



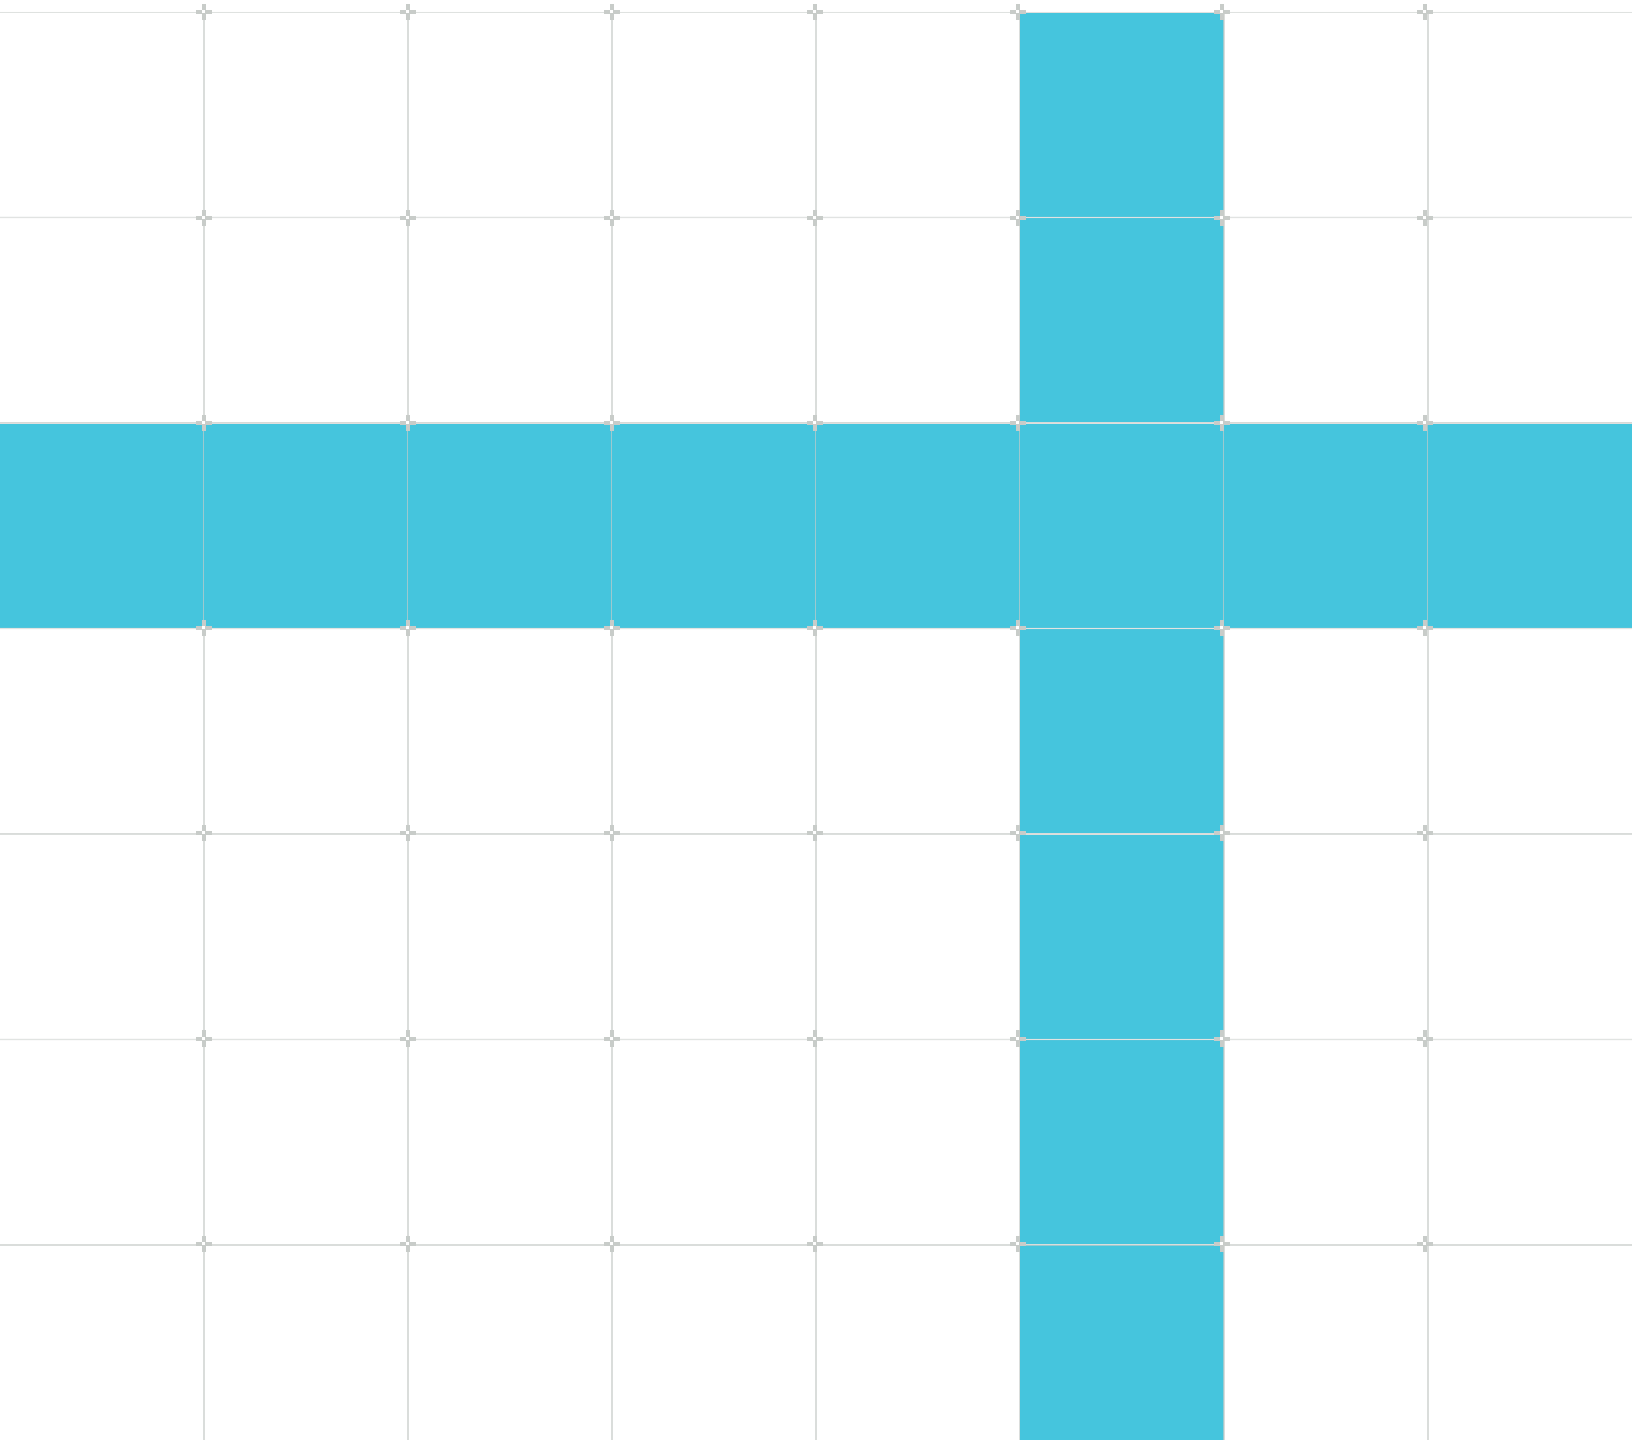
Arm[®] CPU Telemetry Solution

Topdown Methodology Specification

Non-Confidential

Issue 02

Copyright © 2024–2025 Arm Limited (or its affiliates). 109542_02_02_en
All rights reserved.



Arm® CPU Telemetry Solution Topdown Methodology Specification

This document is Non-Confidential.

Copyright © 2024–2025 Arm Limited (or its affiliates). All rights reserved.

This document is protected by copyright and other intellectual property rights.

Arm only permits use of this document if you have reviewed and accepted [Arm's Proprietary Notice](#) found at the end of this document.

This document (109542_02_02_en) was issued on 2025-09-10. There might be a later issue at <https://developer.arm.com/documentation/109542>

See also: [Proprietary notice](#) | [Product and document information](#) | [Useful resources](#)

Start reading

If you prefer, you can skip to [the start of the content](#).

