



# cuDLA API

## API Reference Manual

# Table of Contents

<b>Chapter 1. Modules</b> .....	<b>1</b>
1.1. Data types used by cuDLA driver.....	1
cudlaDevAttribute.....	2
cudlaExternalMemoryHandleDesc_t.....	2
cudlaExternalSemaphoreHandleDesc_t.....	2
CudlaFence.....	2
cudlaModuleAttribute.....	2
cudlaModuleTensorDescriptor.....	2
cudlaSignalEvents.....	2
cudlaTask.....	2
cudlaWaitEvents.....	2
cudlaDevAttributeType.....	2
cudlaFenceType.....	2
cudlaMode.....	3
cudlaModuleAttributeType.....	3
cudlaNvSciSyncAttributes.....	3
cudlaStatus.....	3
cudlaSubmissionFlags.....	5
1.2. cuDLA API.....	5
cudlaCreateDevice.....	5
cudlaDestroyDevice.....	6
cudlaDeviceGetAttribute.....	6
cudlaDeviceGetCount.....	7
cudlaGetLastError.....	7
cudlaGetNvSciSyncAttributes.....	8
cudlaGetVersion.....	9
cudlaImportExternalMemory.....	9
cudlaImportExternalSemaphore.....	10
cudlaMemRegister.....	11
cudlaMemUnregister.....	12
cudlaModuleGetAttributes.....	12
cudlaModuleLoadFromMemory.....	13
cudlaModuleUnload.....	14
cudlaSubmitTask.....	14
<b>Chapter 2. Data Structures</b> .....	<b>17</b>

cudaDevAttribute.....	17
deviceVersion.....	17
unifiedAddressingSupported.....	17
cudaExternalMemoryHandleDesc_t.....	17
extBufObject.....	18
size.....	18
cudaExternalSemaphoreHandleDesc_t.....	18
extSyncObject.....	18
CudlaFence.....	18
fence.....	18
type.....	18
cudaModuleAttribute.....	18
inputTensorDesc.....	19
numInputTensors.....	19
numOutputTensors.....	19
outputTensorDesc.....	19
cudaModuleTensorDescriptor.....	19
cudaSignalEvents.....	19
devPtrs.....	19
eofFences.....	19
numEvents.....	19
cudaTask.....	20
inputTensor.....	20
moduleHandle.....	20
numInputTensors.....	20
numOutputTensors.....	20
outputTensor.....	20
signalEvents.....	20
waitEvents.....	20
cudaWaitEvents.....	20
numEvents.....	20
preFences.....	21
<b>Chapter 3. Data Fields.....</b>	<b>22</b>



---

# Chapter 1. Modules

Here is a list of all modules:

- ▶ [Data types used by cuDLA driver](#)
- ▶ [cuDLA API](#)

## 1.1. Data types used by cuDLA driver

union cudlaDevAttribute

struct cudlaExternalMemoryHandleDesc\_t

struct cudlaExternalSemaphoreHandleDesc\_t

struct CudlaFence

union cudlaModuleAttribute

struct cudlaModuleTensorDescriptor

struct cudlaSignalEvents

struct cudlaTask

struct cudlaWaitEvents

enum cudlaDevAttributeType

Device attribute type.

#### Values

**CUDLA\_UNIFIED\_ADDRESSING = 0**

Flag to check for support for UVA.

**CUDLA\_DEVICE\_VERSION = 1**

Flag to check for DLA HW version.

enum cudlaFenceType

Supported fence types.

#### Values

**CUDLA\_NVSCISYNC\_FENCE = 1**

NvSciSync fence type.

## enum cudlaMode

Device creation modes.

### Values

**CUDLA\_CUDA\_DLA = 0**

Hybrid mode.

**CUDLA\_STANDALONE = 1**

Standalone mode.

## enum cudlaModuleAttributeType

Module attribute types.

### Values

**CUDLA\_NUM\_INPUT\_TENSORS = 0**

Flag to retrieve number of input tensors.

