



Arm Streamline

Version 9.7

Target Setup Guide for Linux

Non-Confidential

Copyright © 2021–2025 Arm Limited (or its affiliates).
All rights reserved.

Issue 03

101814_9.7_03_en



Arm Streamline Target Setup Guide for Linux

This document is Non-Confidential.

Copyright © 2021–2025 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 (101814_9.7_03_en) was issued on 2025-12-11. There might be a later issue at <https://developer.arm.com/documentation/101814>

The product version is 9.7.

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 book is intended for users who need to use Arm® Streamline Performance Analyzer on Linux-based embedded targets. It describes how to prepare your Linux target and how to capture a profile with Streamline.

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. Target Setup on Linux.....	4
1.1 Introduction to using Streamline with a Linux device.....	4
1.2 Compile your Linux application.....	4
1.3 Set up your target.....	5
1.4 System requirements for SPE support.....	7
1.5 System requirements for Mali Timeline support.....	7
1.6 gatord.....	8
1.6.1 gatord command-line options on Linux.....	9
1.7 Support for Arm CoreLink tools.....	14
1.8 Using Streamline with Arm NN.....	14
2. Customize your Streamline report.....	16
2.1 Add a CPU to Streamline.....	16
2.2 Add an uncore peripheral to Streamline.....	16
2.3 Contributing CPU and peripheral support.....	17
3. Advanced target setup information (Linux).....	18
3.1 Kernel configuration menu options for Linux.....	18
3.2 Profile the Linux kernel.....	19
3.3 Build gatord yourself for Linux.....	20
3.4 Add support to gatord for a new CPU or perf PMU.....	20
4. Analytics data collection.....	22
Proprietary notice.....	24
Product and document information.....	26
Product status.....	26
Revision history.....	26
Conventions.....	27
Useful resources.....	30

1. Target Setup on Linux

This chapter explains how to set up your target and host devices ready to use Streamline for application or system profiling.

1.1 Introduction to using Streamline with a Linux device

Streamline supports system-wide profiling of applications running on Linux-based embedded devices. Analyze the behavior of the system hardware by selecting the required Arm CPU, Arm® Mali™ GPU, or Arm® Immortalis™ GPU hardware performance counters for your scenario.

To provide more context to the analysis, you can use software annotations in Arm® software libraries, such as the Mali™ GPU OpenCL device driver.

Use the information in this guide to prepare your Linux target, ready for profiling with Streamline. Then refer to the [Arm Streamline User Guide](#) for details about working with Streamline.

1.2 Compile your Linux application

Before you can profile your executable with Streamline, you must compile your executable. This topic describes the compiler options that you can use when you compile your application.

Recommended compiler options

Arm recommends the following options when you are compiling with GCC or Clang:

-g

Turns on the debug symbols necessary for quality analysis reports.

-fno-inline

Disables inlining and substantially improves the call path quality.

-fno-omit-frame-pointer

Compiles your EABI images and libraries with frame pointers. This option enables Streamline to record the call stack with each sample taken.

Additional compiler options for call stack unwinding

To enable call stack unwinding in Streamline, you can compile your executable with some additional compiler options:

- For AArch64 applications, the `-mno-omit-leaf-frame-pointer` option prevents samples in leaf functions incorrectly listing their grandparent function as their parent.
- For AArch32 applications, you can also use the following options:

-marm

The `marm` option is required if GCC was compiled with the `--with-mode=thumb` option enabled. Using `--with-mode=thumb` without `-marm` breaks call stack unwinding in Streamline.

-mapcs-frame

Generates consistent stack frames that conform to the Arm Procedure Call Standard, which is used by early Arm system configurations. `-mapcs-frame` is supported by GCC only.



Streamline supports call stack unwinding for code that has been generated by Arm® Compiler 6. Streamline does not support call stack unwinding for T32 (Thumb®) code.

Source code annotations

To enable Streamline to provide extra context when profiling your executable, you can add annotations to your source code. Streamline supports two types of annotations:

- User space annotations, for annotating your application
- Kernel annotations, to profile system calls

You can read more information about annotating your code in the [Annotate your code](#) chapter of the Arm Streamline User Guide.

1.3 Set up your target

To profile your software using Streamline, you must have a suitably configured Linux kernel and `gator` running on an Arm®-based hardware target.

