



SANITIZER API

v2020.3.0 | November 2020

Reference Manual



TABLE OF CONTENTS

Chapter 1. Introduction.....	1
1.1. Overview.....	1
Chapter 2. Usage.....	2
2.1. Compatibility and Requirements.....	2
2.2. Callback API.....	2
2.2.1. Driver and Runtime API Callbacks.....	3
2.2.2. Resource Callbacks.....	4
2.2.3. Synchronization Callbacks.....	4
2.2.4. Launch Callbacks.....	5
2.2.5. Memcpy Callbacks.....	5
2.2.6. Memset Callbacks.....	5
2.2.7. Batch Memory Operations Callbacks.....	5
2.3. Patching API.....	5
2.3.1. Writing a Patch.....	6
2.3.2. Insert a Patch.....	6
2.4. Memory API.....	7
Chapter 3. Limitations.....	8
Appendix A. Copyright and Licenses.....	9

Chapter 1.

INTRODUCTION

1.1. Overview

The Compute Sanitizer API enables the creation of sanitizing and tracing tools that target CUDA applications. Examples of such tools are memory and race condition checkers. The Compute Sanitizer API is composed of three APIs: the callback API, the patching API and the memory API. It is delivered as a dynamic library on supported platforms.

Chapter 2.

USAGE

2.1. Compatibility and Requirements

The Compute Sanitizer API requires CUDA 10.1 or newer. Compute Sanitizer API calls will fail with `SANITIZER_ERROR_NOT_INITIALIZED` if the CUDA driver version is not compatible with the Compute Sanitizer version.

2.2. Callback API

The Compute Sanitizer Callback API allows you to register a callback into user code. The callback is invoked when the application calls a CUDA runtime or driver function, or when certain events occur in the CUDA driver. The following terminology is used by the Callback API.

- ▶ **Callback domain:** Callbacks are grouped into domains to make it easier to associate callback functions with groups of related CUDA functions or events. The following callback domains are defined by `Sanitizer_CallbackDomain`.
 1. CUDA driver functions
 2. CUDA runtime functions
 3. CUDA resource tracking
 4. CUDA synchronization notification
 5. CUDA grid launches
 6. CUDA memcpy operations
 7. CUDA memset operations
 8. CUDA batch memory operations
- ▶ **Callback ID:** Each callback is given a unique ID within the corresponding callback domain in order to identify it within the callback function. The CUDA driver API IDs are defined in `sanitizer_driver_cbid.h` and the CUDA runtime API IDs are defined in `sanitizer_runtime_cbid.h`. Other callback IDs are defined in `sanitizer_callbacks.h`. All of these headers are included as part of `sanitizer.h`.

- ▶ **Callback Function:** The callback function must be of the type `Sanitizer_CallbackFunc`. This function type has two arguments that specify the callback: the domain and the ID that identifies why the callback is occurring. The type also has a `cbdata` argument that is used to pass data specific to the callback.
- ▶ **Subscriber:** A subscriber is used to associate each of the callback functions with one or more CUDA API functions. There can be at most one subscriber initialized with `sanitizerSubscribe` at any time. Before initializing a new subscriber, the existing one must be finalized with `sanitizerUnsubscribe`.

Each callback domain is described in detail below. Unless explicitly stated, it is not supported to call any CUDA runtime or driver API from within a callback function. Doing so may cause the application to hang. However, it is supported to call [Compute Sanitizer Memory APIs](#) from within callback functions.

2.2.1. Driver and Runtime API Callbacks

Using the Callback API with the `SANITIZER_CB_DOMAIN_DRIVER_API` or `SANITIZER_CB_DOMAIN_RUNTIME_API` domains, a callback function can be associated with one or more CUDA API functions. When those CUDA functions are called in the application, the callback function is invoked as well. For these domains, the `cbdata` argument to the callback function will be of the type `Sanitizer_CallbackData`.

You can call `cudaDeviceSynchronize`, `cudaStreamSynchronize`, `cuCtxSynchronize` and `cuStreamSynchronize` from within a driver or runtime API callback function.

The following code shows a typical sequence used to associate a callback function with one or more CUDA API functions. For simplicity, error checking code was removed.

```
Sanitizer_SubscriberHandle handle;
MyDataStruct *my_data = ...;
...
sanitizerSubscribe(&handle, my_callback, my_data);
sanitizerEnableDomain(1, handle, SANITIZER_CB_DOMAIN_RUNTIME_API);
```

First, `sanitizerSubscribe` is used to initialize a subscriber with the `my_callback` callback function. Next, `sanitizerEnableDomain` is used to associate that callback with all the CUDA runtime functions. Using this code sequence will cause `my_callback` to be called twice each time any of the CUDA runtime API functions are invoked, once on entry to the CUDA function and once just before the CUDA function exits. Compute Sanitizer callback API functions `sanitizerEnableCallback` and `sanitizerEnableAllDomains` can also be used to associate CUDA API functions with a callback.