**CUDLA\_NUM\_OUTPUT\_TENSORS = 1**

Flag to retrieve number of output tensors.

**CUDLA\_INPUT\_TENSOR\_DESCRIPTOR = 2**

Flag to retrieve all the input tensor descriptors.

**CUDLA\_OUTPUT\_TENSOR\_DESCRIPTOR = 3**

Flag to retrieve all the output tensor descriptors.

## enum cudlaNvSciSyncAttributes

cuDLA NvSciSync attributes.

### Values

**CUDLA\_NVSCISYNC\_ATTR\_WAIT = 1**

Wait attribute.

**CUDLA\_NVSCISYNC\_ATTR\_SIGNAL = 2**

Signal attribute.

## enum cudlaStatus

Error codes.

### Values

**cudlaSuccess = 0**

The API call returned with no errors.

**cudlaErrorInvalidParam = 1**

This indicates that one or more parameters passed to the API is/are incorrect.

**cudaErrorOutOfResources = 2**

This indicates that the API call failed due to lack of underlying resources.

**cudaErrorCreationFailed = 3**

This indicates that an internal error occurred during creation of device handle.

**cudaErrorInvalidAddress = 4**

This indicates that the memory object being passed in the API call has not been registered before.

**cudaErrorOs = 5**

This indicates that an OS error occurred.

**cudaErrorCuda = 6**

This indicates that there was an error in a CUDA operation as part of the API call.

**cudaErrorUmd = 7**

This indicates that there was an error in the DLA runtime for the API call.

**cudaErrorInvalidDevice = 8**

This indicates that the device handle passed to the API call is invalid.

**cudaErrorInvalidAttribute = 9**

This indicates that an invalid attribute is being requested.

**cudaErrorIncompatibleDlaSWVersion = 10**

This indicates that the underlying DLA runtime is incompatible with the current cuDLA version.

**cudaErrorMemoryRegistered = 11**

This indicates that the memory object is already registered.

**cudaErrorInvalidModule = 12**

This indicates that the module being passed is invalid.

**cudaErrorUnsupportedOperation = 13**

This indicates that the operation being requested by the API call is unsupported.

**cudaErrorNvSci = 14**

This indicates that the NvSci operation requested by the API call failed.

**cudaErrorDlaErrInvalidInput = 0x40000001**

DLA HW Error.

**cudaErrorDlaErrInvalidPreAction = 0x40000002**

DLA HW Error.

**cudaErrorDlaErrNoMem = 0x40000003**

DLA HW Error.

**cudaErrorDlaErrProcessorBusy = 0x40000004**

DLA HW Error.

**cudaErrorDlaErrTaskStatusMismatch = 0x40000005**

DLA HW Error.

**cudaErrorDlaErrEngineTimeout = 0x40000006**

DLA HW Error.

**cudaErrorDlaErrDataMismatch = 0x40000007**

DLA HW Error.



**cudaErrorUnknown = 0x7fffffff**

This indicates that an unknown error has occurred.

## enum cudaSubmissionFlags

Task submission flags for [cudaSubmitTask](#).

### Values

**CUDLA\_SUBMIT\_NOOP = 1**

Flag to specify that the submitted task must be bypassed for execution.

## 1.2. cuDLA API

This section describes the application programming interface of the cuDLA driver.

### cudaStatus cudaCreateDevice (const uint64\_t device, const cudaDevHandle \*devHandle, const uint32\_t flags)

Create a device handle.

#### Parameters

**device**

- Device number (can be 0 or 1).

**devHandle**

- Pointer to hold the created cuDLA device handle.

**flags**

- Flags controlling device creation.

#### Returns

[cudaSuccess](#), [cudaErrorOutOfResources](#), [cudaErrorInvalidParam](#), [cudaErrorIncompatibleDlaSWVersion](#), [cudaErrorCreationFailed](#), [cudaErrorCuda](#), [cudaErrorUmd](#), [cudaErrorUnsupportedOperation](#)

#### Description

Creates an instance of a cuDLA device which can be used to submit DLA operations. The application can create the handle in hybrid or standalone mode. In hybrid mode, the current set GPU device is used by this API to decide the association of the created DLA device handle. This function returns [cudaErrorUnsupportedOperation](#) if the current set GPU device is a dGPU as cuDLA is not supported on dGPU presently.