Intended audience

This specification is useful for engineers to collect and analyze CPU telemetry data to gain insights about a system's performance. Architects and system designers can also use it for resource characterization and platform tuning.

Inclusive language commitment

Arm values inclusive communities. Arm recognizes that we and our industry have used language that can be offensive. Arm strives to lead the industry and create change.

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

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>.

Contents

1. About Arm CPU Telemetry Solution.....4

2. Arm Topdown methodology..... 5

2.1 Stage 1: Topdown analysis..... 6

2.2 Stage 2: Microarchitecture analysis.....11

3. Arm Telemetry framework for CPUs.....14

3.1 Data model standardization..... 18

4. Arm Telemetry Solution for software profiling: tools.....19

4.1 Arm telemetry specification and profiling tools..... 19

4.1.1 Arm Topdown tool.....20

4.1.2 Performance analysis using Linux perf tool.....21

4.1.3 Performance analysis using WindowsPerf Tool..... 23

A. Arm Topdown tool example.....27

B. Linux perf data collection..... 29

Proprietary notice.....31

Product and document information.....33

Product status..... 33

Revision history.....33

Conventions..... 34

Useful resources.....37

1. About Arm CPU Telemetry Solution

Arm CPU Telemetry Solution comprises a set of components, that is, a top-down performance analysis methodology, a standardized telemetry framework, and profiling tool. It is designed to use a CPU's telemetry data to help identify performance bottlenecks and improve execution efficiency by using these components.

System-on-Chip (SoC) telemetry enables collection and analysis of hardware execution information from a platform to gain insights about a system's performance. This information can be used to identify performance issues and improve execution efficiency.

A modern CPU contains a hardware performance monitoring unit (PMU) which provides a set of functional events that describe the internal state of the microarchitecture during execution. Multiple events can be used to derive metrics, which provide more meaningful insights by abstracting the hardware details.

Performance analysis is an investigative and diagnostic process. Using a systematic methodology is important to narrow down the bottleneck for root cause analysis and code execution improvements. A well-known example of a systematic methodology is the top-down performance analysis methodology. For more information, see [A. Yasin, IEEE Xplore, "A Top-Down method for performance analysis and counters architecture"](#).

Arm CPU Telemetry Solution is designed to collect, represent, and analyze CPU telemetry data on Arm CPU-based platforms. A supported Arm CPU implementation provides a telemetry specification that defines the hardware PMU events, derived metrics, and Arm Topdown methodology supported by the CPU. Arm Topdown methodology is Arm's implementation of the top-down performance analysis methodology. It is a approach to consume the telemetry events and metrics in a hierarchical decision tree format for hotspot analysis. See [Arm Topdown methodology](#).

The methodology, metrics, and events are represented in a standardized telemetry framework, Arm Telemetry framework that is designed to support the collection and processing of large amounts of CPU telemetry data. In this framework, events are categorized into function groups and metric groups that can be collected in stages, with the help of profiling tools. This framework also helps to manage and represent the telemetry specification of supported CPUs in a standardized machine-readable format, which can be harnessed by profiling tools.

Arm CPU Telemetry Solution also provides the Arm Topdown tool. It is a simple command line tool to profile applications. The tool parses the telemetry machine-readable specification (MRS), which is a JSON file, to collect and process the telemetry data that is supported by the CPU to provide performance insights. Arm Topdown tool is enabled for both Linux and Windows platforms.

For more information about Arm CPU Telemetry Solution, see [Arm® Telemetry on Arm Developer](#).

2. Arm Topdown methodology

Arm Topdown methodology supports performance analysis, workload characterization, and microarchitecture analysis on Arm CPUs that implement version 3 or later of the Performance Monitors Extension, `FEAT_PMU`.

For more information about the Performance Monitors Extension, see [Arm® Architecture Reference Manual for A-profile architecture](#).

This performance analysis methodology provides a systematic and top-down hierarchical approach to use hardware performance monitoring events for analysis use cases. The methodology is formulated with the help of performance metrics derived from the hardware monitoring events, making the hardware monitors accessible for an average software user without extensive microarchitectural knowledge. Computer architects and system designers can also use it for resource characterization and platform tuning by using the specified set of performance metrics.

To identify the application bottleneck for a CPU, the first step is to characterize the execution efficiency of the cycles spent. This workload characterization can be measured by the distribution of execution cycles broken down between:

- Those cycles that worked efficiently.
- Those cycles that are wasted by pipeline stalls and redirections.

The process of breaking down the cycles can be done hierarchically to narrow down the bottleneck in the CPU pipeline. It also helps to identify and ignore the CPU components that are not causing the performance issue.

After identifying the CPU component bottleneck in the hardware, the next step is to measure the microarchitectural metrics of that component for further root cause analysis.

To aid the investigative and diagnosis process, Arm Topdown methodology is conducted in two stages:

Stage 1: Topdown analysis

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: Microarchitecture analysis

The second stage is to conduct Microarchitecture Analysis to further analyze bottlenecked CPU resources. It uses a set of CPU resource effectiveness metrics. For more information, see [Stage 2: Microarchitecture analysis](#).

After completing stage 1, the Arm Topdown methodology provides recommended stage 2 metrics to further analyze the identified bottleneck. Stage 2 metrics can be used directly for a targeted analysis of a CPU resource.


2.1 Stage 1: Topdown analysis

Topdown analysis is the first stage of the methodology to support hotspot detection. A set of pipeline efficiency metrics specify the PMU events to measure, which helps to characterize the distribution of cycles spent by the CPU.

Topdown analysis metrics are formulated as a decision tree of metrics that need to be hierarchically traversed, in a predefined path, to help locate the bottleneck.

An Arm CPU implementation consists of a frontend that fetches instructions from memory, decodes them, and dispatches them to a backend. The backend executes those instructions, including reading and writing data from and to memory. As shown in the pipeline model in the following figure, the frontend is often in order, whereas the backend can be in order or out of order depending on the CPU's microarchitecture.

Figure 2-1: An Arm CPU implementation pipeline model



In a simple scalar CPU implementation, the frontend dispatches a maximum of one instruction for each cycle. In this case, a cycle is classified as follows:

- Useful cycle, whereby an instruction were dispatched
- Stalled cycle, whereby no instruction was dispatched

In a superscalar CPU implementation, the frontend can dispatch many instructions for each cycle using instruction level parallelism, thereby achieving a higher Instructions Per Cycle (IPC). Each instruction takes one of an implementation-defined maximum number of dispatch slots on each cycle. Therefore, the utilization and efficiency of cycles spent is more accurately calculated by counting the number of dispatch slots in which instructions were dispatched and those slots that did not dispatch instructions.

A modern CPU implementation can also break instructions into micro-operations for execution. The frontend of the CPU decodes and decomposes instructions to micro-operations that can be executed by the backend execution units. In this case, the utilization and efficiency of a pipeline is measured by examining cycles or slots on which no operations are dispatched, referred to as stalled cycles or stalled slots.

To keep the CPU backend fed with instructions, the frontend also predicts the future instructions to be executed. The success of future instructions is highly dependent on the control flow of the code that is determined by branched instructions. When these predictions are wrong, the instructions dispatched to the backend are canceled and can cause pipeline bubbles.

On superscalar CPU implementations, the total execution bandwidth of the CPU can be measured in terms of the number of execution slots for operations. The total number of slots supported by the core determines the execution bandwidth of the CPU for top-down accounting and is defined as a microarchitectural parameter. The value of the parameter is used to derive the execution bandwidth metrics for a CPU.

These key characteristics determine efficient pipeline execution. Thus, the Arm Topdown analysis starts by calculating the following measurements to detect inefficiencies. Each measurement is a percentage of the total execution bandwidth of the CPU:

- The percentage of execution bandwidth used by operations that are retired.
- The percentage of execution bandwidth lost to mis-speculation.
- The percentage of execution bandwidth lost to stalls in the frontend.
- The percentage of execution bandwidth lost to stalls in the backend.

These measurements that form the metrics for the Topdown level 1 analysis are defined as follows:

retiring

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

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.

frontend_bound


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

backend_bound

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

The following figure shows an example of Topdown level 1 analysis that was conducted for a microbenchmark designed to stress key CPU blocks.

Figure 2-2: Ustress Workloads Topdown Level 1 metrics (TopdownL1)



After the first level of accounting, Arm Topdown methodology can then support lower levels in the hierarchy, to further break down the stalling events.