The following code shows a typical callback function.

```
void SANITIZERAPI
my_callback(void *userdata,
            Sanitizer_CallbackDomain domain,
            Sanitizer_CallbackId cbid,
            const void *cbdata)
{
    const Sanitizer_CallbackData *cbInfo = (Sanitizer_CallbackData *)cbdata;
    MyDataStruct *my_data = (MyDataStruct *)userdata;

    if ((domain == SANITIZER_CB_DOMAIN_RUNTIME_API) &&
        (cbid == SANITIZER_RUNTIME_TRACE_CBID_cudaMemcpy_v3020) &&
        (cbInfo->callbackSite == SANITIZER_API_ENTER))
    {
        cudaMemcpy_v3020_params *funcParams = (cudaMemcpy_v3020_params *)
(cbInfo->functionParams);
        size_t count = funcParams->count;
        enum cudaMemcpyKind kind = funcParams->kind
        ...
    }
    ...
}
```

In the callback function, **Sanitizer_CallbackDomain** and **Sanitizer_CallbackId** parameters can be used to determine which CUDA API function invocation is triggering this callback. In the example above, we are checking for the CUDA runtime **cudaMemcpy** function. The **cbdata** parameter holds a structure of useful information that can be used within the callback. In this case, we use the **callbackSite** member of the structure to detect that the callback is occurring on entry to **cudaMemcpy**, and we use the **functionParams** member to access the parameters to **cudaMemcpy**. To access the parameters, we first cast **functionParams** to a structure type corresponding to the **cudaMemcpy** function. These parameter structures are contained in **generated_cuda_runtime_api_meta.h**, **generated_cuda_meta.h** and a number of other files.

2.2.2. Resource Callbacks

Using the Callback API with the **SANITIZER_CB_DOMAIN_RESOURCE** domain, a callback function can be associated with some CUDA resource creation and destruction events. For example, when a CUDA context is created, the callback function is invoked with a callback ID equal to **SANITIZER_CBID_RESOURCE_CONTEXT_CREATED**. For this domain, the **cbdata** argument is one of the following types:

- ▶ **Sanitizer_ResourceContextData** for CUDA context creation and destruction
- ▶ **Sanitizer_ResourceStreamData** for CUDA stream creation and destruction
- ▶ **Sanitizer_ResourceModuleData** for CUDA module load and unload
- ▶ **Sanitizer_ResourceMemoryData** for CUDA memory allocation and de-allocation

2.2.3. Synchronization Callbacks

Using the Callback API with the **SANITIZER_CB_DOMAIN_SYNCHRONIZE** domain, a callback function can be associated with CUDA context and stream synchronizations. For example, when a CUDA context is synchronized, the callback function is invoked

with a callback ID equal to `SANITIZER_CBID_SYNCHRONIZE_CONTEXT_SYNCHRONIZED`. For this domain, the `cbdata` argument is of the type `Sanitizer_SynchronizeData`.

2.2.4. Launch Callbacks

Using the Callback API with the `SANITIZER_CB_DOMAIN_LAUNCH` domain, a callback function can be associated with CUDA kernel launches. For example, when a CUDA kernel launch has started, the callback function is invoked with a callback ID equal to `SANITIZER_CBID_LAUNCH_BEGIN`. For this domain, the `cbdata` argument is of the type `Sanitizer_LaunchData`.

2.2.5. Memcpy Callbacks

Using the Callback API with the `SANITIZER_CB_DOMAIN_MEMCPY` domain, a callback function can be associated with CUDA memcpy operations. For example, when a `cudaMemcpy` API function is called, the callback function is invoked with a callback ID equal to `SANITIZER_CBID_MEMCPY_STARTING`. For this domain, the `cbdata` argument is of the type `Sanitizer_MemcpyData`.

2.2.6. Memset Callbacks

Using the Callback API with the `SANITIZER_CB_DOMAIN_MEMSET` domain, a callback function can be associated with CUDA memset operations. For example, when a `cudaMemset` API function is called, the callback function is invoked with a callback ID equal to `SANITIZER_CBID_MEMSET_STARTING`. For this domain, the `cbdata` argument is of the type `Sanitizer_MemsetData`.

2.2.7. Batch Memory Operations Callbacks

Using the Callback API with the `SANITIZER_CB_DOMAIN_BATCH_MEMOP` domain, a callback function can be associated with CUDA batch memory operations. For example, when a `cuStreamWriteValue` API function is called, the callback function is invoked with a callback ID equal to `SANITIZER_CBID_BATCH_MEMOP_WRITE`. For this domain, the `cbdata` argument is of the type `Sanitizer_BatchMemopData`.

2.3. Patching API

The Compute Sanitizer Patching API allows you to load patch functions and insert them into user code. Patch functions will be invoked when the application's CUDA code executes certain instructions or calls certain CUDA device functions. The following terminology is used by the Patching API:

- ▶ **Instruction ID:** Each patchable event is given a unique ID than can be passed to patching API functions to specify that these events should be patched. Instruction IDs are defined by `Sanitizer_InstructionId`.
- ▶ **Instrumentation point:** A location in the original CUDA code that is being instrumented by the Compute Sanitizer API. Upon execution, the user code path is

modified so that a patch gets executed either before or after the patched event. All patches are executed prior to the event, with the exception of device-side malloc.

- ▶ **Patch:** A CUDA `__device__` function that the Compute Sanitizer will insert into another existing CUDA code. Patch function signatures must match the one expected by the API (see below for the expected signature types).

2.3.1. Writing a Patch

The patch must follow the function signature required by the Compute Sanitizer API for a given instruction ID. The mapping of instruction ID to function signature is documented in the comments of `Sanitizer_InstructionId` in `sanitizer_patching.h`. For instance, if we wish to patch a memory access using the instruction ID `SANITIZER_INSTRUCTION_MEMORY_ACCESS`, we need to use the `SanitizerCallbackMemoryAccess` type.

```
extern "C" __device__
SanitizerPatchResult SANITIZERAPI my_memory_access_callback(
    void* userdata,
    uint64_t pc,
    void* ptr,
    uint32_t accessSize,
    uint32_t flags)
{
    MyDeviceDataStruct *my_data = (MyDeviceDataStruct *)userdata

    if ((flags & SANITIZER_MEMORY_DEVICE_FLAG_WRITE) != 0)
        // log write
    else
        // log read

    return SANITIZER_PATCH_SUCCESS;
}
```

In this patch, we log write and read accesses to a structure we allocated previously. `extern "C"` ensures that the patch name will not be mangled, allowing us to use its name as a string directly in calls to `sanitizerPatchInstructions` (see [below](#)).

