



# Installation Guide

## Nsight Systems Installation Guide

# Table of Contents

<b>Chapter 1. Overview.....</b>	<b>1</b>
<b>Chapter 2. System Requirements.....</b>	<b>3</b>
Supported Platforms.....	3
CUDA Version.....	4
Requirements for Android-Based Devices.....	4
Requirements for x86_64, Power, and ARM SBSA Targets on Linux.....	5
x86_64 Windows Target Device Requirements.....	6
Host Application Requirements.....	6
<b>Chapter 3. Getting Started Guide.....</b>	<b>7</b>
3.1. Finding the Right Package.....	7
3.2. Installing GUI on the Host System.....	8
3.3. Optional: Setting up the CLI.....	8
3.4. Launching the GUI.....	9

---

# Chapter 1. Overview

**Nsight Systems** is a statistical sampling profiler with tracing features. It is designed to work with devices and devkits based on NVIDIA Tegra SoCs (system-on-chip), ARM SBSA (server based system architecture) systems, IBM Power systems, and systems based on the x86\_64 processor architecture that also include NVIDIA GPU(s).

Throughout this document we will refer to the device on which profiling happens as the **target**, and the computer on which the user works and controls the profiling session as the **host**. Note that for x86\_64 based systems these may be on the same device, whereas with Tegra, ARM, or IBM Power based systems they will always be separate.

Furthermore, three different activities are distinguished as follows:

- ▶ **Profiling** — The process of collecting any performance data. A profiling session in Nsight Systems typically includes sampling and tracing.
- ▶ **Sampling** — The process of periodically stopping the *profilee* (the application under investigation during the profiling session), typically to collect backtraces (call stacks of active threads), which allows you to understand statistically how much time is spent in each function. Additionally, hardware counters can also be sampled. This process is inherently imprecise when a low number of samples have been collected.
- ▶ **Tracing** — The process of collecting precise information about various activities happening in the profilee or in the system. For example, profilee API execution may be traced providing the exact time and duration of a function call.

Nsight Systems supports multiple generations of Tegra SoCs, NVIDIA discrete GPUs, and various CPU architectures, as well as various target and host operating systems. This documentation describes the full set of features available in any version of Nsight Systems. In the event that a feature is not available in all versions, that will be noted in the text. In general, Nsight Systems Embedded Platforms Edition indicates the package that supports Tegra processors for the embedded and automotive market and Nsight Systems Workstation Edition supports x86\_64, IBM Power, and ARM server (SBSA) processors for the workstation and cluster market.

Common features that are supported by Nsight Systems on most platforms include the following:

- ▶ Sampling of the profilee and collecting backtraces using multiple algorithms (such as frame pointers or DWARF data). Building top-down, bottom-up, and flat views as appropriate. This information helps identify performance bottlenecks in CPU-intensive code.

- ▶ Sampling or tracing system power behaviors, such as CPU frequency.
- ▶ (Only on Nsight Systems Embedded Platforms Edition) Sampling counters from ARM PMU (Performance Monitoring Unit). Information such as cache misses gets statistically correlated with function execution.
- ▶ Support for multiple windows. Users with multiple monitors can see multiple reports simultaneously, or have multiple views into the same report file.

With Nsight Systems, a user could:

- ▶ Identify call paths that monopolize the CPU.
- ▶ Identify individual functions that monopolize the CPU (across different call paths).
- ▶ For Nsight Systems Embedded Platforms Edition, identify functions that have poor cache utilization.
- ▶ If platform supports CUDA, see visual representation of CUDA Runtime and Driver API calls, as well as CUDA GPU workload. Nsight Systems uses the CUDA Profiling Tools Interface (CUPTI), for more information, see: [CUPTI documentation](#).
- ▶ If the user annotates with NVIDIA Tools Extension (NVTX), see visual representation of NVTX annotations: ranges, markers, and thread names.
- ▶ For Windows targets, see visual representation of D3D12: which API calls are being made on the CPU, graphic frames, stutter analysis, as well as GPU workloads (command lists and debug ranges).
- ▶ For x86\_64 targets, see visual representation of Vulkan: which API calls are being made on the CPU, graphic frames, stutter analysis, as well as Vulkan GPU workloads (command buffers and debug ranges).