A relatively high `frontend_bound` metric shows that execution cycles are being wasted due to pipeline stalls in the in-order frontend of the CPU. Many causes can exist for these stalls, such as:

- Inefficiency in the branch prediction unit
- Fetch latency due to instruction cache misses
- Translation delays caused by instruction Translation Lookaside Buffer (TLB) walks

For example, if the frontend is stalled and there is an instruction cache miss in progress, then the stall might be due to the cache miss and hence attributed to it.

A relatively high `backend_bound` metric shows that execution cycles are wasted due to pipeline stalls in the backend of the CPU. Many causes can exist for these stalls such as:

- Inefficiency in the backend execution units
- Data cache misses
- Translation delays caused by data TLB walks

A relatively high `bad_speculation` metric shows that the pipeline stalls can be due pipeline flushes or machine clears that break the pipeline requiring a control flow change. The major causes for these stalls are branch mis-predictions and exceptions.

A relatively high `retiring` metric shows that the pipelines were utilized. However, this metric can indicate inefficiency due to underutilization of the microarchitectural capabilities. For example, scalar execution of a code that should have been performed more efficiently using vector operations.

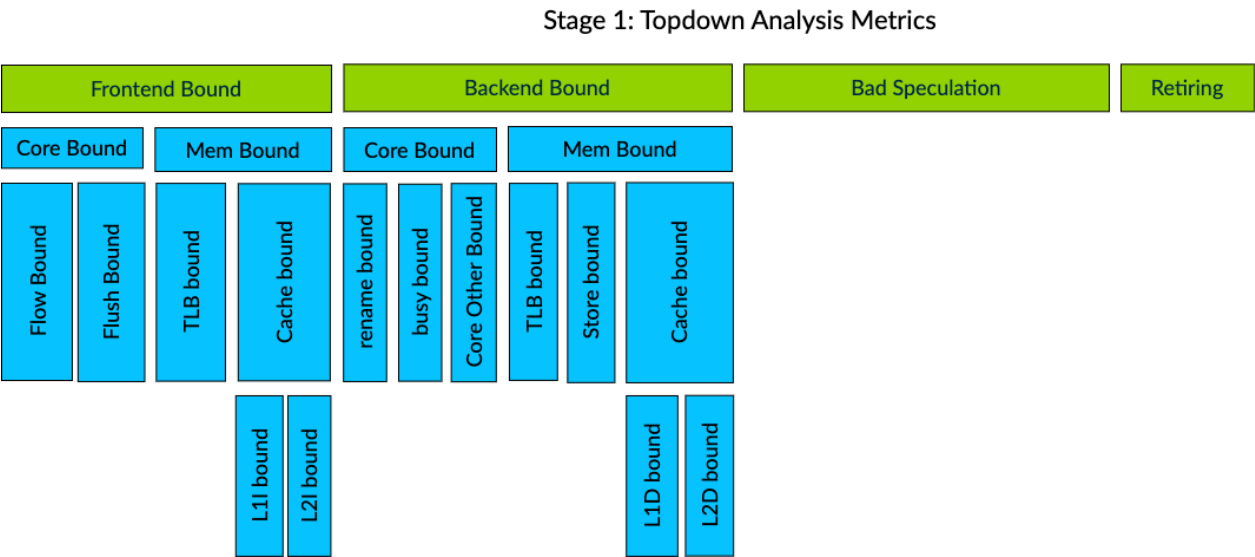
After the level 1 metrics, second level analysis mainly determines whether a pipeline stall is due to memory or processor effects, for example:

- An example of a memory effect is a cache miss. If memory effects dominate, then the program is referred to as memory bound.
- An example of a processor effect is when a backend is full of operations waiting for execution units or results of the previous operations to become available. If processor effects dominate, then the program is referred to as CPU bound.

These stalls can also be broken down into third and fourth-level metrics, such as the specific cases outlined in the earlier memory and processor effects examples. Support for these deeper levels may depend on the underlying microarchitecture.

The following figure shows a Topdown methodology for an Arm CPU that provides support for four levels of metrics to enable more detailed analysis.

Figure 2-3: Arm CPU Topdown Methodology: Four Levels



Refer to Topdown analysis tree of the Arm CPU implementation to understand what the CPU supports, for example, in an Arm core telemetry specification. For more information about available Arm core telemetry specifications, see [Arm® Telemetry on Arm Developer](#).

Topdown event implementations for Arm CPU microarchitecture

It is not always practical or feasible to continue down the levels of top-down accounting, especially to track slot level. The lower the level, the higher the number of events that occur simultaneously.

The lower the level, the smaller it becomes, and any signal causing the bottleneck can get lost in the noise. Thus, to precisely attribute the causes to effects becomes more difficult. Tracking cause to effect can involve a lot of hardware which is:


- Difficult to design
- Difficult to validate
- Consumes area and power that can be otherwise spent on additional processing resources

As a result, the accounting of events takes the form of: the counter counts each cycle by `STALL_<other event>` where `<a condition applies>`. That is, these events count the coincidence of a stall with the applied condition.

Although described as being coincidental, CPU implementations might choose to delay some event signals while deciding whether the event occurs. Because of pipelining, a stall that results from a condition might not be immediately possible once the condition becomes active. To improve the quality of the event data, the condition might be delayed to determine whether it is coincidental with a stall. That is, `<a condition applies>` becomes `<a condition applied after a fixed, implementation-specific, number of cycles>`.

For the purpose of stall accounting in Arm A-profile “big” CPUs, the boundary between frontend and backend is at the point of rename and dispatch as shown in the following figure. It is the point in the pipeline where the core switches from an in-order pipeline to an out-of-order pipeline. It also marks the point where resource contention that causes a stall condition is associated with the throughput in the out-of-order division of the Arm A-profile “big” CPU. Examples of resource contention are execution bandwidth and physical register renames.

Figure 2-4: Frontend and backend division in an Arm “big” CPU with an out-of-order backend



Arm architecture defines some key Topdown events that support the Topdown metrics. The following table show some example events used in the Topdown level 1 metric definitions of Arm processor IP that have adopted Arm CPU Telemetry Solution, referred to as Arm cores in this specification. Additional levels of Topdown analysis are present in Arm processor IP.

Table 2-1: Topdown analysis event support in Arm cores

Topdown event category	Stall events	Supported Arm cores
Overall Stalls	STALL	Neoverse™ N2, N3, V1, V2, V3, Arm® C1-Ultra, Arm C1-Premium, Arm C1-Pro, and Arm C1-Nano

Topdown event category	Stall events	Supported Arm cores
-	STALL_SLOT	Neoverse N2, N3, V1, V2, V3, C1-Ultra, C1-Premium, C1-Pro, and C1-Nano
Frontend Stall Events	STALL_FRONTEND	Neoverse N1, N2, N3, V1, V2, V3, C1-Ultra, C1-Premium, C1-Pro, and C1-Nano
-	STALL_SLOT_FRONTEND	Neoverse N2, N3, V1, V2, V3, C1-Ultra, C1-Premium, C1-Pro, and C1-Nano
Backend Stall Events	STALL_BACKEND	Neoverse N1, N2, N3, V1, V2, V3, C1-Ultra, C1-Premium, C1-Pro, and C1-Nano
	STALL_SLOT_BACKEND	Neoverse N2, N3, V1, V2, V3, C1-Ultra, C1-Premium, C1-Pro, and C1-Nano



Note

Refer to the Topdown methodology tree in each supported Arm core for the complete list of events, hierarchy of metrics derived from these events, and event relationships. A Topdown methodology may include common architectural events across Arm processor families and a set of events which can be microarchitecture dependent. For more information about available Arm core telemetry specifications, see [Arm® Telemetry on Arm Developer](#).