Valid values for `flags` are:

- ▶ [CUDLA\\_CUDA\\_DLA](#) - In this mode, cuDLA serves as a programming model extension of CUDA wherein DLA work can be submitted using CUDA constructs.
- ▶ [CUDLA\\_STANDALONE](#) - In this mode, cuDLA works standalone without any interaction with CUDA.

## cudaStatus cudaDestroyDevice (const cudaDevHandle devHandle)

Destroy device handle.

### Parameters

#### **devHandle**

- A valid device handle.

### Returns

[cudaSuccess](#), [cudaErrorInvalidDevice](#), [cudaErrorCuda](#), [cudaErrorUmd](#)

### Description

Destroys the instance of the cuDLA device which was created with `cudaCreateDevice`. Before destroying the handle, it is important to ensure that all the tasks submitted previously to the device are completed. Failure to do so can lead to application crashes.

In hybrid mode, cuDLA internally performs memory allocations with CUDA using the primary context. As a result, before destroying or resetting a CUDA primary context, it is mandatory that all cuDLA device initializations are destroyed.

## cudaStatus cudaDeviceGetAttribute (const cudaDevHandle devHandle, const cudaDevAttributeType attrib, const cudaDevAttribute \*pAttribute)

Get cuDLA device attributes.

### Parameters

#### **devHandle**

- The input cuDLA device handle.

#### **attrib**

- The attribute that is being requested.

#### **pAttribute**

- The output pointer where the attribute will be available.

## Returns

[cudaSuccess](#), [cudaErrorInvalidParam](#), [cudaErrorInvalidDevice](#), [cudaErrorUmd](#), [cudaErrorInvalidAttribute](#)

## Description

UVA addressing between CUDA and DLA requires special support in the underlying kernel mode drivers. Applications are expected to query the cuDLA runtime to check if the current version of cuDLA supports UVA addressing.

## cudaStatus cudaDeviceGetCount (const uint64\_t \*pNumDevices)

Get device count.

## Parameters

### **pNumDevices**

- The number of DLA devices will be available in this variable upon successful completion.

## Returns

[cudaSuccess](#), [cudaErrorInvalidParam](#), [cudaErrorUmd](#), [cudaErrorIncompatibleDlaSWVersion](#)

## Description

Get number of DLA devices available to use.

## cudaStatus cudaGetLastError (const cudaDevHandle devHandle)

Gets the last asynchronous error in task execution.

## Parameters

### **devHandle**

- A valid device handle.

## Returns

[cudaSuccess](#), [cudaErrorInvalidDevice](#), [cudaErrorDlaErrInvalidInput](#), [cudaErrorDlaErrInvalidPreAction](#), [cudaErrorDlaErrNoMem](#), [cudaErrorDlaErrProcessorBusy](#), [cudaErrorDlaErrTaskStatusMismatch](#), [cudaErrorDlaErrEngineTimeout](#), [cudaErrorDlaErrDataMismatch](#), [cudaErrorUnknown](#)

## Description

The DLA tasks execute asynchronously on the DLA HW. As a result, the status of the task execution is not known at the time of task submission. The status of the task executed by the DLA HW most recently for the particular device handle can be queried using this interface.

Note that a return code of [`cudaSuccess`](#) from this function does not necessarily imply that most recent task executed successfully. Since this function returns immediately, it can only report the status of the tasks at the snapshot of time when it is called. To be guaranteed of task completion, applications must synchronize on the submitted tasks in hybrid or standalone modes and then call this API to check for errors.

## `cudaStatus cudaGetNvSciSyncAttributes (uint64_t *attrList, const uint32_t flags)`

Get cuDLA's NvSciSync attributes.