There can be multiple patches defined in a single CUDA file. This file must then be compiled using the following nvcc options:

```
$ nvcc --cubin --keep-device-functions -Xptxas --compile-as-tools-patch
  MySanitizerPatches.cu -o MySanitizerPatches.cubin
```

The `--cubin` option can be replaced by `--fatbin` if a fatbin is preferred over a cubin as the output file.

2.3.2. Insert a Patch

Once the patch has been generated, it can be inserted into user code by using the following procedure:

1. **Load the patch.** There are two APIs used to load the patch: `sanitizerAddPatchesFromFile` and `sanitizerAddPatches`. They use the same input format as `cuModuleLoad` and `cuModuleLoadData`, respectively.

2. **Select which instructions to patch** by using the `sanitizerPatchInstructions` API.
3. **Patch user code** by using the `sanitizerPatchModule` API.
4. **Optionally, set the callback data for patches** by using the `sanitizerSetCallbackData` API.

The following code shows a typical sequence using these APIs. For simplicity, error checking was removed.

```
CUcontext ctx = ... // current CUDA context
sanitizerAddPatchesFromFile("MySanitizerPatches.cubin", ctx);

CUmodule module = ... // module containing the user code
sanitizerPatchInstructions(SANITIZER_INSTRUCTION_MEMORY_ACCESS, module,
    "my_memory_access_callback");

sanitizerPatchModule(module);

MyDeviceDataTracker *deviceDataTracker;
cudaMalloc(&deviceDataTracker, sizeof(*deviceDataTracker));

CUfunction function = ... // kernel to be launched for which we want to set the
    callbackdata for the patches
sanitizerSetCallbackData(function, deviceDataTracker);
```

All subsequent launches using code from this CUDA module will be instrumented and `my_memory_access_callback` will be invoked before every memory access. However, the callback data is only set for all subsequent launches of the given kernel. An easy way to have a kernel `CUfunction`, is through the Sanitizer launch callbacks. Instrumentation can be removed by using the `sanitizerUnpatchModule` API.

2.4. Memory API

The Compute Sanitizer Memory API provides replacement functions for the CUDA Memory API that can be safely called from within Compute Sanitizer [callbacks](#).

- ▶ `sanitizerAlloc` is a replacement for `cudaMalloc`.
- ▶ `sanitizerFree` is a replacement for `cudaFree`.
- ▶ `sanitizerMemcpyHostToDeviceAsync` is a replacement for `cudaMemcpyAsync` for host-to-device copies.
- ▶ `sanitizerMemcpyDeviceToHost` is a replacement for `cudaMemcpy` for device-to-host copies.
- ▶ `sanitizerMemset` is a replacement for `cudaMemset`.

These functions can also be called in normal user code, where they can be mixed with the CUDA API. For instance, memory allocated with `sanitizerAlloc` can be freed with `cudaFree`. However, since only CUDA API calls will cause [callbacks](#) to be invoked, this can lead to an incoherent tracking state and should be avoided.

When called from a launch domain callback, `sanitizerMemcpyDeviceToHost` may only be used with host pinned memory as destination.

Chapter 3. LIMITATIONS

No known issues at this time.

Appendix A.

COPYRIGHT AND LICENSES

NVIDIA CORPORATION

NVIDIA SOFTWARE LICENSE AGREEMENT

IMPORTANT - READ BEFORE DOWNLOADING, INSTALLING, COPYING OR USING THE LICENSED SOFTWARE

This Software License Agreement ("SLA"), made and entered into as of the time and date of click through action ("Effective Date"), is a legal agreement between you and NVIDIA Corporation ("NVIDIA") and governs the use of the NVIDIA computer software and the documentation made available for use with such NVIDIA software. By downloading, installing, copying, or otherwise using the NVIDIA software and/or documentation, you agree to be bound by the terms of this SLA. If you do not agree to the terms of this SLA, do not download, install, copy or use the NVIDIA software or documentation. IF YOU ARE ENTERING INTO THIS SLA ON BEHALF OF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE LEGAL AUTHORITY TO BIND THE ENTITY TO THIS SLA, IN WHICH CASE "YOU" WILL MEAN THE ENTITY YOU REPRESENT. IF YOU DON'T HAVE SUCH AUTHORITY, OR IF YOU DON'T ACCEPT ALL THE TERMS AND CONDITIONS OF THIS SLA, THEN NVIDIA DOES NOT AGREE TO LICENSE THE LICENSED SOFTWARE TO YOU, AND YOU MAY NOT DOWNLOAD, INSTALL, COPY OR USE IT.

1. LICENSE.

1.1 License Grant. Subject to the terms of the AGREEMENT, NVIDIA hereby grants you a non-exclusive, non-transferable license, without the right to sublicense (except as expressly set forth in a Supplement), during the applicable license term unless earlier terminated as provided below, to have Authorized Users install and use the Software, including modifications (if expressly permitted in a Supplement), in accordance with the Documentation. You are only licensed to activate and use Licensed Software for which you have a valid license, even if during the download or installation you are presented with other product options. No Orders are binding on NVIDIA until accepted by NVIDIA. Your Orders are subject to the AGREEMENT.