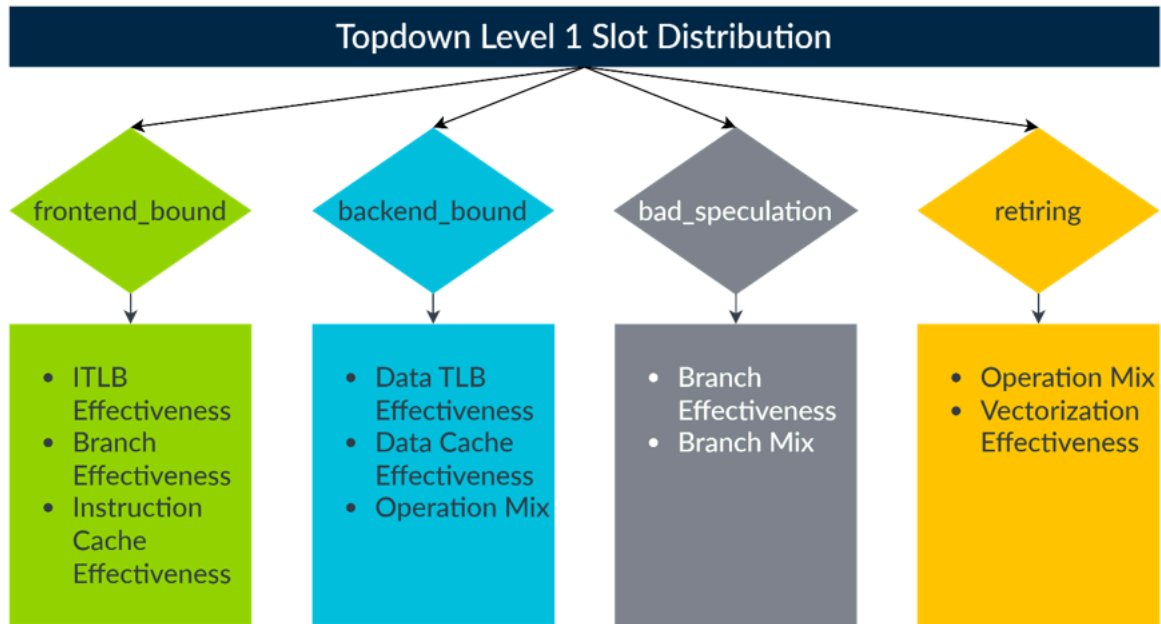
2.2 Stage 2: Microarchitecture analysis

After the potential hotspot in the CPU pipeline is identified in stage 1, the next stage is to conduct a microarchitectural analysis of the CPU pipeline resource causing the bottleneck.

Stage 2 is defined as the Microarchitecture Analysis stage for which a set of CPU resource effectiveness metrics are defined under metric groups for each resource. Industry-standard metrics such as Misses Per Kilo Instructions (MPKI) and Miss Ratios are metric groups that are also defined in this stage.

A relatively high frontend stall rate indicates that cycles are wasted due to pipeline stalls in the in-order frontend of the CPU. A relatively high backend stall rate indicates that cycles are wasted due to pipeline stalls in the backend. This breakdown helps to narrow down the dominating CPU blocks that can be further analyzed to identify performance bottlenecks as shown in the following figure.

Figure 2-5: Stage 2 metrics for Microarchitecture Analysis



The list of major CPU resource effectiveness metrics to further analyze a `frontend_bound` workload are as follows:

- Instruction TLB (ITLB) Effectiveness metrics
- Instruction Cache Effectiveness metrics
- Branch Effectiveness metrics

The list of major CPU resource effectiveness metrics to further analyze a `backend_bound` workload is as follows:

- Data TLB Effectiveness metrics
- Data Cache Effectiveness metrics
- Operation Mix metrics

The list of major CPU resource effectiveness metrics to further analyze a `bad_speculation` workload is as follows:

- Branch Effectiveness metrics
- Branch Mix metrics

The list of major CPU resource effectiveness metrics to further analyze a `retiring` workload is as follows:

- Operation Mix metrics
- Vectorization Effectiveness metrics



Performance analysis metrics supported by each Arm core depend on the hardware events implemented by the core. For more information about the metric groups and the list of metrics within each group, see the relevant Arm core telemetry specification. For more information about available Arm core telemetry specifications, see [Arm® Telemetry on Arm Developer](#).

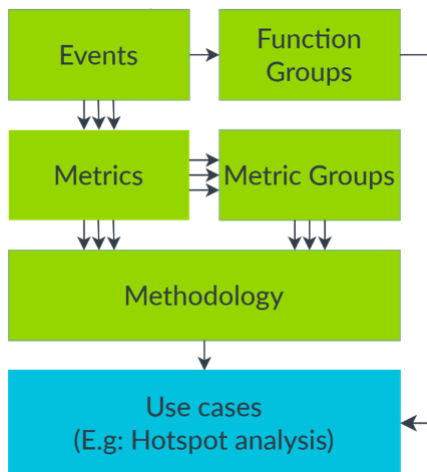
3. Arm Telemetry framework for CPUs

Arm Telemetry framework for CPUs has key elements that follow a standardized data model to produce and consume the CPU's telemetry data. The data model helps specify the telemetry features for the end users and profiling tools.

This framework supports all collaterals for telemetry, that is, written specifications for end users, for example, software analysts or system designers, and machine-readable specifications for profiling tools.

The main elements of the framework are events, metrics, metric groups, and methodology. The following figure shows those elements as a high-level data model.

Figure 3-1: Data model for Arm Telemetry framework for CPUs



PMU events implemented by an Arm core are either:

- Common, that is, the events are defined by the Arm architecture.
- **IMPLEMENTATION DEFINED**, that is, the events are specific to an Arm CPU.

Events are also either:

- Architectural - the events give the same result for the same workload on all Arm cores implementing the Arm architecture, to within a reasonable degree of accuracy, as described by the Arm architecture.
- Microarchitectural - the events give different results for the same workload on different Arm cores.

All events supported by an Arm core are grouped by function for each CPU resource in which they are counted. Metrics are derived from events that support the Topdown methodology and are grouped into metric groups for analysis. Key metrics can be standardized for an Arm processor product family. Event implementations that form a metric can vary for each microarchitectural

requirement. This design approach standardizes and abstracts metrics such that an analyst with software expertise can consume the events directly without knowing the hardware design in depth.



Note

Each definition of the framework element contains an example from the Arm Machine Readable Specification schema. All Arm cores that support the Arm Telemetry framework publish the associated Arm core telemetry specification with the MRS source data to the Arm CPU Telemetry Solution data repository in GitLab: [Arm® Telemetry Solution GitLab repository](#)

Events

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

Data fields

- Event Code
- Event Mnemonic
- Event Title
- Description
- Functional Category

Example event

```
"L1I_CACHE_REFILL": {
  "code": "0x0001",
  "title": "Level 1 instruction cache refill",
  "description": "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.",
  "accesses": [
    "PMU",
    "ETE"
  ],
  "architecture_defined": true,
  "product_defined": false
}
```

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.

Data fields

- Metric Name
- Metric Title
- Metric Description
- Metric Formula
- Metric Unit
- Metric Events

Example metric

```
"l1i_cache_mпки": {
  "title": "L1I Cache MPKI",
  "formula": "((L1I_CACHE_REFILL / INST_RETIRED) * 1000)",
  "description": "This metric measures the number of level 1 instruction cache accesses missed per thousand instructions executed.",
  "units": "MPKI",
  "events": [
    "INST_RETIRED",
    "L1I_CACHE_REFILL"
  ]
}
```

Metric groups

Groups of metrics that are analyzed together for performance correlation during a specific investigation. A group can be dependent on a usage model. Thus, some metrics can belong to multiple metric groups.

Data fields

- Metric Group Name
- Metric Group Title
- Metric Group Description
- Metric Group Metrics

Example metric group

```
"Topdown_L1": {
  "title": "Topdown Level 1",
  "description": "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.",
  "metrics": [
    "frontend_bound",
    "backend_bound",
    "retiring",
    "bad_speculation"
  ]
}
```

Events required for the metrics can be obtained from the metric data schema.

Methodology

Methodology is actionable guidance to explain how to consume the different metrics and events for a specific usage model, for example, hotspot analysis. Methodology can be a metric group or collection of metric groups. It provides actions and guidance notes on how to consume the data to take steps or derive actionable insights from the data.