### Parameters

#### **attrList**

- Attribute list created by the application.

#### **flags**

- Applications can use this flag to specify how they intend to use the NvSciSync object created from the `attrList`.

### Returns

[`cudaSuccess`](#), [`cudaErrorInvalidParam`](#), [`cudaErrorUnsupportedOperation`](#), [`cudaErrorInvalidAttribute`](#), [`cudaErrorNvSci`](#)

## Description

Gets the NvSciSync's attributes in the attribute list created by the application.

In the event of failed NvSci initialization this function would return [`cudaErrorUnsupportedOperation`](#). This function can return [`cudaErrorNvSci`](#) or [`cudaErrorInvalidAttribute`](#) in certain cases when the underlying NvSci operation fails.

The valid values of `flags` can be one of the following (or an OR of these values):

- ▶ [`CUDA\_NVSCISYNC\_ATTR\_WAIT`](#), specifies that the application intends to use the NvSciSync object created using this attribute list as a waiter in cuDLA and therefore needs cuDLA to fill waiter specific NvSciSyncAttr.
- ▶ [`CUDA\_NVSCISYNC\_ATTR\_SIGNAL`](#), specifies that the application intends to use the NvSciSync object created using this attribute list as a signaler in cuDLA and therefore needs cuDLA to fill signaler specific NvSciSyncAttr.

## cudaStatus cudaGetVersion (const uint64\_t \*version)

Returns the version number of the library.

### Parameters

#### version

- cuDLA library version will be available in this variable upon successful execution.

### Returns

[cudaSuccess](#), [cudaErrorInvalidParam](#)

### Description

cuDLA is semantically versioned. This function will return the version as 1000000\*major + 1000\*minor + patch.

## cudaStatus cudaImportExternalMemory (const cudaDevHandle devHandle, const cudaExternalMemoryHandleDesc \*desc, const uint64\_t \*\*devPtr, const uint32\_t flags)

Imports external memory into cuDLA.

### Parameters

#### devHandle

- A valid device handle.

#### desc

- Contains description about allocated external memory.

#### devPtr

- The output pointer where the mapping will be available.

#### flags

- Reserved for future. Must be set to 0.

### Returns

[cudaSuccess](#), [cudaErrorInvalidParam](#), [cudaErrorInvalidDevice](#),  
[cudaErrorUnsupportedOperation](#), [cudaErrorNvSci](#), [cudaErrorInvalidAttribute](#),  
[cudaErrorMemoryRegistered](#), [cudaErrorUmd](#)

### Description

Imports the allocated external memory by registering it with DLA. After successful registration, the returned pointer can be used in a task submit.

On Tegra, cuDLA supports importing NvSciBuf objects in standalone mode only. In the event of failed NvSci initialization (either due to usage of this API in hybrid mode or an issue in the NvSci library initialization), this function would return [cudaErrorUnsupportedOperation](#). This function can return [cudaErrorNvSci](#) or [cudaErrorInvalidAttribute](#) in certain cases when the underlying NvSci operation fails.

**Note:**

This API can return task execution errors from previous DLA task submissions.

```
cudaStatus cudlaImportExternalSemaphore
(const cudaDevHandle devHandle, const
cudaExternalSemaphoreHandleDesc *desc, const
uint64_t **devPtr, const uint32_t flags)
```

Imports external semaphore into cuDLA.

### Parameters

**devHandle**

- A valid device handle.

**desc**

- Contains semaphore object.

**devPtr**

- The output pointer where the mapping will be available.

**flags**

- Reserved for future. Must be set to 0.

### Returns

[cudaSuccess](#), [cudaErrorInvalidParam](#), [cudaErrorInvalidDevice](#),  
[cudaErrorUnsupportedOperation](#), [cudaErrorNvSci](#), [cudaErrorInvalidAttribute](#),  
[cudaErrorMemoryRegistered](#)

### Description

Imports the allocated external semaphore by registering it with DLA. After successful registration, the returned pointer can be used in a task submit to signal synchronization objects.