---

# Chapter 2. System Requirements

Nsight Systems supports multiple platforms. For simplicity, think of these as Nsight Systems Embedded Platforms Edition and Nsight Systems Workstation Edition, where Nsight Systems Workstation Edition supports desktops, workstations, and clusters with x86\_64, IBM Power, and ARM SBSA CPUs on Linux and Windows OSs, while Nsight Systems Embedded Platforms Edition supports NVIDIA Tegra products for the embedded and gaming space on Android, Linux for Tegra, and QNX OSs.

## Supported Platforms

Depending on your OS, different GPUs are supported.

### Android

- ▶ NVIDIA SHIELD Android TV
- ▶ NVIDIA SHIELD Tablet
- ▶ Various devkits

### L4T (Linux for Tegra)

- ▶ Jetson AGX Xavier
- ▶ Jetson TX2
- ▶ Jetson TX2i
- ▶ Jetson TX
- ▶ Jetson Nano

### x86\_64, IBM Power (from Power 9), or ARM SBSA

- ▶ NVIDIA GPU architectures starting with Pascal
- ▶ OS (64 bit only)
  - ▶ Ubuntu 14.04, 16.04, and 18.04
  - ▶ CentOS and RedHat Enterprise Linux 7.4+ with kernel version 3.10.0-693 or later.
  - ▶ Windows 10

## CUDA Version

- ▶ Nsight Systems supports CUDA 10.0, 10.1, 10.2, 11.0, and 11.1 for most platforms
- ▶ Nsight Systems on ARM SBSA supports 10.2, 11.0, and 11.1

Note that CUDA version and driver version must be compatible.

CUDA Version	Driver minimum version
11.0	450
10.2	440.30
10.1	418.39
10.0	410.48

## Requirements for Android-Based Devices

To use Nsight Systems with Android-based target devices, you must first:

1. Have a compatible Android device.
2. Have a compatible Android OS image.

If installing from CODEWORKS, note that devices from NVIDIA typically ship with support, while OEMs or ODMs producing their own Nsight Systems choose whether or not to support this tool.

To check for the version number of the kernel support of Nsight Systems on a target device, run the following command from ADB shell:

```
cat /proc/quadd/version
```

Minimal supported version is 1.82.

3. Install the Android SDK Platform tools.
4. Install the ADB USB Driver.
5. Have your target application compiled with the correct flags and packaged with the correct permissions.

Items 3-4 can be handled automatically for you by installing the latest NVIDIA CodeWorks for Android. Item 2 may already be on your device, if it is a retail device starting with the Tegra 4 (minimum), but not all devices are required to come with OS images that support this tool. Item 2 can also be found for devkits by downloading the latest OS Image for your device. CodeWorks for Android can be found at: [NVIDIA CodeWorks SDK](#).

Using CodeWorks for Android is the recommended path for getting your Android device and system prepared to use Nsight Systems.

## Requirements for x86\_64, Power, and ARM SBSA Targets on Linux

When attaching to x86\_64, Power, or ARM SBSA Linux-based target from the GUI on the host, the connection is established through SSH.

**Use of Linux Perf:** To collect thread scheduling data and IP (instruction pointer) samples, the Perf paranoid level on the target system must be 2 or less. Use the following command to check:

```
cat /proc/sys/kernel/perf_event_paranoid
```

If the output is >2, then do the following to temporarily adjust the paranoid level (note that this has to be done after each reboot):

```
sudo sh -c 'echo 2 >/proc/sys/kernel/perf_event_paranoid'
```

To make the change permanent, use the following command:

```
sudo sh -c 'echo kernel.perf_event_paranoid=2 > /etc/sysctl.d/local.conf'
```

**Kernel version:** To collect thread scheduling data and IP (instruction pointer) samples and backtraces, the kernel version must be:

- ▶ 3.10.0-693 or later for CentOS and RedHat Enterprise Linux 7.4+
- ▶ 4.3 or greater for all other distros including Ubuntu

To check the version number of the kernel on a target device, run the following command on the device:

```
uname -a
```

Note that only CentOS, RedHat, and Ubuntu distros are tested/confirmed to work correctly.