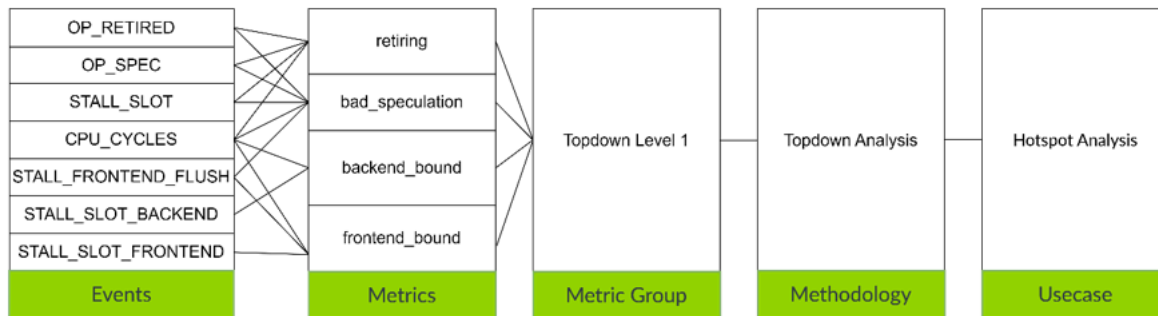
An example methodology is Arm Topdown methodology. It is a decision tree with root and leaf nodes of metrics that belong to specific metric groups. Metric relationships are key to represent the Topdown tree. Metric groups are required for the analysis process. They provide guidance for profiling data collection tools. Guidance includes how many hardware monitoring events to count, and in which order, because hardware resources are limited. Relationships between metrics and metric groups is specified for a methodology. Stage 1 and stage 2 metric relationships are captured in the methodology specification schema.

Topdown Methodology schema

```
"topdown_methodology": {
  "title": "Topdown Methodology",
  "description": "Topdown Performance Analysis Methodology",
  "metric_grouping": {
    "stage_1": [
      "Topdown_L1"
    ],
    "stage_2": [
      "Cycle_Accounting",
      "General",
      "MPKI",
      "Miss_Ratio",
      "Branch_Effectiveness",
      "ITLB_Effectiveness",
      "DTLB_Effectiveness",
      "L1I_Cache_Effectiveness",
      "L1D_Cache_Effectiveness",
      "L2_Cache_Effectiveness",
      "LL_Cache_Effectiveness",
      "Operation_Mix"
    ]
  },
  "decision_tree": {
    "root_nodes": [
      "frontend_bound",
      "backend_bound",
      "retiring",
      "bad_speculation"
    ],
  },
  "metrics": {
    "name": "frontend_bound",
    "group": "Topdown_L1",
    "next_items": [
      "Branch_Effectiveness",
      "ITLB_Effectiveness",
      "L1I_Cache_Effectiveness",
      "L2_Cache_Effectiveness",
      "LL_Cache_Effectiveness"
    ],
    "sample_events": [
      "STALL_SLOT_FRONTEND",
      "STALL_FRONTEND"
    ]
  }
}
```

The following figure shows how the different elements of the Arm Telemetry framework are used to construct the level 1 metrics of the Topdown methodology, grouped as the Topdown level 1 metric group.

Figure 3-2: An example demonstration of Telemetry framework relationship



Because hardware counters are limited, it is important that the profiling data collection tool has sufficient counters to collect the data required for an analysis stage or a metric group. For the Topdown analysis level 1 metrics, as shown in the previous figure, a total of seven events are required for full coverage. The number of counters required will vary for each CPU, depending on the events required to derive the metrics. Software uses multiplexing when collecting metrics in a group that requires more events than the number of available hardware counters. This discrepancy can result in accuracy issues. Thus, hardware requirements for supported counters must be derived for each methodology requirement to achieve a reliable solution.

3.1 Data model standardization

All components of the Arm CPU Telemetry Solution are designed for scalability to enable an Arm CPU implementation to adopt this solution. The elements of the Arm Telemetry framework can be used across individual Arm cores or Arm processor families.

The elements also provide the flexibility for the characteristics of an Arm processor family or an individual Arm core, for example, a core's events implementation. Thus, the metric formulae can vary as required. An Arm processor family refers to a family of Arm processor IP that implement features of the A-profile architecture. Each family supports different markets and their requirements, for example, Cortex®-A, Cortex-X, Neoverse, and Arm processors in Arm Lumex™ platforms.

For example, Topdown analysis defines a collection of metric groups that form a decision tree for a hotspot analysis use case. The nodes can be designed as a breadth first search tree, such that the traversal direction is decided by the metric weights that are defined for tree nodes at a single level. Metric groups, metric formulae, and relations can be updated for each Arm CPU implementation. The elements of the Arm Telemetry framework let Arm processor IP teams easily define their microarchitectural events and metric formulae to derive the Topdown analysis relationships. However, the interfaces remain standardized for tooling and collateral generation.

Arm Telemetry framework can be easily adopted by any Arm CPU implementation, thereby achieving a standard telemetry solution for the Arm ecosystem. For more information about how the solution enables the profiling tools in both Linux and Windows environments, see [Arm Telemetry Solution for software profiling: tools](#).

4. Arm Telemetry Solution for software profiling: tools

Performance monitoring features in a system have two main usage models, profiling and software optimization and run-time monitoring. Both models are enabled by established profiling tools.

Each usage model uses counting and sampling to read the hardware telemetry data supported by the platform hardware and the software telemetry data supported by system software.

Profiling and software optimization

Profiling tools that are used in software optimization get the telemetry data from system software and hardware performance monitors, which identify bottlenecks or areas of improvement. The main goal of optimization is primarily to reduce the elapsed time required to execute the program. While reducing power consumption can be a goal, it is usually a side effect of reducing the elapsed time.

Run-time monitoring

Monitoring tools that are used in system operations get the telemetry data from system software and hardware performance monitors across the system. The goal of performance monitoring is to provide a system operator with answers that can improve the efficiency of the system. Run-time monitoring helps to:

- Understand the current performance of the system.
- Evaluate whether the system executes efficiently or requires any further tuning or added capacity to meet the load demands to meet the Service Level Agreement (SLA). An SLA in a system meets sustained throughput or latency requirements of the deployed application.

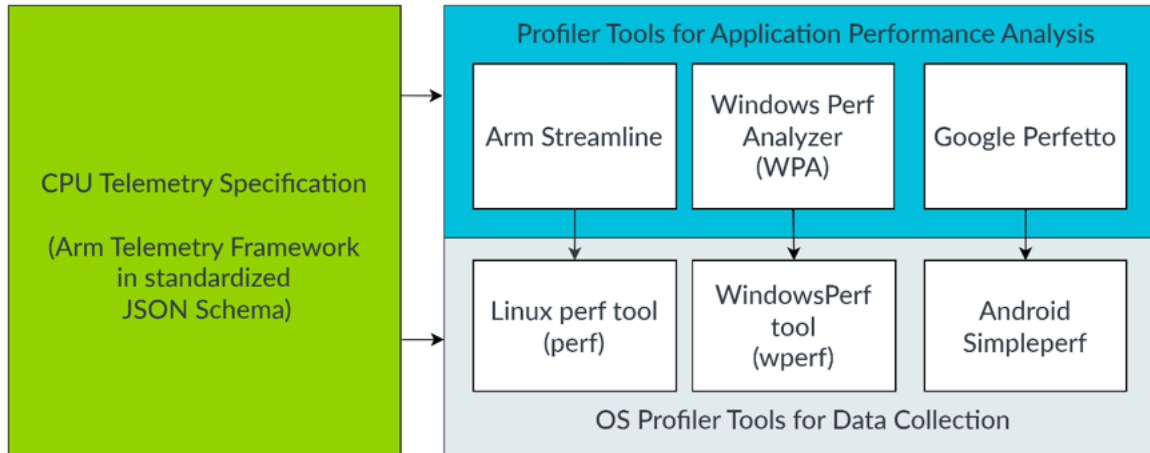
In a data center environment, observability and continuous run-time monitoring can cause the operator to do effective capacity planning and, for example, adjust capacity or rearrange load on the available machines. The job can be automated by software to monitor the system and take the required actions to guarantee the application SLA. Telemetry data from the SoC hardware is heavily used in this run-time monitoring, along with system software and application software provided metrics.

4.1 Arm telemetry specification and profiling tools

The Arm cores that support Arm CPU Telemetry Solution provide a telemetry specification that contains the methodology, metric groups, metrics, and hardware performance monitoring events which are supported by the core.

This specification is also provided in a standardized machine-readable format, that is, an Arm Telemetry JSON file that follows the schemas for the elements of the Arm Telemetry framework. See [Arm Telemetry framework for CPUs](#). These files can be consumed by any profiling or monitoring tool as shown in in the following figure.

Figure 4-1: Arm Telemetry Solution and profiling tool enablement



Common profiling tools are Google Perfetto which is heavily used in the Android ecosystem and Microsoft Windows Performance Analyzer (WPA) which is heavily used in the Windows ecosystem. These tools either rely on the performance data collection capabilities provided by the underlying operating system or might support the tools' drivers for data collection. They can directly use the Arm Telemetry JSON files for telemetry data collection from the hardware.