Before you begin

- Download and install the studio package appropriate to your host platform (Windows, Linux, or macOS):
 - Download Arm® Performance Studio from [Arm Performance Studio Downloads](#), and install using the instructions in the [Arm Streamline User Guide](#).
 - Download Arm® Development Studio from [Arm Development Studio Downloads](#), and install using the instructions in the [Arm Development Studio Getting Started Guide](#).
- If you intend to use the Statistical Profiling Extension (SPE) in Streamline, there are some additional system requirements. Refer to [System requirements for SPE support](#).

Procedure

1. Validate the Linux kernel configuration.

The kernel configuration must include the options that are described in [Kernel configuration menu options for Linux](#). You can usually find your kernel configuration in `/proc/config.gz`. If this file is not visible, depending how your system is configured, you can create it by running:

```
sudo modprobe configs
```

The following command confirms whether an option is enabled:

```
zcat /proc/config.gz | grep <OPTION>
```

For example:

```
zcat /proc/config.gz | grep CONFIG_PROFILING
```

2. Install `gator` on your target.

Two pre-built `gator` binaries are available:

- For Arm®v7 targets, and Arm®v8 targets that support AArch32 execution state.
- For Arm®v8, or later, AArch64 targets.

The pre-built `gator` binaries are available from the following locations:

- The Arm® Performance Studio installation directory `<install_directory>/streamline/bin/linux/`
- The Arm® Development Studio installation directory `<install_directory>/sw/streamline/bin/linux/`

The source code for `gator` is available from the following locations:

- The Arm® Performance Studio installation directory `<install_directory>/streamline/gator/daemon/`
- The Arm® Development Studio installation directory `<install_directory>/sw/streamline/gator/daemon/`
- <https://github.com/ARM-software/gator>. This site is the official distribution channel for all `gator` releases, and contains the latest source updates between suite releases.

The version of `gator` running on the target must be compatible with the version of Streamline that you are using. Use the same version of `gator` as the version of Streamline.

For more information, see [gator](#).

3. To ensure `gator` has execute permission, enter the following command:

```
chmod +x gator
```

4. Ensure that you have root privileges, if necessary, then enter the following to execute `gator`:

```
./gator
```

By default, `gator` uses port 8080 for communication with the host, but you can specify the port by launching `gator` with the port number as a parameter. For example:

```
./gator -p 5050
```

If the target runs a firewall that prevents Streamline from connecting to `gator`, Arm recommends that you change the firewall settings to allow connections to `gator`.

Next steps

- Capture a profile of your application as described in [Capture a Streamline profile](#).
- To stop `gator` at the end of a capture, press Ctrl+C.

1.4 System requirements for SPE support

When using the Statistical Profiling Extension (SPE) support in Streamline, there are some more system requirements.

- Use Linux kernel 4.16 or later. To confirm that support is available, check for the path `/sys/bus/event_source/devices/arm_spe_xx`.
- Disable kernel page table isolation for the target. To ensure that kernel page table isolation is disabled, boot the device with the command-line argument `kpti=off`.
- For kernel versions 4.20 and later, apply the patch at <http://lkml.iu.edu/hypermail/linux/kernel/1903.3/06760.html>, or upgrade to a kernel version greater than or equal to 5.1-RC5.

For more information about SPE, see [Configure SPE counters](#) in the Arm Streamline User Guide.

1.5 System requirements for Mali Timeline support

This topic details some system requirements to support Mali™ Timeline Events.



Note

Arm®Mali™ Timeline Events are supported on Linux from DDK version r43p0. To use the GPU Timeline layer driver, you must have a Mali™ device driver version r51p0 or later.

Enabling support for the Mali™ Timeline requires the following actions:

- Manually set the environment variable `MALI_GPU_RENDERSTAGES_ENABLE=1`
- Compile the driver with instrumentation and profiling support, and ensure the Perfetto `traced` daemon is running.

- To start profiling, configure the environment variable before starting the target application. As it starts, the driver checks the variable to see if it needs to start a profiling session and publish information to the Perfetto daemon.

For more information about Mali™ Timeline Events, see [Enable Mali Timeline events using the Streamline GUI](#) or [Enable Mali Timeline events for headless capture](#) in the Arm Streamline User Guide.