**glibc version:** To check the glibc version on a target device, run the following command:

```
ldd --version
```

Nsight Systems requires glibc 2.14 or more recent.

**CUDA:** See above for supported CUDA versions in this release. Use the deviceQuery command to determine the CUDA driver and runtime versions on the system. the deviceQuery command is available in the CUDA SDK. It is normally installed at:

```
/usr/local/cuda/samples/1_Utilities/deviceQuery
```

Only pure 64-bit environments are supported. In other words, 32-bit systems or 32-bit processes running within a 64-bit environment are not supported.

Nsight Systems requires write permission to the ``/var/lock`` directory on the target system.

**Docker:** See Collecting Data within a Docker section of [Profiling in a Docker on Linux Devices](#) for more information.

## x86\_64 Windows Target Device Requirements

### DX12 Requires:

- ▶ Windows 10 with NVIDIA Driver 411.63 or higher for DX12 trace
- ▶ Windows 10 April 2018 Update (version 1803, AKA Redstone 4) with NVIDIA Driver 411.63 or higher for DirectX Ray Tracing, and tracing DX12 Copy command queues.

## Host Application Requirements

The Nsight Systems host application runs on the following host platforms:

- ▶ Windows 7 and higher, Windows Server 2008 R2 and higher. Only 64-bit versions are supported.
- ▶ Linux Ubuntu 14.04 and higher are known to work, running on other modern distributions should be possible as well. Only 64-bit versions are supported.
- ▶ OS X 10.10 "Yosemite" and higher.



---

# Chapter 3. Getting Started Guide

## 3.1. Finding the Right Package

Nsight Systems is available for multiple targets and multiple host OSs. To choose the right package, first consider the target system to be analyzed.

- ▶ For Tegra target systems, select Nsight Systems for Tegra available as part of [NVIDIA JetPack SDK](#) and [NVIDIA CodeWorks SDK](#).
- ▶ For x86\_64, IBM Power target systems, or ARM SDSA select from the target packages from Nsight Systems for Workstations, available from <https://developer.nvidia.com/nsight-systems>. This web release will always contain the latest and greatest Nsight Systems features.
- ▶ The x86\_64, IBM Power, and ARM SDSA target versions of Nsight Systems are also available in the [CUDA Toolkit](#).

Each package is limited to one architecture. For example, Tegra packages do not contain support for profiling x86 targets, and x86 packages do not contain support for profiling Tegra targets.

After choosing an appropriate target version, select the package corresponding to the host OS, the OS on the system where results will be viewed. These packages are in the form of common installer types: .msi for Windows; .run, .rpm, and .deb for x86 Linux; .deb and .rpm for Linux on IBM Power; and .dmg for the MacOS installer.

Note: the IBM Power and ARM SBSA packages do not have a GUI for visualization of the result. If you wish to visualize your result, please download and install the GUI available for MacOS, x86\_64 Linux, or Windows systems.

### Tegra packages

- ▶ Windows host – Install .msi on Windows machine. Enables remote access to Tegra device for profiling.
- ▶ Linux host – Install .run on Linux system. Enables remote access to Tegra device for profiling.
- ▶ MacOS host – Install .dmg on MacOS machine. Enables remote access to Tegra device for profiling.