Operating systems support performance data collection using utilities, for example, the Linux perf tool on Linux OS, WindowsPerf tool on Windows OS, and simpleperf tool on Android. These utilities also have support to collect performance metrics from the hardware monitoring units. They can use the Arm Telemetry specification to get the hardware information for event collection which can be post processed to derive metrics.

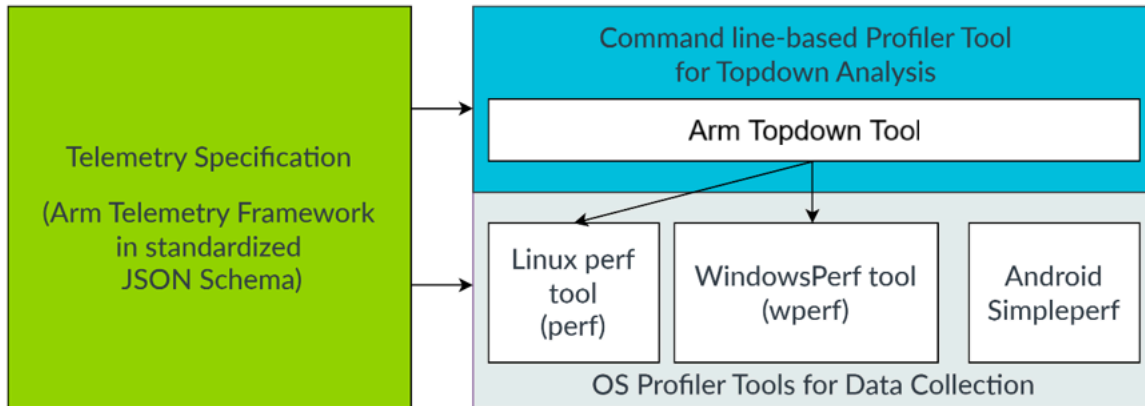
Arm recommends collecting all metrics that are in stage 1 and stage 2 Topdown analysis for workload characterization. For further analysis, an Arm core telemetry specification provides a recommended set of Microarchitecture Analysis metric groups against some hotspots detected in stage 1. All stage 2 metrics can be used to derive further insights into the overall microarchitecture behavior during the execution of the application under investigation. These metrics can be used independently of stage 1.

4.1.1 Arm Topdown tool

Arm CPU Telemetry Solution is implemented by the Arm Topdown tool. This tool supports collecting the recommended set of the telemetry data and processing it to derive the relevant set of metrics for application Topdown analysis.

It is a command line tool that supports using Linux perf and WindowsPerf to profile applications on Linux and Windows platforms respectively. The tool parses the machine-readable telemetry specifications provided as JSON files. It collects the data that is required to help the user with the Arm Topdown methodology stages, as shown in the following figure.

Figure 4-2: Arm Topdown tool



The tool is available to download from the Arm CPU Telemetry Solution: [Arm® Telemetry Solution GitLab repository](#)

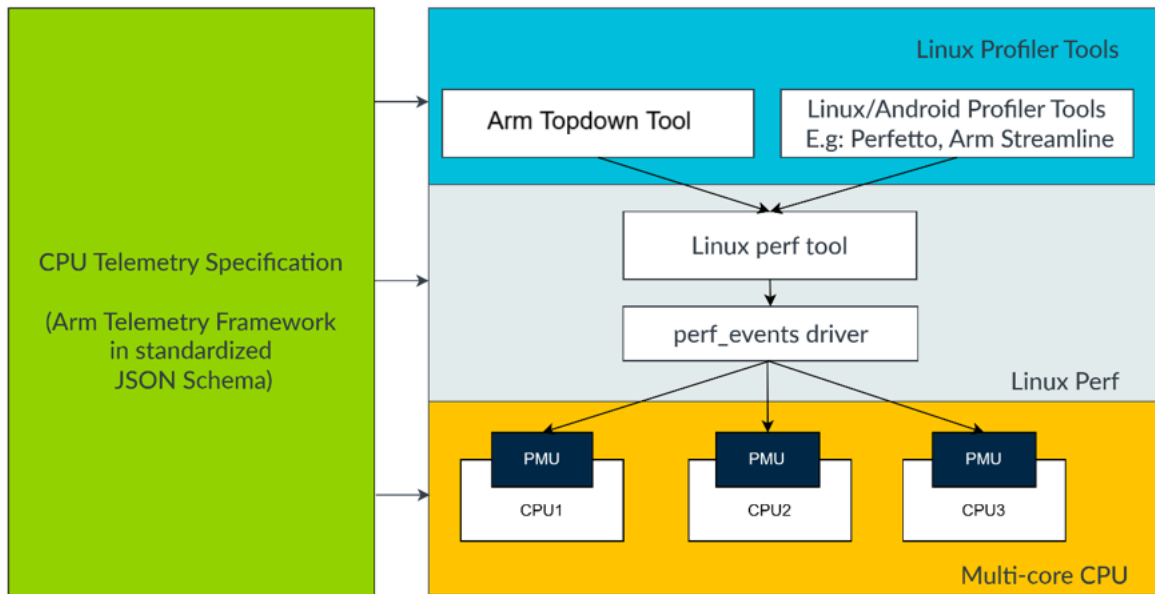
For introductory guidance about how to use the tool, see [Arm Topdown tool example](#).

4.1.2 Performance analysis using Linux perf tool

The Linux perf tool is a widely used open source performance analysis tool for collecting hardware and software performance events from different sources in the hardware and system software.

The kernel employs a `perf_event` subsystem to collect measurable events including the hardware PMU events from the CPU. As shown in the following figure, each CPU has its dedicated PMU hardware and the kernel perf driver collects events from each CPU's PMU separately.

Figure 4-3: Performance monitoring architecture for CPUs in a Linux based platform



The Linux `perf_event` subsystem provides an interface between the Linux kernel and user space performance monitoring tools to collect the raw hardware events as needed. As an open source tool, Linux perf can be used for performance monitoring, which supports the following types of performance measurement techniques:

Counting

Counting method collects overall statistics of an event during a workload's execution, where the counter assigned with each event produces an aggregate of the overall event count. These event statistics help to characterize the overall workload execution behavior, without providing any details on where a particular event occurred in the program. This method is the best approach for an initial workload characterization exercise to identify performance limitations of the workload.

Event sampling

Event sampling is a profiling method where each event is sampled, by configuring the PMU counter to overflow after a preset number of events. This overflow interrupt records the event count and also the program counter address and register information. Such sampled data is used to construct profiling information about the application, including stack trace and function level annotations. With this data, it is easy to locate the libraries and code portions that contribute to the large portion of the sampled event.

These measurement techniques or modes are available through the following perf commands:

perf-stat

Provide performance counter statistics for overall execution of the program. For more information, see [Linux perf-stat](#).

perf-record

Record the execution performance with the percentage of samples for each event for all libraries and functions. For more information, see [Linux perf-record](#).

perf-report

Generate a report of the recorded sample using record. For more information, see [Linux perf-report](#).

perf-annotate

Annotate a report with the samples' percentage on the disassembly of the code. For more information, see [Linux perf-annotate](#).

When high accuracy is needed, for example, when profiling hot loops or significant portions of code, the Counting mode should be preferred for its accuracy. It might require multiple profiling iterations when many different events must be logged.

When the number of events is greater than the total number of available counters, the counter is time multiplexed between events, and the final count is scaled for the total time period. This multiplexed counting can cause accuracy issues, but it is sufficient unless a precise measurement is needed.

The event sampling mode is extremely useful for hotspot analysis which relies on a statistical approach to sample different events over a large portion of time or code. This method has limitations that cause accuracy issues, such as:

- Sampling delay, that is, between the counter overflow and interrupt handler, which causes skid in the data obtained, that is, data stored during the sampling process and may not be the exact point where the event occurred.
- Speculative execution style of the CPU, whereby some instructions that executed and triggered events might not be valid if they were on the wrong code path.

While the event sampling mode has accuracy limitations, it is the best way to advance identification of hotspots in code execution. Linux perf allows tuning the sampling frequency, which helps to study variations in the event counts if the data shows large inconsistencies across runs.

For introductory guidance about how to use Linux perf to collect hardware counters on Arm cores in both counting and sampling modes, see [Linux perf data collection](#).

4.1.3 Performance analysis using WindowsPerf Tool

WindowsPerf is a Windows on Arm performance profiling tool, which is `Linux perf` inspired. Profiling is based on the Performance Monitors Extension, `FEAT_PMU`, and its hardware counters. It supports both counting and sampling models.