1.6 gatord

To communicate with the target, Streamline requires a daemon, called `gatord`, to be installed and running on the target. `gatord` must be running before you can capture trace data.

`gatord` reads and processes data from applications that are running on the target. When running a headless capture, it then creates a directory, whose name ends in `.apc`, containing the capture data.

`gatord` uses the `perf` API to collect performance information. You can run `gatord` on a device with or without root permissions. Running `gatord` on a device with root permissions supports system-wide data collection, including profiling all applications and the kernel. When run as non-root user, `gatord` is limited to capturing processes that you have permission to profile. These processes are usually running as the same user.



System-wide data collection using non-root `gatord` falls back to `/proc` polling, which is deprecated. To perform system-wide data collection, use root `gatord`.

`gatord` has some restrictions, for example:

- It polls the following Linux counters every 100ms, instead of every 1ms or when they change because files in the `/proc` or `/sys` filesystem are read:
 - Memory
 - Disk I/O
 - Network

This rate is fixed and overrides the sample rate that is specified in the **Capture Settings** dialog box.

- When using `gatord`, the `Memory: Used` counter does not contain per-process information. As a result, memory statistics are not available in **Processes** mode.

1.6.1 gatord command-line options on Linux

`gatord` must be running before you can capture trace data. The command-line options configure how `gatord` captures events and how it communicates with Streamline running on your host.

`gatord` has two modes of operation:

Daemon mode (the default mode)

Sends captured events to a host running Streamline.

Local capture mode

Writes the capture to a file then exits.

To enable this mode, specify an output directory with the `--output` flag.

Arguments available in all modes

The following `gatord` arguments are available for all modes:

-h, --help

Lists all the available `gatord` command-line options.

-c, --config-xml <config_xml>

Specify the path and filename of the `configuration.xml` file that defines the capture options. In daemon mode, the list of counters is written to this file. In local capture mode, the list of counters is read from this file.

-e, --events-xml <events_xml>

Specify the path and filename of the `events.xml` file. `events.xml` defines all the counters that Streamline collects during the capture session.

-E, --append-events-xml <events_xml>

Specify the path and filename of the `events.xml` file to append.

-P, --pmu-xml <pmu_xml>

Specify the path and filename of the `pmu.xml` file to append.

-v, --version

Print version information.

-d, --debug

Enable debug messages.

-A, --app <cmd> <args...>

Specify the command to execute when the capture starts. This argument must be the last argument that is passed to `gatord`. All subsequent arguments are passed to the launched application.

-D, --disable-kernel-annotations

Disable collection of kernel annotations.

-k, --exclude-kernel <yes|no>

Specify whether to filter out kernel events from the perf results.

-S, --system-wide <yes|no>

Specify whether to capture the whole system.

In daemon mode, `no` is only applicable when `--allow-command` is specified. In this mode, you must enter a command in the **Start** view.

Defaults to `yes`, unless `--app`, `--pid`, or `--wait-process` is specified.

-u, --call-stack-unwinding <yes|no>

Enable or disable call stack unwinding. Defaults to `yes`.

-r, --sample-rate <none|low|normal|high>

Specify the sample rate for the capture. The frequencies for each sample rate are:

- `high` = 10kHz
- `normal` = 1kHz
- `low` = 100Hz
- `none` = the lowest possible rate

-t, --max-duration <s>

Specify the maximum duration that the capture can run for in seconds. Defaults to 0, which means unlimited.

-f, --use-efficient-ftrace <yes|no>

Enable efficient ftrace data collection mode. Defaults to `yes`.

-w, --app-cwd <path>

Specify the working directory for the application that `gator` launches. Defaults to the current directory.

-x, --stop-on-exit <yes|no>

Stop the capture when the launched application exits. Defaults to `no`, unless `--app`, `--pid`, or `--wait-process` is specified.

-Q, --wait-process <command>

Wait for a process that matches the specified command to launch before starting the capture. Attach to the specified process and profile it.

-Z, --mmap-pages <n>

The maximum number of pages to map per mmaped perf buffer is equal to `<n+1>`. `n` must be a power of two.

-O, --disable-cpu-onlining <yes|no>

To not switch on CPU cores that are offline to read their information. This option is useful for kernels that fail to handle this action correctly, for example they reboot the system. Defaults to `no`.