SLA Supplements: Certain Licensed Software licensed under this SLA may be subject to additional terms and conditions that will be presented to you in a Supplement for acceptance prior to the delivery of such Licensed Software under this SLA and the

applicable Supplement. Licensed Software will only be delivered to you upon your acceptance of all applicable terms.

1.2 Limited Purpose Licenses. If your license is provided for one of the purposes indicated below, then notwithstanding contrary terms in Section 1.1 or in a Supplement, such licenses are for internal use and do not include any right or license to sub-license and distribute the Licensed Software or its output in any way in any public release, however limited, and/or in any manner that provides third parties with use of or access to the Licensed Software or its functionality or output, including (but not limited to) external alpha or beta testing or development phases. Further:

(i) Evaluation License. You may use evaluation licenses solely for your internal evaluation of the Licensed Software for broader adoption within your Enterprise or in connection with a NVIDIA product purchase decision, and such licenses have an expiration date as indicated by NVIDIA in its sole discretion (or ninety days from the date of download if no other duration is indicated).

(ii) Educational/Academic License. You may use educational/academic licenses solely for educational purposes and all users must be enrolled or employed by an academic institution. If you do not meet NVIDIA's academic program requirements for educational institutions, you have no rights under this license.

(iii) Test/Development License. You may use test/development licenses solely for your internal development, testing and/or debugging of your software applications or for interoperability testing with the Licensed Software, and such licenses have an expiration date as indicated by NVIDIA in its sole discretion (or one year from the date of download if no other duration is indicated).

1.3 Pre-Release Licenses. With respect to alpha, beta, preview, and other pre-release Software and Documentation ("Pre-Release Licensed Software") delivered to you under the AGREEMENT you acknowledge and agree that such Pre-Release Licensed Software (i) may not be fully functional, may contain errors or design flaws, and may have reduced or different security, privacy, accessibility, availability, and reliability standards relative to commercially provided NVIDIA software and documentation, and (ii) use of such Pre-Release Licensed Software may result in unexpected results, loss of data, project delays or other unpredictable damage or loss. THEREFORE, PRE-RELEASE LICENSED SOFTWARE IS NOT INTENDED FOR USE, AND SHOULD NOT BE USED, IN PRODUCTION OR BUSINESS-CRITICAL SYSTEMS. NVIDIA has no obligation to make available a commercial version of any Pre-Release Licensed Software and NVIDIA has the right to abandon development of Pre-Release Licensed Software at any time without liability.

1.4 Enterprise and Contractor Usage. You may allow your Enterprise employees and Contractors to access and use the Licensed Software pursuant to the terms of the AGREEMENT solely to perform work on your behalf, provided further that with respect to Contractors: (i) you obtain a written agreement from each Contractor which contains terms and obligations with respect to access to and use of Licensed Software no less protective of NVIDIA than those set forth in the AGREEMENT, and (ii) such Contractor's access and use expressly excludes any sublicensing or distribution rights for the Licensed Software. You are responsible for the compliance with the terms and conditions of the AGREEMENT by your Enterprise and Contractors. Any act or omission that, if committed by you, would constitute a breach of the AGREEMENT shall

be deemed to constitute a breach of the AGREEMENT if committed by your Enterprise or Contractors.

1.5 Services. Except as expressly indicated in an Order, NVIDIA is under no obligation to provide support for the Licensed Software or to provide any patches, maintenance, updates or upgrades under the AGREEMENT. Unless patches, maintenance, updates or upgrades are provided with their separate governing terms and conditions, they constitute Licensed Software licensed to you under the AGREEMENT.

2. LIMITATIONS.

2.1 License Restrictions. Except as expressly authorized in the AGREEMENT, you agree that you will not (nor authorize third parties to): (i) copy and use Software that was licensed to you for use in one or more NVIDIA hardware products in other unlicensed products (provided that copies solely for backup purposes are allowed); (ii) reverse engineer, decompile, disassemble (except to the extent applicable laws specifically require that such activities be permitted) or attempt to derive the source code, underlying ideas, algorithm or structure of Software provided to you in object code form; (iii) sell, transfer, assign, distribute, rent, loan, lease, sublicense or otherwise make available the Licensed Software or its functionality to third parties (a) as an application services provider or service bureau, (b) by operating hosted/virtual system environments, (c) by hosting, time sharing or providing any other type of services, or (d) otherwise by means of the internet; (iv) modify, translate or otherwise create any derivative works of any Licensed Software; (v) remove, alter, cover or obscure any proprietary notice that appears on or with the Licensed Software or any copies thereof; (vi) use the Licensed Software, or allow its use, transfer, transmission or export in violation of any applicable export control laws, rules or regulations; (vii) distribute, permit access to, or sublicense the Licensed Software as a stand-alone product; (viii) bypass, disable, circumvent or remove any form of copy protection, encryption, security or digital rights management or authentication mechanism used by NVIDIA in connection with the Licensed Software, or use the Licensed Software together with any authorization code, serial number, or other copy protection device not supplied by NVIDIA directly or through an authorized reseller; (ix) use the Licensed Software for the purpose of developing competing products or technologies or assisting a third party in such activities; (x) use the Licensed Software with any system or application where the use or failure of such system or application can reasonably be expected to threaten or result in personal injury, death, or catastrophic loss including, without limitation, use in connection with any nuclear, avionics, navigation, military, medical, life support or other life critical application ("Critical Applications"), unless the parties have entered into a Critical Applications agreement; (xi) distribute any modification or derivative work you make to the Licensed Software under or by reference to the same name as used by NVIDIA; or (xii) use the Licensed Software in any manner that would cause the Licensed Software to become subject to an Excluded License. Nothing in the AGREEMENT shall be construed to give you a right to use, or otherwise obtain access to, any source code from which the Software or any portion thereof is compiled or interpreted. You acknowledge that NVIDIA does not design, test, manufacture or certify the Licensed Software for use in the context of a Critical Application and NVIDIA shall not be liable to you or any third party, in whole or in part, for any claims or damages arising from such use. You agree to defend, indemnify and hold harmless NVIDIA and its Affiliates, and their respective employees, contractors, agents, officers and directors, from and