WindowsPerf can instrument Arm CPU performance counters. Currently, it can collect:

- Core PMU counters for all or specified processor cores
- unCore PMU counters such as:
 - Arm DynamIQ™ Shared Unit PMU
 - Arm CoreLink™ DMC-520 PMU
- Arm Statistical Profiling Extension (SPE)

- Arm Neoverse CMN-700
- Arm Neoverse CMN S3

The tool supports two models.

Counting model

Obtains aggregate counts of special event occurrences.

For example usage of the counting model, see [Counting model](#) in `wperf` [README.md](#).

Sampling model

Determines the frequencies of event occurrences that are produced by program locations at different levels such as functional, basic block, and instruction.

Sampling model features include:

- Sampling mode initial merge. See [Merge request 111: Merge sampling model to the project](#).
- Support for DLL symbol resolution. See [Merge request 132: Add support for DLLs symbols printout](#).
- Deduce from the command line image name and program debug or PDB file name. See [Merge request 134: Deduce from command line image and PDB files for sampled executable](#).
- Stop sampling when sampled process ends. See [Merge request 135: Stop sampling when sampled process exits](#).

For example usage of the sample model, see [Sampling model](#) in `wperf` [README.md](#).

Arm Telemetry Solution integration

The integration of WindowsPerf and Arm Telemetry Solution is a significant advancement in performance analysis on Windows On Arm. This integration is primarily based on PMU events, which provide a detailed insight into the system's performance.

This methodology is tailored for each Arm CPU's microarchitecture. It involves the use of PMU events, metrics, and groups of metrics to provide a comprehensive analysis of the system's performance. Furthermore, the WindowsPerf Tool is capable of platform microarchitecture detection, including Neoverse N1, N2, N3, V1, V2 and V3 cores, and processors in Arm Lumex platforms.

The Arm Telemetry Solution also includes a `topdown-tool` that leverages the WindowsPerf as a backend for Windows On Arm. This tool applies the Arm Topdown methodology to break down the CPU and CMN mesh performance into different hierarchical levels. It provides a detailed and systematic approach to performance analysis. The `topdown-tool` uses the WindowsPerf to access the core and CMN mesh PMU events and metrics. See [Arm Topdown tool example](#).

WindowsPerf offers human readable console output and a standardized JSON output to file, console, or CSV output for timeline mode.

WindowsPerf architecture

The WindowsPerf architecture comprises a driver and CLI.

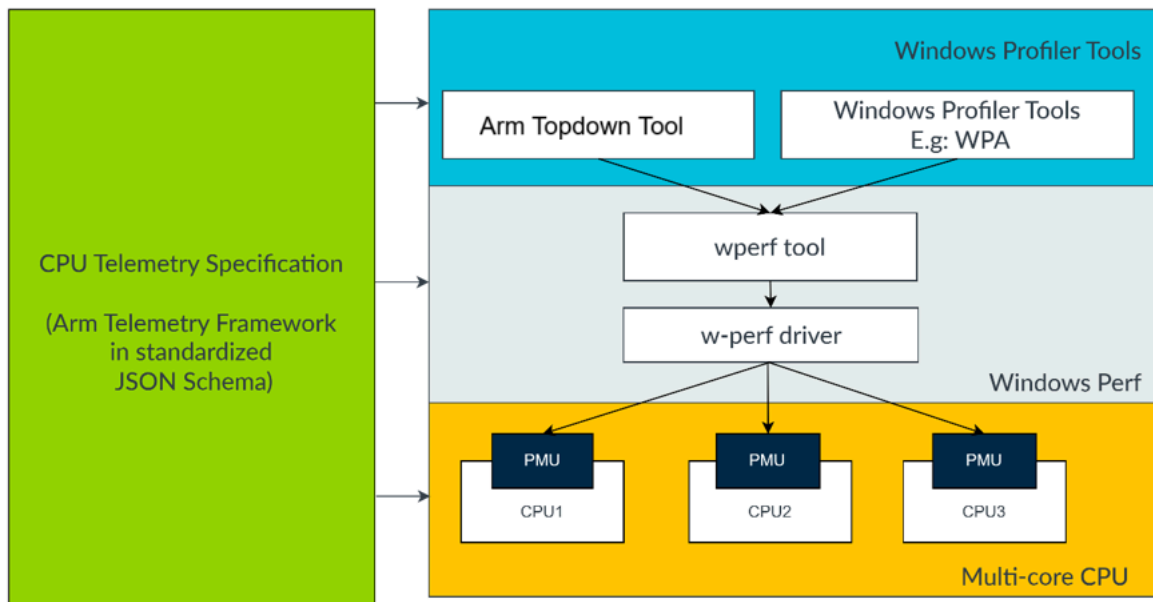
wperf-driver - a signed Windows kernel mode driver



WindowsPerf Kernel driver is signed only for Windows 11.

- wperf - a command line interface like Linux perf, as shown in the following figure.

Figure 4-4: Performance monitoring architecture for CPUs in a Windows based platform



Install WindowsPerf

You can install WindowsPerf using the Windows Package Manager, winget.



Run WinGet with administrator privileges in the Command Prompt or PowerShell window.

To install, run:

```
winget install WindowsPerf
```



Run `winget1 --help` to list all commands and options.

For more information about installing WindowsPerf and the source and package binaries, including `wperf-driver` and `wperf`, see [INSTALL.md](#) and [Releases](#) respectively.

WindowsPerf counting and sampling modes

WindowsPerf supports the counting and sampling modes through different commands.

wperf stat

Provides performance counter statistics for the overall execution of Arm Neoverse processors. It supports multiplexing capabilities for events and can profile between 1 and *n* cores as specified. In counting mode, timeline mode is also supported, which refers to consecutive counting of events and a CSV file format output.

You can count:

- [Core PMU events](#)
- [CMN mesh PMU and Watchpoint events](#) for CMN-700 and CMN-S3 r0p1

wperf sample and wperf record

WindowsPerf provides `sample` and `record` commands to support:

- PMU event sampling for core events, which records the execution performance by using the percentage of samples for each event associated with libraries and functions in a Windows on Arm application. In sampling mode, users can sample Windows on Arm applications that are pinned to a specified core. WindowsPerf offers sampling output with hot function names annotated with:
 - Source code file name + line number
 - Optional disassembly
- Sampling with Arm Statistical Profiling Extension (SPE). For more information, see [Sampling with SPE](#).

For more information about how to use WindowsPerf to collect hardware counters in both counting and sampling modes, see [WindowsPerf tool data collection].

Appendix A Arm Topdown tool example

Use the following steps to conduct application hotspot analysis using the Arm Topdown tool.



Arm Topdown tool, `topdown-tool` provides rich options to fine tune the profiling process. Run `topdown-tool --help` to list all options.

Step 1, Collect the stage 1 Topdown analysis metrics

In this step, we specify the Topdown methodology stage to profile, for example, `topdown` for stage 1, and `uarch` for stage 2.

```
topdown-tool --cpu-stages topdown ./a.out
```

Results for Topdown level 1:

```
CPU Neoverse V2 metrics
├─ Stage 1 (Topdown metrics)
│   └─ Topdown Level 1 (Topdown_L1)
│       ┌──────────┴──────────┐
│       │ Metric              │ Value │ Unit │
│       └──────────┴──────────┘
│       ┌──────────┴──────────┐
│       │ Backend Bound       │ 98.00 │ %    │
│       │ Bad Speculation     │ 0.01  │ %    │
│       │ Frontend Bound      │ 0.05  │ %    │
│       │ Retiring            │ 1.82  │ %    │
```

Based on these results, this workload shows a very high backend bound metric.

Step 2, Collect the stage 2 Microarchitecture Analysis metrics

In this step, we can specify the metric groups of interest for further analysis. In the following example, the cache effectiveness is checked for a high backend bound workload.

```
topdown-tool --cpu-metric-group L1D_Cache_Effectiveness,L2_Cache_Effectiveness ./a.out
```

Results for Stage 2:

```
CPU Neoverse V2 metrics
├─ Stage 2 (uarch metrics)
│   └─ L1 Data Cache Effectiveness (L1D_Cache_Effectiveness)
│       └─ Follows
│           └─ Backend Bound (backend_bound)
│               ┌──────────┴──────────┐
│               │ Metric              │ Value │ Unit │
│               └──────────┴──────────┘
│               ┌──────────┴──────────┐
│               │ L1D Cache Miss Ratio │ 0.972 │ per cache access │
│               │ L1D Cache MPKI       │ 476.998 │ misses per 1,000 instructions │
```

└─ L2 Unified Cache Effectiveness (L2_Cache_Effectiveness)

└─ Follows

└─ Backend Bound (backend_bound)

└─ Frontend Bound (frontend_bound)

Metric	Value	Unit
L2 Cache Miss Ratio	0.014	per cache access
L2 Cache MPKI	13.093	misses per 1,000 instructions

Based on these results, there are very high L1D cache misses.