On Tegra, cuDLA supports importing NvSciSync objects in standalone mode only. In the event of failed NvSci initialization (either due to usage of this API in hybrid mode or an issue in the NvSci library initialization), this function would return [cudaErrorUnsupportedOperation](#). This

function can return [cudaErrorNvSci](#) or [cudaErrorInvalidAttribute](#) in certain cases when the underlying NvSci operation fails.



**Note:**

This API can return task execution errors from previous DLA task submissions.

```
cudaStatus cudaMemRegister (const
  cudaDevHandle devHandle, const uint64_t *ptr, const
  size_t size, const uint64_t **devPtr, const uint32_t
  flags)
```

Registers the CUDA memory to DLA engine.

### Parameters

**devHandle**

- A valid cuDLA device handle create by a previous call to [cudaCreateDevice](#).

**ptr**

- The CUDA pointer to be registered.

**size**

- The size of the mapping i.e the number of bytes from ptr that must be mapped.

**devPtr**

- The output pointer where the mapping will be available.

**flags**

- Reserved for future. Must be set to 0.

### Returns

[cudaSuccess](#), [cudaErrorInvalidDevice](#), [cudaErrorInvalidParam](#), [cudaErrorInvalidAddress](#), [cudaErrorCuda](#), [cudaErrorUmd](#), [cudaErrorOutOfResources](#), [cudaErrorMemoryRegistered](#)

### Description

As part of registration, a system mapping is created whereby the DLA HW can access the underlying CUDA memory. The resultant mapping is available in devPtr and applications must use this mapping while referring this memory in submit operations.

This function will return [cudaErrorInvalidAddress](#) if the pointer or size to be registered is invalid. In addition, if the input pointer was already registered, then this function will return [cudaErrorMemoryRegistered](#). Attempting to re-register memory does not cause any irrecoverable error in cuDLA and applications can continue to use cuDLA APIs even after this error has occurred.

**Note:**

This API can return task execution errors from previous DLA task submissions.

## `cudaStatus cudaMemUnregister (const cudaDevHandle devHandle, const uint64_t *devPtr)`

Unregisters the input memory from DLA engine.

### Parameters

#### **devHandle**

- A valid cuDLA device handle create by a previous call to [cudaCreateDevice](#).

#### **devPtr**

- The pointer to be unregistered.

### Returns

[cudaSuccess](#), [cudaErrorInvalidDevice](#), [cudaErrorInvalidAddress](#), [cudaErrorUmd](#)

### Description

The system mapping that enables the DLA HW to access the memory is removed. This mapping could have been created by a previous call to [cudaMemRegister](#) , [cudaImportExternalMemory](#) or [cudaImportExternalSemaphore](#).

**Note:**

This API can return task execution errors from previous DLA task submissions.

## `cudaStatus cudaModuleGetAttributes (const cudaModule hModule, const cudaModuleAttributeType attrType, const cudaModuleAttribute *attribute)`

Get DLA module attributes.

### Parameters

#### **hModule**

- The input DLA module.

#### **attrType**

- The attribute type that is being requested.



**attribute**

- The output pointer where the attribute will be available.

**Returns**

[cudaSuccess](#), [cudaErrorInvalidParam](#), [cudaErrorInvalidModule](#), [cudaErrorInvalidDevice](#), [cudaErrorUmd](#), [cudaErrorInvalidAttribute](#)

**Description**

Get module attributes from the loaded module. This API returns [cudaErrorInvalidDevice](#) if the module is not loaded in any device.

```
cudaStatus cudaModuleLoadFromMemory (const
cudaDevHandle devHandle, const uint8_t *pModule,
const size_t moduleSize, const cudaModule
*hModule, const uint32_t flags)
```

Load a DLA module.

**Parameters****devHandle**

- The input cuDLA device handle. The module will be loaded in the context of this handle.

**pModule**

- A pointer to an in-memory module.

**moduleSize**

- The size of the module.