against any and all claims, damages, obligations, losses, liabilities, costs or debt, fines, restitutions and expenses (including but not limited to attorney's fees and costs incident to establishing the right of indemnification) arising out of or related to you and your Enterprise, and their respective employees, contractors, agents, distributors, resellers, end users, officers and directors use of Licensed Software outside of the scope of the AGREEMENT or any other breach of the terms of the AGREEMENT.

2.2 Third Party License Obligations. The Licensed Software may come bundled with, or otherwise include or be distributed with, third party software licensed by an NVIDIA supplier and/or open source software provided under an open source license (collectively, "Third Party Software"). Notwithstanding anything to the contrary herein, Third Party Software is licensed to you subject to the terms and conditions of the software license agreement accompanying such Third Party Software whether in the form of a discrete agreement, click-through license, or electronic license terms accepted at the time of installation and any additional terms or agreements provided by the third party licensor ("Third Party License Terms"). Use of the Third Party Software by you shall be governed by such Third Party License Terms, or if no Third Party License Terms apply, then the Third Party Software is provided to you as-is, without support or warranty or indemnity obligations, for use in or with the Licensed Software and not otherwise used separately. Copyright to Third Party Software is held by the copyright holders indicated in the Third Party License Terms.

Audio/Video Encoders and Decoders. You acknowledge and agree that it is your sole responsibility to obtain any additional third party licenses required to make, have made, use, have used, sell, import, and offer for sale your products or services that include or incorporate any Third Party Software and content relating to audio and/or video encoders and decoders from, including but not limited to, Microsoft, Thomson, Fraunhofer IIS, Sisvel S.p.A., MPEG-LA, and Coding Technologies as NVIDIA does not grant to you under the AGREEMENT any necessary patent or other rights with respect to audio and/or video encoders and decoders.

2.3 Limited Rights. Your rights in the Licensed Software are limited to those expressly granted under the AGREEMENT and no other licenses are granted whether by implication, estoppel or otherwise. NVIDIA reserves all rights, title and interest in and to the Licensed Software not expressly granted under the AGREEMENT.

3. CONFIDENTIALITY. Neither party will use the other party's Confidential Information, except as necessary for the performance of the AGREEMENT, nor will either party disclose such Confidential Information to any third party, except to personnel of NVIDIA and its Affiliates, you, your Enterprise, your Enterprise Contractors, and each party's legal and financial advisors that have a need to know such Confidential Information for the performance of the AGREEMENT, provided that each such personnel, employee and Contractors are subject to a written agreement that includes confidentiality obligations consistent with those set forth herein. Each party will use all reasonable efforts to maintain the confidentiality of all of the other party's Confidential Information in its possession or control, but in no event less than the efforts that it ordinarily uses with respect to its own Confidential Information of similar nature and importance. The foregoing obligations will not restrict either party from disclosing the other party's Confidential Information or the terms and conditions of the AGREEMENT as required under applicable securities regulations or pursuant to the order or requirement of a court, administrative agency, or other governmental body,

provided that the party required to make such disclosure (i) gives reasonable notice to the other party to enable it to contest such order or requirement prior to its disclosure (whether through protective orders or otherwise), (ii) uses reasonable effort to obtain confidential treatment or similar protection to the fullest extent possible to avoid such public disclosure, and (iii) discloses only the minimum amount of information necessary to comply with such requirements.

NVIDIA Confidential Information under the AGREEMENT includes output from Licensed Software developer tools identified as "Pro" versions, where the output reveals functionality or performance data pertinent to NVIDIA hardware or software products.

4. OWNERSHIP. You are not obligated to disclose to NVIDIA any modifications that you, your Enterprise or your Contractors make to the Licensed Software as permitted under the AGREEMENT. As between the parties, all modifications are owned by NVIDIA and licensed to you under the AGREEMENT unless otherwise expressly provided in a Supplement. The Licensed Software and all modifications owned by NVIDIA, and the respective Intellectual Property Rights therein, are and will remain the sole and exclusive property of NVIDIA or its licensors. You shall not engage in any act or omission that would impair NVIDIA's and/or its licensors' Intellectual Property Rights in the Licensed Software or any other materials, information, processes or subject matter proprietary to NVIDIA. NVIDIA's licensors are intended third party beneficiaries with the right to enforce provisions of the AGREEMENT with respect to their Confidential Information and/or Intellectual Property Rights.

5. FEEDBACK. You may, but you are not obligated, to provide Feedback to NVIDIA. You hereby grant NVIDIA and its Affiliates a perpetual, non-exclusive, worldwide, irrevocable license to use, reproduce, modify, license, sublicense (through multiple tiers of sublicensees), distribute (through multiple tiers of distributors) and otherwise commercialize any Feedback that you voluntarily provide without the payment of any royalties or fees to you. NVIDIA has no obligation to respond to Feedback or to incorporate Feedback into the Licensed Software.