-F, --spe-sample-rate <n>

Specify the SPE periodic sampling rate. The rate, `<n>`, is the number of operations between each sample, and must be a nonzero positive integer. The hardware specifies the minimum rate. Values below this threshold are ignored and the hardware minimum is used instead.

-L, --capture-log

Enable to generate a log file for the capture in the capture's directory, as well as sending the logs to `stderr`. If you are using the `streamline_me.py` script, gator is launched with `--capture-log` enabled by default.

--smmuv3-model <model_id>|<iidr>

Specify the SMMUV3 model. You can specify the model ID string directly, such as `mmu-600`, or the hex value representation for the model's IIDR number either fully, such as `4832243b`, or partially, such as `483_43b`.

-Y, --off-cpu-time <yes|no>

Collect Off-CPU time statistics. Detailed statistics require `root` permission.

-I, --inherit <yes|no|poll|experimental>

When profiling an application, gator monitors all subsequent threads and child processes. Specify `no` to monitor only the initial thread of the application. Specify `poll` to periodically look for new processes or threads. Specify `experimental` if you have applied the kernel patches provided by Arm for top-down profiling.



Per-function metrics are only supported in system-wide mode, or when `--inherit` is set to `no`, `poll` or `experimental`. The default is `yes`.

-N, --num-pmu-counters <n>

Override the number of programmable PMU counters that are available.

This option reduces the number of programmable PMU counters available for profiling. Use this option when the default is incorrect, or because some programmable counters are unavailable because they are consumed by the OS, or other processes, or by a hypervisor.



The Arm PMU typically exposes 6 programmable counters, and one fixed function cycle counter. This argument assumes the fixed cycle counter is not part of the reduced set of counters. If your target exposes 2 programmable counters and the fixed cycle counter, then 2 is passed for the value of `<n>`. However, if your target exposes 2 programmable counters and no fixed cycle counter, then 1 is passed for the value of `<n>`.

-g, --gpu-timeline <yes|no|auto>

Controls GPU Timeline data collection. The options are:

- `yes` enables collection and produces an error if the `MaliTimeline_Perfetto` counter is not enabled.
- `no` disables collection.
- `auto`, the default, collects data if the counter is enabled but otherwise disables collection without error.

**Note**

Timeline data is provided by a layer driver loaded into your application.

--gpu-timeline-layer-path <path>

The GPU Timeline layer library path. The default uses the standard layer.

Arguments available in daemon mode

The following arguments are available in daemon mode only:

-p, --port <port_number>|uds

Set the port number that `gator` uses to communicate with the host. The default is 8080.

If you use the argument `uds`, the TCP socket is disabled and an abstract Unix domain socket is created. This socket is named `streamline-data`. If you use Android, creating a Unix domain socket is useful because `gator` is usually prevented from creating a TCP server socket.

Alternatively, you can connect to `localhost:<local_port>` in Streamline using:

```
adb forward tcp:<local_port> localabstract:streamline-data
```

-a, --allow-command

Allows you to run a command on the device during profiling. The command is specified in the **Start** view.

**Caution**

If you use this option, an unauthenticated user could run arbitrary commands on the device using Streamline.

Arguments available in local capture mode

The following arguments are available in local capture mode only:

-s, --session-xml <session_xml>

Specify the `session.xml` file that the configuration is taken from. Any additional arguments override values that are specified in this file.

-o, --output <apc_dir>

Specify the path and filename of the output directory for a local capture. The directory path will be appended with the extension `.apc` if it is not already the case.

-i, --pid <pids...>

A comma-separated list of process IDs to profile.

-C, --counters <counters>

A comma-separated list of counters to enable. You can specify this option multiple times. An event code and a slot identify most hardware counters. To specify the counter for a particular slot, pass:

```
--counters <device>_cnt<s>:<e>
```

Where:

- <device> is the prefix that identifies the device type.
- <s> is the slot number.
- <e> is the event code.

```
-X, --spe <id>[:events=<indexes>][:ops=<types>][:min_latency=<lat>][:inv]
```

- Enable the Statistical Profiling Extension (SPE).

