatively executed predicated SVE operations with at least one non active predicate elements.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x80BC SVE_LDFF_SPEC, Operation speculatively executed, SVE first-fault load, event

Counts speculatively executed SVE first fault or non-fault load operations.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x80BD SVE_LDFF_FAULT_SPEC, Operation speculatively executed, SVE first-fault load which set FFR bit to 0b0, event

Counts speculatively executed SVE first fault or non-fault load operations that clear at least one bit in the FFR.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x80E3 ASE_SVE_INT8_SPEC, Integer operation speculatively executed, Advanced SIMD or SVE 8-bit, event

Counts speculatively executed Advanced SIMD or SVE integer operations with the largest data type an 8-bit integer.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x80E7 ASE_SVE_INT16_SPEC, Integer operation speculatively executed, Advanced SIMD or SVE 16-bit, event

Counts speculatively executed Advanced SIMD or SVE integer operations with the largest data type a 16-bit integer.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x80EB ASE_SVE_INT32_SPEC, Integer operation speculatively executed, Advanced SIMD or SVE 32-bit, event

Counts speculatively executed Advanced SIMD or SVE integer operations with the largest data type a 32-bit integer.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

0x80EF ASE_SVE_INT64_SPEC, Integer operation speculatively executed, Advanced SIMD or SVE 64-bit, event

Counts speculatively executed Advanced SIMD or SVE integer operations with the largest data type a 64-bit integer.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

SVE

7.18 TRACE (TRACE) events for Neoverse V2

Trace related events.

Summary of events in TRACE:

- Total implemented Common events: 9
- Total Implemented Product ImpDef events: 0
- PMU Only events : 0
- ETE Only events : 0

Table 7-18: TRACE events summary

Code	Mnemonic	Name	Description
0x400C	TRB_WRAP	Trace buffer current write pointer wrapped	This event is generated each time the current write pointer is wrapped to the base pointer.
0x4010	TRCEXTOUT0	Trace unit external output 0	This event is generated each time an event is signaled by ETE external event 0.
0x4011	TRCEXTOUT1	Trace unit external output 1	This event is generated each time an event is signaled by ETE external event 1.
0x4012	TRCEXTOUT2	Trace unit external output 2	This event is generated each time an event is signaled by ETE external event 2.
0x4013	TRCEXTOUT3	Trace unit external output 3	This event is generated each time an event is signaled by ETE external event 3.
0x4018	CTI_TRIGOUT4	Cross-trigger Interface output trigger 4	This event is generated each time an event is signaled on CTI output trigger 4.
0x4019	CTI_TRIGOUT5	Cross-trigger Interface output trigger 5	This event is generated each time an event is signaled on CTI output trigger 5.
0x401A	CTI_TRIGOUT6	Cross-trigger Interface output trigger 6	This event is generated each time an event is signaled on CTI output trigger 6.
0x401B	CTI_TRIGOUT7	Cross-trigger Interface output trigger 7	This event is generated each time an event is signaled on CTI output trigger 7.

For a complete list of the events in Neoverse V2, see PMU events cheat sheet for Neoverse V2 and PMU events lookup table for Neoverse V2.

0x400C TRB_WRAP, Trace buffer current write pointer wrapped, event

This event is generated each time the current write pointer is wrapped to the base pointer.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TRACE

0x4010 TRCEXTOUT0, Trace unit external output 0, event

This event is generated each time an event is signaled by ETE external event 0.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TRACE

0x4011 TRCEXTOUT1, Trace unit external output 1, event

This event is generated each time an event is signaled by ETE external event 1.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TRACE

0x4012 TRCEXTOUT2, Trace unit external output 2, event

This event is generated each time an event is signaled by ETE external event 2.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TRACE

0x4013 TRCEXTOUT3, Trace unit external output 3, event

This event is generated each time an event is signaled by ETE external event 3.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TRACE

`0x4018` CTI_TRIGOUT4, Cross-trigger Interface output trigger 4, event

This event is generated each time an event is signaled on CTI output trigger 4.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TRACE

`0x4019` CTI_TRIGOUT5, Cross-trigger Interface output trigger 5, event

This event is generated each time an event is signaled on CTI output trigger 5.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TRACE

`0x401A` CTI_TRIGOUT6, Cross-trigger Interface output trigger 6, event

This event is generated each time an event is signaled on CTI output trigger 6.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TRACE

`0x401B` CTI_TRIGOUT7, Cross-trigger Interface output trigger 7, event

This event is generated each time an event is signaled on CTI output trigger 7.

Related telemetry artifacts

There are no related metrics or metric groups because this event is not used in the Neoverse V2 Methodology Specification.

Functional groups

TRACE

8. Supplemental performance debug PMU events

The Neoverse V2 core implements an additional set of events that are available for use in debugging the performance behaviors of the processor. These events are not guaranteed to have the same level of accuracy as the architected PMU counters. Any use of these events must take this variable accuracy into account.

The following table shows the events.

Table 8-1: Supplemental performance debug PMU events forthe Neoverse V2 core

Code	Mnemonic	Description
0x0E1	IMP_STALL_FRONTEND_MEM	No operation issued due to the frontend, cache miss
0x0E2	IMP_STALL_FRONTEND_TLB	No operation issued due to the frontend, TLB miss
0x108	IMP_L2_CACHE_IF_REFILL	Level 2 cache refill, fetch
0x10B	IMP_L2_CACHE_PF_LATE_REFILL	Level 2 prefetch requests, late
0x120	IMP_CT_FLUSH	Flushes including architectural, microarchitectural, and branch redirects
0x121	IMP_CT_FLUSH_MEM	Flushes due to memory hazards
0x122	IMP_CT_FLUSH_BAD_BRANCH	Flushes due to non-branch instruction predicted as a branch
0x123	IMP_CT_FLUSH_PREDECODE_ERR	Flushes due to bad predecode
0x124	IMP_CT_FLUSH_ISB	Flushes due to ISB or similar side-effects
0x125	IMP_CT_FLUSH_OTHER	Flushes due to other hazards
0x127	IMP_LS_RAR	Loadstore detected nuke due to read-after-read ordering hazard
0x128	IMP_LS_RAW	Loadstore detected nuke due to read-after-write ordering hazard
0x15B	IMP_STALL_FRONTED_FLUSH	No operation sent for execution due to the frontend flush recovery
0x158	IMP_STALL_BACKEND_RENAME_FRF	RN dispatch stall due to flag registers
0x159	IMP_STALL_BACKEND_RENAME_GRF	RN dispatch stall due to general registers
0x15A	IMP_STALL_BACKEND_RENAME_VRF	RN dispatch stall due to vector registers
0x15C	IMP_STALL_BACKEND_IQ_SX	RN dispatch stall due to SX IQ entries
0x15D	IMP_STALL_BACKEND_IQ_MX	RN dispatch stall due to MX IQ entries
0x15E	IMP_STALL_BACKEND_IQ_LS	RN dispatch stall due to LS IQ entries
0x15F	IMP_STALL_BACKEND_IQ_VX	RN dispatch stall due to VX IQ entries
0x160	IMP_STALL_BACKEND_MCQ	RN dispatch stall due to MCQ full
0x17B	IMP_NEAR_CAS	Near atomics: compare and swap
0x17C	IMP_NEAR_CAS_PAS	Near atomics: compare and swap pass