6. NO WARRANTIES. THE LICENSED SOFTWARE AND ANY CONFIDENTIAL INFORMATION AND/OR SERVICES ARE PROVIDED BY NVIDIA "AS IS" AND "WITH ALL FAULTS," AND NVIDIA AND ITS AFFILIATES EXPRESSLY DISCLAIM ALL WARRANTIES OF ANY KIND OR NATURE, WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF OPERABILITY, CONDITION, VALUE, ACCURACY OF DATA, OR QUALITY, AS WELL AS ANY WARRANTIES OF MERCHANTABILITY, SYSTEM INTEGRATION, WORKMANSHIP, SUITABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, NON-INFRINGEMENT, OR THE ABSENCE OF ANY DEFECTS THEREIN, WHETHER LATENT OR PATENT. NO WARRANTY IS MADE ON THE BASIS OF TRADE USAGE, COURSE OF DEALING OR COURSE OF TRADE. WITHOUT LIMITING THE FOREGOING, NVIDIA AND ITS AFFILIATES DO NOT WARRANT THAT THE LICENSED SOFTWARE OR ANY CONFIDENTIAL INFORMATION AND/OR SERVICES PROVIDED UNDER THE AGREEMENT WILL MEET YOUR REQUIREMENTS OR THAT THE OPERATION THEREOF WILL BE UNINTERRUPTED OR ERROR-FREE, OR THAT ALL ERRORS WILL BE CORRECTED.

7. LIMITATION OF LIABILITY. TO THE MAXIMUM EXTENT PERMITTED BY LAW, NVIDIA AND ITS AFFILIATES SHALL NOT BE LIABLE FOR ANY SPECIAL,

INCIDENTAL, PUNITIVE OR CONSEQUENTIAL DAMAGES, OR ANY LOST PROFITS, LOSS OF USE, LOSS OF DATA OR LOSS OF GOODWILL, OR THE COSTS OF PROCURING SUBSTITUTE PRODUCTS, ARISING OUT OF OR IN CONNECTION WITH THE AGREEMENT OR THE USE OR PERFORMANCE OF THE LICENSED SOFTWARE AND ANY CONFIDENTIAL INFORMATION AND/OR SERVICES PROVIDED UNDER THE AGREEMENT, WHETHER SUCH LIABILITY ARISES FROM ANY CLAIM BASED UPON BREACH OF CONTRACT, BREACH OF WARRANTY, TORT (INCLUDING NEGLIGENCE), PRODUCT LIABILITY OR ANY OTHER CAUSE OF ACTION OR THEORY OF LIABILITY. IN NO EVENT WILL NVIDIA'S AND ITS AFFILIATES TOTAL CUMULATIVE LIABILITY UNDER OR ARISING OUT OF THE AGREEMENT EXCEED THE NET AMOUNTS RECEIVED BY NVIDIA OR ITS AFFILIATES FOR YOUR USE OF THE PARTICULAR LICENSED SOFTWARE DURING THE TWELVE (12) MONTHS BEFORE THE LIABILITY AROSE or up to US \$10.00 if you acquired the Licensed Software for no charge). THE NATURE OF THE LIABILITY, THE NUMBER OF CLAIMS OR SUITS OR THE NUMBER OF PARTIES WITHIN YOUR ENTERPRISE THAT ACCEPTED THE TERMS OF THE AGREEMENT SHALL NOT ENLARGE OR EXTEND THIS LIMIT. THE FOREGOING LIMITATIONS SHALL APPLY REGARDLESS OF WHETHER NVIDIA, ITS AFFILIATES OR ITS LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES AND REGARDLESS OF WHETHER ANY REMEDY FAILS ITS ESSENTIAL PURPOSE. YOU ACKNOWLEDGE THAT NVIDIA'S OBLIGATIONS UNDER THE AGREEMENT ARE FOR THE BENEFIT OF YOU ONLY. The disclaimers, exclusions and limitations of liability set forth in the AGREEMENT form an essential basis of the bargain between the parties, and, absent any such disclaimers, exclusions or limitations of liability, the provisions of the AGREEMENT, including, without limitation, the economic terms, would be substantially different.

8. TERM AND TERMINATION.

8.1 AGREEMENT, Licenses and Services. This SLA shall become effective upon the Effective Date, each Supplement upon their acceptance, and both this SLA and Supplements shall continue in effect until your last access or use of the Licensed Software and/or services hereunder, unless earlier terminated as provided in this "Term and Termination" section. Each Licensed Software license ends at the earlier of (a) the expiration of the applicable license term, or (b) termination of such license or the AGREEMENT. Each service ends at the earlier of (x) the expiration of the applicable service term, (y) termination of such service or the AGREEMENT, or (z) expiration or termination of the associated license and no credit or refund will be provided upon the expiration or termination of the associated license for any service fees paid.