Where:

- <id> is the name of the SPE properties that are specified in the `events.xml` or `pmus.xml` file. It uniquely identifies the available events and counters for the SPE hardware.
- <indexes> is a comma-separated list of event indexes to filter the sampling by. A sample is only recorded if all events are present.
- <types> is a comma-separated list of operation types to filter the sampling by. If a sample is any of the types in <types>, it is recorded. Valid types are `LD` for load, `ST` for store and `B` for branch.
- <lat> is the minimum latency. A sample is only recorded if its latency is greater than or equal to this value. The valid range is [0,4096].
- `:inv` include this flag to invert the SPE event filter. This value is ignored if the device does not support SPE 1.2. Default is disabled.

Example 1-1: Add support for a new PMU without rebuilding gator

`--pmus-xml` specifies an XML file that defines a new PMU to add to the list of PMUs that `gator` has built-in support for. The list of built-in PMUs is defined in `:file:pmus.xml`, which is in the `gator` source directory.

`--append-events-xml` specifies an XML file that defines one or more event counters to append to the `events.xml` file. This option allows you to add new events to `gator` without having to rebuild `gator` or to entirely replace `events.xml`.

The `events.xml` file must include the XML header and elements that are shown in the following example:

```
<events>
  <category name="Filesystem">
    <event counter="filesystem_loginuid" path="/proc/self/loginuid"
      title="loginuid" name="loginuid" description="loginuid"/>
  </category>
```

```
</events>
```

Example 1-2: Configure various counters

The Instructions Executed counter is configured in slot 0 as:

```
--counters ARMv8_Cortex_A53_cnt0:0x08
```

To configure the cycle counter, specify `--counters <device>_ccnt`. For example:

```
--counters ARMv8_Cortex_A53_ccnt
```

Other counters do not have event codes and are identified only by name. For example:

```
--counters PERF_COUNT_SW_PAGE_FAULTS
```

1.7 Support for Arm CoreLink™ tools

Streamline supports various Arm®CoreLink™ interconnect parts, including: CCI-400, CCI-500, CCN, CMN-600, L2C-310, SMMUv3, and DSU.

Profiling these devices relies on support for the correct PMU driver being present and enabled in the Linux kernel. Ensure you are using a version that supports your hardware, and that you have the correct configuration options enabled in the kernel.

1.8 Using Streamline with Arm® NN

Streamline supports capturing profiling information from Arm NN enabled applications. This feature enables you to collect performance counters from supported Arm®Ethos™ NPU and performance data from the Arm NN library. To profile an Arm NN application, `gator` must first see which events are available in the running application. Streamline can then capture a profile as normal.

Before you begin

- Build your application against Arm NN version 20.08 or later. See <https://github.com/ARM-software/armnn> for more information.
- When you create the runtime in your application code, you must set `options.m_ProfilingOptions.m_EnableProfiling = true` for the options passed to `IRuntime::Create`.
- Enable timeline profiling by setting `options.m_ProfilingOptions.m_TimelineEnabled = true`.

```
// Create Arm NN runtime  
  
IRuntime::CreationOptions options; // default options  
options.m_ProfilingOptions.m_EnableProfiling = true;
```

```
options.m_ProfilingOptions.m_TimelineEnabled = true;  
IRuntimePtr run = IRuntime::Create(options);
```

Procedure

1. Start `gator` in daemon mode as described in [Set up your target](#).



Local capture mode is not available because `gator` does not know which events are available yet.

2. Run your application.
3. Configure counters and [Capture a Streamline profile](#) as described in the Arm Streamline User Guide.



To get the timeline data, you must start your application again after starting the capture.

2. Customize your Streamline report

This chapter explains how to add support for a new CPU, or an uncore peripheral outside of the CPU, to Streamline.

2.1 Add a CPU to Streamline

Customize your Streamline reports with support for custom CPU Performance Monitoring Units (PMUs).

Procedure

1. Create an `events-<processor_name>.xml` file and save it in `<DS_install_directory>/sw/streamline/gator/daemon/`.
2. Include the following in the file:

```
<counter_set name="<ID>_cnt" count="<NEVENTS>" />
<category name="<NAME>" counter_set="<ID>_cnt" per_cpu="no">
  <event .../ >
</category>
```

`<ID>` is a string that can contain letters, numbers, and underscores.

`<NEVENTS>` is the maximum number of counters in the hardware.

If the number of events is unlimited, the `<counter_set>` element and the `counter_set="<ID>_cnt"` attribute are not required.