### x86\_64 packages

- ▶ Windows host – Install .msi on Windows machine. Enables remote access to Linux x86\_64 or Windows devices for profiling as well as running on local system.
- ▶ Linux host – Install .run, .rpm, or .deb on Linux system. Enables remote access to Linux x86\_64 or Windows devices for profiling or running collection on localhost.
- ▶ Linux CLI only – The Linux CLI is shipped in all x86 packages, but if you just want the CLI, we have a package for that. Install .deb on Linux system. Enables only CLI collection, report can be imported or opened in x86\_64 host.
- ▶ MacOS host – Install .dmg on MacOS machine. Enables remote access to Linux x86\_64 device for profiling.

### IBM Power packages

- ▶ Power CLI only - The IBM Power support does not include a host GUI. Install .deb or .rpm on your Power system. Enables only CLI collection, report can be imported or opened in GUI on any supported host platform.

### ARM SBSA packages

- ▶ ARM SBSA CLI only - ARM SBSA support does not include a host GUI. Install .deb or .rpm on your ARM SBSA system. Enables only CLI collection, report can be imported or opened in GUI on any supported host platform.

## 3.2. Installing GUI on the Host System

Copy the appropriate file to your host system in a directory where you have write and execute permissions. Run the install file, accept the EULA, and Nsight Systems will install on your system.

On Linux, there are special options to enable automated installation. Running the installer with the `--accept` flag will automatically accept the EULA, running with the `--accept` flag and the `--quiet` flag will automatically accept the EULA without printing to stdout. Running with `--quiet` without `--accept` will display an error.

The installation will create a Host directory for this host and a Target directory for each target this Nsight Systems package supports.

All binaries needed to collect data on a target device will be installed on the target by the host on first connection to the device. There is no need to install the package on the target device.

If installing from the CUDA Toolkit, see the [CUDA Toolkit documentation](#).

## 3.3. Optional: Setting up the CLI

All Nsight Systems targets can be profiled using the CLI. IBM Power and ARM SBSA targets can only be profiled using the CLI. The CLI is especially helpful when scripts are used to run unattended collections or when access to the target system via ssh is not possible. In particular, this can be used to enable collection in a Docker container.

The CLI can be found in the Target directory of the Nsight Systems installation. Users who want to install the CLI as a standalone tool can do so by copying the files within the Target directory to the location of their choice.

If you wish to run the CLI without root (recommended mode) you will want to install in a directory where you have full access.

Once you have the CLI set up, you can use the `nsys status -e` command to check your environment.

```
~$ nsys status -e

Sampling Environment Check
Linux Kernel Paranoid Level = 1: OK
Linux Distribution = Ubuntu
Linux Kernel Version = 4.15.0-109-generic: OK
Linux perf_event_open syscall available: OK
Sampling trigger event available: OK
Intel(c) Last Branch Record support: Available
Sampling Environment: OK
```

This status check allows you to ensure that the system requirements for CPU sampling using Nsight Systems are met in your local environment. If the Sampling Environment is not OK, you will still be able to run various trace operations.

Intel(c) Last Branch Record allows tools, including Nsight Systems to use hardware to quickly get limited stack information. Nsight Systems will use this method for stack resolution by default if available.

For information about changing these environment settings, see System Requirements section in the Installation Guide. For information about changing the backtrace method, see Profiling from the CLI in the User Guide.

To get started using the CLI, run `nsys --help` for a list of options or see Profiling Applications from the CLI in the User Guide for full documentation.

## 3.4. Launching the GUI

### Installed from Product Download Page

Depending on your OS, Nsight Systems will have installed an icon on your host desktop that you can use to launch the GUI. To launch the GUI directly, run the `nsight-sys` executable in the Host sub-directory of your installation.

### Installed as Part of the CUDA Toolkit

Nsight Systems is installed by the CUDA Toolkit installer in `/usr/local/cuda-[version]/NsightSystems-[version]`. To launch the GUI, run the `nsight-sys` executable in the **Host** sub-directory there.

In a future version of the CUDA Toolkit, this installation will be moved under the `CUDA/bin` directory.