8.2 Termination and Effect of Expiration or Termination. NVIDIA may terminate the AGREEMENT in whole or in part: (i) if you breach any term of the AGREEMENT and fail to cure such breach within thirty (30) days following notice thereof from NVIDIA (or immediately if you violate NVIDIA's Intellectual Property Rights); (ii) if you become the subject of a voluntary or involuntary petition in bankruptcy or any proceeding relating to insolvency, receivership, liquidation or composition for the benefit of creditors, if that petition or proceeding is not dismissed with prejudice within sixty (60) days after filing, or if you cease to do business; or (iii) if you commence or participate in any legal proceeding against NVIDIA, with respect to the Licensed Software that is the subject of the proceeding during the pendency of such legal proceeding. If you or your authorized

NVIDIA reseller fail to pay license fees or service fees when due then NVIDIA may, in its sole discretion, suspend or terminate your license grants, services and any other rights provided under the AGREEMENT for the affected Licensed Software, in addition to any other remedies NVIDIA may have at law or equity. Upon any expiration or termination of the AGREEMENT, a license or a service provided hereunder, (a) any amounts owed to NVIDIA become immediately due and payable, (b) you must promptly discontinue use of the affected Licensed Software and/or service, and (c) you must promptly destroy or return to NVIDIA all copies of the affected Licensed Software and all portions thereof in your possession or control, and each party will promptly destroy or return to the other all of the other party's Confidential Information within its possession or control. Upon written request, you will certify in writing that you have complied with your obligations under this section. Upon expiration or termination of the AGREEMENT all provisions survive except for the license grant provisions.

9. CONSENT TO COLLECTION AND USE OF INFORMATION.

You hereby agree and acknowledge that the Software may access and collect data about your Enterprise computer systems as well as configures the systems in order to (a) properly optimize such systems for use with the Software, (b) deliver content through the Software, (c) improve NVIDIA products and services, and (d) deliver marketing communications. Data collected by the Software includes, but is not limited to, system (i) hardware configuration and ID, (ii) operating system and driver configuration, (iii) installed applications, (iv) applications settings, performance, and usage data, and (iv) usage metrics of the Software. To the extent that you use the Software, you hereby consent to all of the foregoing, and represent and warrant that you have the right to grant such consent. In addition, you agree that you are solely responsible for maintaining appropriate data backups and system restore points for your Enterprise systems, and that NVIDIA will have no responsibility for any damage or loss to such systems (including loss of data or access) arising from or relating to (a) any changes to the configuration, application settings, environment variables, registry, drivers, BIOS, or other attributes of the systems (or any part of such systems) initiated through the Software; or b) installation of any Software or third party software patches initiated through the Software. In certain systems you may change your system update preferences by unchecking "Automatically check for updates" in the "Preferences" tab of the control panel for the Software.

In connection with the receipt of the Licensed Software or services you may receive access to links to third party websites and services and the availability of those links does not imply any endorsement by NVIDIA. NVIDIA encourages you to review the privacy statements on those sites and services that you choose to visit so that you can understand how they may collect, use and share personal information of individuals. NVIDIA is not responsible or liable for: (i) the availability or accuracy of such links; or (ii) the products, services or information available on or through such links; or (iii) the privacy statements or practices of sites and services controlled by other companies or organizations.

To the extent that you or members of your Enterprise provide to NVIDIA during registration or otherwise personal data, you acknowledge that such information will be collected, used and disclosed by NVIDIA in accordance with NVIDIA's privacy policy, available at URL http://www.nvidia.com/object/privacy_policy.html.

10. GENERAL.

This SLA, any Supplements incorporated hereto, and Orders constitute the entire agreement of the parties with respect to the subject matter hereto and supersede all prior negotiations, conversations, or discussions between the parties relating to the subject matter hereto, oral or written, and all past dealings or industry custom. Any additional and/or conflicting terms and conditions on purchase order(s) or any other documents issued by you are null, void, and invalid. Any amendment or waiver under the AGREEMENT must be in writing and signed by representatives of both parties.

The AGREEMENT and the rights and obligations thereunder may not be assigned by you, in whole or in part, including by merger, consolidation, dissolution, operation of law, or any other manner, without written consent of NVIDIA, and any purported assignment in violation of this provision shall be void and of no effect. NVIDIA may assign, delegate or transfer the AGREEMENT and its rights and obligations hereunder, and if to a non-Affiliate you will be notified.

Each party acknowledges and agrees that the other is an independent contractor in the performance of the AGREEMENT, and each party is solely responsible for all of its employees, agents, contractors, and labor costs and expenses arising in connection therewith. The parties are not partners, joint ventures or otherwise affiliated, and neither has any authority to make any statements, representations or commitments of any kind to bind the other party without prior written consent.

Neither party will be responsible for any failure or delay in its performance under the AGREEMENT (except for any payment obligations) to the extent due to causes beyond its reasonable control for so long as such force majeure event continues in effect.

The AGREEMENT will be governed by and construed under the laws of the State of Delaware and the United States without regard to the conflicts of law provisions thereof and without regard to the United Nations Convention on Contracts for the International Sale of Goods. The parties consent to the personal jurisdiction of the federal and state courts located in Santa Clara County, California. You acknowledge and agree that a breach of any of your promises or agreements contained in the AGREEMENT may result in irreparable and continuing injury to NVIDIA for which monetary damages may not be an adequate remedy and therefore NVIDIA is entitled to seek injunctive relief as well as such other and further relief as may be appropriate. If any court of competent jurisdiction determines that any provision of the AGREEMENT is illegal, invalid or unenforceable, the remaining provisions will remain in full force and effect. Unless otherwise specified, remedies are cumulative.

The Licensed Software has been developed entirely at private expense and is "commercial items" consisting of "commercial computer software" and "commercial computer software documentation" provided with RESTRICTED RIGHTS. Use, duplication or disclosure by the U.S. Government or a U.S. Government subcontractor is subject to the restrictions set forth in the AGREEMENT pursuant to DFARS 227.7202-3(a) or as set forth in subparagraphs (c)(1) and (2) of the Commercial Computer Software - Restricted Rights clause at FAR 52.227-19, as applicable. Contractor/manufacturer is NVIDIA, 2788 San Tomas Expressway, Santa Clara, CA 95051.