3. Add a `<pmu>` element to `pmus.xml`.
See the comment at the top of the `pmus.xml` file for the required attributes.
4. Rebuild `gator`.

2.2 Add an uncore peripheral to Streamline

Customize your Streamline reports with support for custom uncore peripheral Performance Monitoring Units (PMUs).

Procedure

1. Create an `events-<peripheral_name>.xml` file and save it in `<DS_install_directory>/sw/streamline/gator/daemon/`.
2. Include the following in the file:

```
<counter_set name="<ID>_cnt" count="<NEVENTS>" />
<category name="<NAME>" counter_set="<ID>_cnt" per_cpu="no">
  <event .../ >
</category>
```


<ID> is a string that can contain letters, numbers, and underscores.

<NEVENTS> is the maximum number of counters in the hardware.

If the number of events is unlimited, the <counter_set> element and the counter_set="<ID>_cnt" attribute are not required.

3. Add an <uncore_pmu> element to pmus.xml.
See the comment at the top of the pmus.xml file for the required attributes.
4. Rebuild gator.

2.3 Contributing CPU and peripheral support

The gator daemon is an open-source component, licensed under the GPLv2 license, and published on GitHub. We are willing to accept pull requests from external developers to add support for custom Arm cores and supporting uncore peripherals.

The gator daemon is in this repository <https://github.com/ARM-software/gator/>, where you can also find the [contributing guidelines](#).

3. Advanced target setup information (Linux)

This appendix provides extra configuration information beyond the standard setup.

3.1 Kernel configuration menu options for Linux

You must enable certain kernel configuration options to run Streamline.

The following `menuconfig` menus have options that are required for Streamline:



- If these options are not set correctly, you must change them and rebuild your kernel. If they are set correctly, you are ready to build and install the gator driver.
- The location of these options might change between releases. If so, use the search option in `menuconfig` to find them.
- Extra options are required to enable Arm®Mali™ GPU support.

General Setup

Enable the **Profiling Support** option `CONFIG_PROFILING`, and the **Kernel performance events and counters** option `CONFIG_PERF_EVENTS`. `CONFIG_PERF_EVENTS` is required for kernel versions 3.0 and later. Enable the **Timers subsystem > High Resolution Timer Support** option `CONFIG_HIGH_RES_TIMERS`.

Kernel Features

The **Enable hardware performance counter support for perf events** option `CONFIG_HW_PERF_EVENTS`. `CONFIG_HW_PERF_EVENTS` is required for kernel versions 3.0 and later. If you are using Symmetric MultiProcessing (SMP), enable the **Use local timer interrupts** option `CONFIG_LOCAL_TIMERS`. If you are running on Linux version 3.12 or later, the `CONFIG_LOCAL_TIMERS` option is not necessary.

CPU Power Management

Optionally enable the **CPU Frequency scaling** option `CONFIG_CPU_FREQ` to enable the CPU Freq **Timeline** view chart. `gator` requires kernel version 2.6.38 or greater to enable this chart.

Kernel hacking

If other trace configuration options are enabled, the **Trace process context switches and events** option `CONFIG_ENABLE_DEFAULT_TRACERS` might not be visible in `menuconfig` as an option. Enabling one of these other trace configurations, for example `CONFIG_GENERIC_TRACER`, `CONFIG_TRACING`, or `CONFIG_CONTEXT_SWITCH_TRACER`, is sufficient to enable tracing. Optionally enable the **Compile the kernel with debug info** option `CONFIG_DEBUG_INFO`. This option is only required for profiling the Linux kernel.

**Caution**

Kernel versions before 4.6, with `CONFIG_CPU_PM` enabled, produce invalid results. For example, counters not showing any data, large spikes, and non-sensible values for counters. This issue is due to the kernel PMU driver not saving state when the processor powers down, or not restoring state when it powers up. To avoid this issue, upgrade to the latest version of the kernel, or apply the patch found at <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=da4e4f18afe0f3729d68f3785c5802f786d36e34>. This patch applies cleanly to version 4.4, and it might also be possible to back port it to other versions. If you apply the patch, you might also require the patch at <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=cbcc72e037b8a3eb1fad3c1ae22021df21c97a51>.

3.2 Profile the Linux kernel

To ensure the statistics that the kernel generates align with source code in the Analysis Reports, include the kernel in the **Program Images** section in the **Capture Settings** dialog box. Before you can include the kernel in the **Program Images** section, you must build a version of `vmlinux` with kernel debug information enabled.