There are many metric groups to explore. For example, we can check the proportions of different kinds of instructions, for example, the Operation Mix metric.

```
topdown-tool --cpu-metric-group Operation_Mix ./a.out
```

Results for Speculative Operation Mix:

CPU Neoverse V2 metrics

└─ Stage 2 (uarch metrics)

└─ Speculative Operation Mix (Operation_Mix)

└─ Follows

└─ Backend Bound (backend_bound)

└─ Retiring (retiring)

Metric	Value	Unit
Barrier Operations Percentage	0.02	%
Branch Operations Percentage	48.10	%
Crypto Operations Percentage	0.00	%
Integer Operations Percentage	2.85	%
Load Operations Percentage	47.61	%
Floating Point Operations Percentage	0.00	%
Advanced SIMD Operations Percentage	0.00	%
Store Operations Percentage	0.79	%
SVE Operations (Load/Store Inclusive) Percentage	0.00	%

Appendix B Linux perf data collection

To enable PMU event collection, the Linux Kernel must be built with `CONFIG_HW_PERF_EVENTS` enabled in the kernel config.

Most of the production builds have this config option enabled. However, ensure that this option is enabled when building a custom kernel.

For more information about `perf_event` and unprivileged users, see [Linux perf_event and tool security, Unprivileged users](#).

There are also two system settings which need to be configured as root user to obtain kernel symbols and add extra privileges as follows:

```
echo -1 > /proc/sys/kernel/perf_event_paranoid
echo 0 > /proc/sys/kernel/kptr_restrict
```

perf_event_paranoid

This setting affects privilege checks in the kernel. If set to `-1`, it permits enabling of events that might reveal sensitive information or can impact the stability of the system. For more information about this setting, see [Linux perf_event_paranoid](#).

kptr_restrict

This setting affects whether kernel addresses are exposed, that is, through `/proc/kallsyms`. Some developers use this technique to get kernel symbol resolution when they do not have the `vmlinux` to hand, or where `KASLR` is in use. For more information about this setting, see [Linux kptr-restrict](#).

In both cases, there are potential security implications, so it is advised to check the official kernel documentation and consult the system administrator before enabling them.

A quick test to verify that PMU events are being counted properly is to use the `perf stat` functionality of the Linux perf tool to count instructions and cycles. `perf stat` counts the total count of a specified event, provided as the hex register code. `0x8` is the hex code for instructions that are retired and `0x11` is the hex code for CPU cycles specified by the Arm architecture common across all Arm CPU implementations.

These events are provided to the `perf stat -e` option with a prefix `'r'` to it. For a code example, see [Collect hardware PMU events using counting mode on Linux perf](#).

The `perf stat` command counts the total count of instructions and cpu cycles on all CPUs for 10 seconds. Linux perf allows to count for a particular CPU, for each process, for each thread and so on.

`perf` can silently fail if an event is not supported or enabled on an Arm core. For more information about supported events, see the specific Arm core telemetry specification or its Technical Reference Manual.

Collect hardware PMU events using counting mode on Linux perf

To count all events for characterization, a typical solution is to capture events in batches of the total counter registers available in the platform.

One method to determine the number of counters available is to successively increase the number of counters that are requested in a single group until the last counter reads <"not supported">. If this method is performed through the Linux perf tool, it opens N+1 counters, because perf always opens the group leader event in disabled mode, and the kernel does not count this event towards the number of available events.

This example uses the `-e` option to pass the instructions and cycles events to `perf stat` by using their hex codes, that is, `'r8'` and `'r11'` respectively:

```
ubuntu@linux-1:~$ perf stat -e r8,r11 -- sleep 10
Performance counter stats for 'sleep 10':

    1242692      r8
    1000747      r11

.001721954 seconds time elapsed

0.000772000 seconds user
0.000000000 seconds sys
```

Collect hardware PMU events using sampling mode on Linux perf

For sampling events, use the `perf record` command from Linux perf tool.

For more information about how to conduct sampling and analyze the sampled data with these command lines, see the following Linux perf examples.

```
ubuntu@linux-1:~$ perf record -e instructions, cycles -- sleep
ubuntu@linux-1:~$ perf report
ubuntu@linux-1:~$ perf annotate
```

Proprietary Notice

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

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

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

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

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

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

This document consists solely of commercial items. You shall be responsible for ensuring that any permitted use, duplication, or disclosure of this document complies fully with any relevant

export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws. Use of the word “partner” in reference to Arm’s customers is not intended to create or refer to any partnership relationship with any other company. Arm may make changes to this document at any time and without notice.

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

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

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

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

PRE-1121-V1.0

Product and document information

Read the information in these sections to understand the release status of the product and documentation, and the conventions used in Arm documents.

Product status

All products and services provided by Arm require deliverables to be prepared and made available at different levels of completeness. The information in this document indicates the appropriate level of completeness for the associated deliverables.

Product completeness status

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

Revision history

These sections can help you understand how the document has changed over time.

Document release information

The Document history table gives the issue number and the released date for each released issue of this document.

Document history

Issue	Date	Confidentiality	Change
02-02	10 September 2025	Non-Confidential	Second issue to support new processors in Arm® Lumex™ platforms and enhancements to WindowsPerf and Arm Topown tool for new products.
01-01	23 September 2024	Non-Confidential	First issue for all revisions of the product.

Change history

The Change history tables describe the technical changes between released issues of this document in reverse order. Issue numbers match the revision history in [Document release information](#) on page 33.

Table 2: Issue 02-02

Change	Location
Updated to support new processors in Arm® Lumex™ platforms	Throughout the document

Change	Location
Enhancements to WindowsPerf tool to support Arm new processors in Lumex platforms and Neoverse™ CMN-700 and S3 products Performance analysis using WindowsPerf Tool	
Enhancements to Arm Topdown tool to support Arm new processors in Lumex platforms and Neoverse™ CMN-700 and S3 products Arm Topdown tool example	

Table 3: Issue 01-01

Change	Location
First release	-

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
<i>italic</i>	Citations.
bold	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 <u>underline</u>	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: <pre>MRC p15, 0, <Rd>, <CRn>, <CRm>, <Opcode_2></pre>
SMALL CAPITALS	Terms that have specific technical meanings as defined in the <i>Arm® Glossary</i> . 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 of harming yourself.



This information is important and needs your attention.



This information might help you perform a task in an easier, better, or faster way.




This information reminds you of something important relating to the current content.

Timing diagrams

The following figure explains the components used in timing diagrams. Variations, when they occur, have clear labels. You must not assume any timing information that is not explicit in the diagrams.

Shaded bus and signal areas are undefined, so the bus or signal can assume any value within the shaded area at that time. The actual level is unimportant and does not affect normal operation.

Figure 1: Key to timing diagram conventions



Signals

The signal conventions are:

Signal level

The level of an asserted signal depends on whether the signal is active-HIGH or active-LOW. Asserted means:

- HIGH for active-HIGH signals.
- LOW for active-LOW signals.

Lowercase n

At the start or end of a signal name, n denotes an active-LOW signal.

Useful resources

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

Arm documents are available on developer.arm.com/documentation.

Confidential documents are only available to licensees, when logged in. Each document link in the tables below provides direct access to the online version of the document.

Arm product resources	Document ID	Confidentiality
Arm® Telemetry Solution GitLab repository	–	Non-Confidential
Arm® Telemetry on Arm Developer	–	Non-Confidential

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

Non-Arm resources	Document ID	Organization
Linux kptr-restrict	–	Linux kptr-restrict
Linux perf-annotate	–	Linux perf-annotate
Linux perf-record	–	Linux perf-record
Linux perf-report	–	Linux perf-report
Linux perf-stat	–	Linux perf-stat
Linux perf_event and tool security, Unprivileged users	–	Linux perf_event and tool security, Unprivileged users
Linux perf_event_paranoid	–	Linux perf_event_paranoid