You acknowledge that the Licensed Software described under the AGREEMENT is subject to export control under the U.S. Export Administration Regulations (EAR) and

economic sanctions regulations administered by the U.S. Department of Treasury's Office of Foreign Assets Control (OFAC). Therefore, you may not export, reexport or transfer in-country the Licensed Software without first obtaining any license or other approval that may be required by BIS and/or OFAC. You are responsible for any violation of the U.S. or other applicable export control or economic sanctions laws, regulations and requirements related to the Licensed Software. By accepting this SLA, you confirm that you are not a resident or citizen of any country currently embargoed by the U.S. and that you are not otherwise prohibited from receiving the Licensed Software.

Any notice delivered by NVIDIA to you under the AGREEMENT will be delivered via mail, email or fax. Please direct your legal notices or other correspondence to NVIDIA Corporation, 2788 San Tomas Expressway, Santa Clara, California 95051, United States of America, Attention: Legal Department.

GLOSSARY OF TERMS

Certain capitalized terms, if not otherwise defined elsewhere in this SLA, shall have the meanings set forth below:

- a. "Affiliate" means any legal entity that Owns, is Owned by, or is commonly Owned with a party. "Own" means having more than 50% ownership or the right to direct the management of the entity.
- b. "AGREEMENT" means this SLA and all associated Supplements entered by the parties referencing this SLA.
- c. "Authorized Users" means your Enterprise individual employees and any of your Enterprise's Contractors, subject to the terms of the "Enterprise and Contractors Usage" section.
- d. "Confidential Information" means the Licensed Software (unless made publicly available by NVIDIA without confidentiality obligations), and any NVIDIA business, marketing, pricing, research and development, know-how, technical, scientific, financial status, proposed new products or other information disclosed by NVIDIA to you which, at the time of disclosure, is designated in writing as confidential or proprietary (or like written designation), or orally identified as confidential or proprietary or is otherwise reasonably identifiable by parties exercising reasonable business judgment, as confidential. Confidential Information does not and will not include information that: (i) is or becomes generally known to the public through no fault of or breach of the AGREEMENT by the receiving party; (ii) is rightfully known by the receiving party at the time of disclosure without an obligation of confidentiality; (iii) is independently developed by the receiving party without use of the disclosing party's Confidential Information; or (iv) is rightfully obtained by the receiving party from a third party without restriction on use or disclosure.
- e. "Contractor" means an individual who works primarily for your Enterprise on a contractor basis from your secure network.
- f. "Documentation" means the NVIDIA documentation made available for use with the Software, including (without limitation) user manuals, datasheets, operations instructions, installation guides, release notes and other materials provided to you under the AGREEMENT.

g. "Enterprise" means you or any company or legal entity for which you accepted the terms of this SLA, and their subsidiaries of which your company or legal entity owns more than fifty percent (50%) of the issued and outstanding equity.

h. "Excluded License" includes, without limitation, a software license that requires as a condition of use, modification, and/or distribution that software be (i) disclosed or distributed in source code form; (ii) licensed for the purpose of making derivative works; or (iii) redistributable at no charge.

i. "Feedback" means any and all suggestions, feature requests, comments or other feedback regarding the Licensed Software, including possible enhancements or modifications thereto.

j. "Intellectual Property Rights" means all patent, copyright, trademark, trade secret, trade dress, trade names, utility models, mask work, moral rights, rights of attribution or integrity service marks, master recording and music publishing rights, performance rights, author's rights, database rights, registered design rights and any applications for the protection or registration of these rights, or other intellectual or industrial property rights or proprietary rights, howsoever arising and in whatever media, whether now known or hereafter devised, whether or not registered, (including all claims and causes of action for infringement, misappropriation or violation and all rights in any registrations and renewals), worldwide and whether existing now or in the future.

k. "Licensed Software" means Software, Documentation and all modifications owned by NVIDIA.

l. "Order" means a purchase order issued by you, a signed purchase agreement with you, or other ordering document issued by you to NVIDIA or a NVIDIA authorized reseller (including any on-line acceptance process) that references and incorporates the AGREEMENT and is accepted by NVIDIA.

m. "Software" means the NVIDIA software programs licensed to you under the AGREEMENT including, without limitation, libraries, sample code, utility programs and programming code.

n. "Supplement" means the additional terms and conditions beyond those stated in this SLA that apply to certain Licensed Software licensed hereunder.

MICROSOFT DETOURS

Microsoft Detours is used under the Professional license (<http://research.microsoft.com/en-us/projects/detours/>).

NVIDIA agrees to include in all copies of the NVIDIA Applications a proprietary rights notice that includes a reference to Microsoft software being included in such applications. NVIDIA shall not remove or obscure, but shall retain in the Software, any copyright, trademark, or patent notices that appear in the Software.

BREAKPAD

Copyright 2006, Google Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- ▶ Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- ▶ Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- ▶ Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

PROTOBUF

Copyright 2014, Google Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- ▶ Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- ▶ Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- ▶ Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF

THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Code generated by the Protocol Buffer compiler is owned by the owner of the input file used when generating it. This code is not standalone and requires a support library to be linked with it. This support library is itself covered by the above license.

Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2019-2020 NVIDIA Corporation. All rights reserved.

This product includes software developed by the Syncro Soft SRL (<http://www.sync.ro/>).