Procedure

1. Navigate to the kernel build directory.
2. Enter the following command to enable you to change `menuconfig` options:

```
make ARCH=arm CROSS_COMPILE=<cross_compiler_directory>/bin/arm-linux-gnueabi-hf-menuconfig
```
3. In the **Kernel Hacking** menu, select the **Compile the kernel with debug info** option. This option enables the `CONFIG_DEBUG_INFO` kernel option.
4. Enter the following command to build a new `vmlinux` image:

```
make -j5 ARCH=arm CROSS_COMPILE=<cross_compiler_directory>/bin/arm-linux-gnueabi-hf- uImage
```

**Note**

You can profile a driver by either statically linking it into the kernel image or by adding the module as an image in the **Capture Settings** dialog box.

5. Open the **Capture Settings** dialog box.
6. Click the **Add image** button in the **Program Images** section.
7. Navigate to your `vmlinux` file and select it.
8. Click **Open**.
9. Start a new capture session.

Related information

[Kernel configuration menu options for Linux](#) on page 18

3.3 Build gatord yourself for Linux

Build `gatord` yourself to apply patches for bug fixes or add support for new features.

About this task

It is not possible to build `gatord` on a Windows host.

Procedure

1. Either download the `gatord` source from the daemon directory in <https://github.com/ARM-software/gator>, or copy the source that is supplied in `<DS_install_directory>/sw/streamline/gator/daemon/`.
2. Follow the instructions in the `README.md` file in the `gator` directory.

3.4 Add support to gatord for a new CPU or perf PMU

There are two methods for adding support for a new CPU or perf PMU to `gatord`. One that requires rebuilding `gatord`, and one that does not.

- Perf support in Linux for the PMU is required because `gatord` uses `perf` to read the hardware counters.
- Check the perf PMUs supported by your kernel by running:



Note

```
ls /sys/bus/event_source/devices/
```

If `ARMv7_Cortex_A<xx>`, `CCI_400`, or `cnn` are listed, then `A<xx>`, `CCI-400`, or `CNN` respectively are supported.

- Only XML changes are required, no code changes are necessary.
- Rebuilding `gatord` after the XML changes is recommended but not required because you can pass PMUs and events to `gatord` on the command line.

Make the following changes, then rebuild `gatord`:

1. Add a line to `<install_directory>/sw/streamline/gator/daemon/pmus.xml` describing the new PMU. For CPUs, the following information is required:
 - The CPU Implementer and Primary part number from the Main ID Register.
 - The number of generic hardware counters that can be selected simultaneously.
 - Optionally, set the perf PMU name of the CPU to ensure correct operation in multi-PMU heterogeneous CPU implementations, such as `big.LITTLE™` or `DynamiQ™`.
2. Create an events XML file, named `events-<xxx>.xml` in the `gatord` source directory that defines the events that the new PMU generates. Exclude the XML header and `<events>`

element. See the Cortex®-A15 events XML file, `<DS_install_directory>/sw/streamline/gator/daemon/events-Cortex-A15.xml` for an example.

Alternatively, to add support without having to rebuild `gator`, do the following:

1. Create an events XML file that defines the events that the new PMU generates. This file must include the XML header and `<events>` element, for example:

```
<events>
  <counter_set name="ARMv7_Cortex_A9_cnt" count="6"/>
  <category name="Cortex-A9" counter_set="ARMv7_Cortex_A9_cnt" per_cpu="yes"
    supports_event_based_sampling="yes">
    <event counter="ARMv7_Cortex_A9_ccnt" event="0xff" title="Clock"
      name="Cycles" display="hertz" units="Hz" \
        average_selection="yes" average_cores="yes" description="The number of core
        clock cycles"/>
    <event event="0x00" title="Software" name="Increment"
      description="Incremented only on writes to the \
        Software Increment Register"/>
    <event event="0x01" title="Cache" name="Instruction refill"
      description="Instruction fetch that causes \
        a refill of at least the level of instruction or unified cache closest to
        the processor"/>

    <!-- ... -->

  </category>
</events>
```

2. Create an XML file that defines information about the new PMU. For the required format, see the `gator pmus.xml`. For example:

```
<pmus>
  <pmu pmnc_name="ARMv7_Cortex_A9" cpuid="0x41c09" core_name="Cortex-A9"
    pmnc_counters="6"/>
</pmus>
```

3. Copy these files to the target and restart `gator` using the following flags:
 - `-E` to specify the location of the events XML file. This flag causes the events to be appended to the list of events that `gator` supports.
 - `-P` to specify the location of the PMU XML file. This option causes the new PMU to be added to the list of PMUs defined in `pmus.xml` that `gator` has built-in support for.

Related information

[Build gator yourself for Linux](#) on page 20

4. Analytics data collection

Arm periodically collects anonymous information to understand how you use our tools and their features. We use this information to improve your experience with our tools. We do not collect any personally identifiable information.

Analytics event data

This section describes the types of analytics event data that we collect.

Host system information

We collect the following host system information to track product usage and manage our support policy:

host_os_name

The name of the host OS.

host_os_version

The version of the host OS.

host_os_cpu

The CPU architecture of the host. For example, x86-64, or arm64.

Target information

We collect the following information to track product usage and manage our support policy:

target_os_name

The architecture of the target. For example, Windows, Linux, or Android.

Tool version information

We collect the following information from the tool:

tool_name

The name of the tool.

tool_version

The version of the tool.

Tool usage information

We collect the following information to track product usage:

user_id

An opaque ID generated by hashing your user name and host machine name.

session_id

An opaque ID generated by hashing your `user_id`, and the current date and time. This information also helps us to understand how you use features together.

Tool feature usage information

We collect the following information about the tool features:

capture_time_*

Reports how long the tool takes to complete different stages of data capture and analysis. We use this information to prioritize performance improvements in the tool roadmap.

feature_gpu_api_*

Reports which graphics API your target application is using. We use this information to prioritize API support in the tool roadmap.

Disabling analytics

You can disable analytics by using any of the following methods:

- Set the `ARM_DISABLE_ANALYTICS` environment variable to any value, including an empty string.
- Pass the `--disable_analytics` command-line option to any tools that emit analytics.
- In the Streamline menu, click **Window -> Preferences -> Streamline -> Confirmations**, and clear the **Allow collection of anonymous analytics data** checkbox.

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
0907-03	11 December 2025	Non-Confidential	Updated document for v9.7
0907-00	31 July 2025	Non-Confidential	New document for v9.7
0906-00	1 May 2025	Non-Confidential	New document for v9.6
0905-01	20 March 2025	Non-Confidential	Updated document for v9.5
0905-00	6 February 2025	Non-Confidential	New document for v9.5
0904-00	28 November 2024	Non-Confidential	New document for v9.4
0903-00	5 September 2024	Non-Confidential	New document for v9.3
0902-00	7 June 2024	Non-Confidential	New document for v9.2
0901-00	12 April 2024	Non-Confidential	New document for v9.1
0900-00	15 February 2024	Non-Confidential	New document for v9.0

Issue	Date	Confidentiality	Change
0809-00	23 November 2023	Non-Confidential	New document for v8.9
0808-00	28 September 2023	Non-Confidential	New document for v8.8
0807-00	3 August 2023	Non-Confidential	New document for v8.7
0806-00	8 June 2023	Non-Confidential	New document for v8.6
0805-00	20 April 2023	Non-Confidential	New document for v8.5
0804-00	14 February 2023	Non-Confidential	New document for v8.4
0803-00	17 November 2022	Non-Confidential	New document for v8.3
0802-00	19 August 2022	Non-Confidential	New document for v8.2
0801-00	20 May 2022	Non-Confidential	New document for v8.1
0800-00	18 February 2022	Non-Confidential	New document for v8.0
0709-00	18 November 2021	Non-Confidential	New document for v7.9
0708-00	20 August 2021	Non-Confidential	New document for v7.8

Change history

For information about the functional changes to Streamline, see the [Arm Performance Studio Release Notes](#).

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.

Convention	Use
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 <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: <div>MRC p15, 0, <Rd>, <CRn>, <CRm>, <Opcode_2></div>
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 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.



Remember

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

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 following tables provides direct access to the online version of the document.

Arm product resources	Document ID	Confidentiality
Arm Development Studio Getting Started Guide	101469	Non-Confidential
Arm Streamline User Guide	101816	Non-Confidential