**hModule**

- The address in which the loaded module handle will be available upon successful execution.

**flags**

- Reserved for future. Must be set to 0.

**Returns**

[cudaSuccess](#), [cudaErrorInvalidDevice](#), [cudaErrorInvalidParam](#), [cudaErrorOutOfResources](#), [cudaErrorUnsupportedOperation](#), [cudaErrorUmd](#)

**Description**

Loads the module into the current device handle. Currently, DLA supports only 1 loadable per device handle. So, attempting to load another loadable in the same device handle would return with an error code of [cudaErrorUnsupportedOperation](#).

## cudaStatus cudaModuleUnload (const cudaModule hModule, const uint32\_t flags)

Unload a DLA module.

### Parameters

#### hModule

- Handle to the loaded module.

#### flags

- Reserved for future. Must be set to 0.

### Returns

[cudaSuccess](#), [cudaErrorInvalidParam](#), [cudaErrorInvalidDevice](#), [cudaErrorInvalidModule](#), [cudaErrorUmd](#)

### Description

Unload the module from the device handle that it was loaded into. This API returns [cudaErrorInvalidDevice](#) if the module is not loaded into a valid device.



#### Note:

This API can return task execution errors from previous DLA task submissions.

## cudaStatus cudaSubmitTask (const cudaDevHandle devHandle, const cudaTask \*ptrToTasks, const uint32\_t numTasks, const void \*stream, const uint32\_t flags)

Submits the inference operation on DLA.

### Parameters

#### devHandle

- A valid cuDLA device handle.

#### ptrToTasks

- A list of inferencing tasks.

#### numTasks

- The number of tasks.

#### stream

- The stream on which the DLA task has to be submitted.

**flags**

- Reserved for future. Must be set to 0.

**Returns**

[cudaSuccess](#), [cudaErrorInvalidParam](#), [cudaErrorInvalidDevice](#), [cudaErrorInvalidModule](#), [cudaErrorCuda](#), [cudaErrorUmd](#), [cudaErrorOutOfResources](#), [cudaErrorInvalidAddress](#), [cudaErrorUnsupportedOperation](#), [cudaErrorInvalidAttribute](#), [cudaErrorNvSci](#)

**Description**

This operation takes in a sequence of tasks and submits them to the DLA HW for execution in the same sequence as they appear in the input task array. The input and output tensors are assumed to be pre-registered using [cudaMemRegister](#) (in hybrid mode) or [cudaImportExternalMemory](#) (in standalone mode). Failure to do so can result in this function returning [cudaErrorInvalidAddress](#).

The `stream` parameter must be specified as the CUDA stream on which the DLA task is submitted for execution in hybrid mode. In standalone mode, this parameter must be passed as NULL and failure to do so will result in this function returning [cudaErrorInvalidParam](#).

The [cudaTask](#) structure has a provision to specify wait and signal events that cuDLA must wait on and signal respectively as part of [cudaSubmitTask\(\)](#). Each submitted task will wait for all its wait events to be signaled before beginning execution and will provide a signal event (if one is requested for during [cudaSubmitTask](#)) that the application (or any other entity) can wait on to ensure that the submitted task has completed execution. In cuDLA 1.0, only NvSciSync fences are supported as part of wait events. Furthermore, only NvSciSync objects (registered as part of [cudaImportExternalSemaphore](#)) can be signaled as part of signal events and the fence corresponding to the signaled event is returned as part of [cudaSubmitTask](#).

For wait events, applications are expected to

- ▶ register their synchronization objects using [cudaImportExternalSemaphore](#).
- ▶ create the required number of fence placeholders using [CudlaFence](#).
- ▶ fill in the placeholders with the relevant fences from the application.
- ▶ list out all the fences in [cudaWaitEvents](#).

For signal events, applications are expected to

- ▶ register their synchronization objects using [cudaImportExternalSemaphore](#).
- ▶ create the required number of placeholder fences using [CudlaFence](#).
- ▶ place the registered objects and the corresponding fences in [cudaSignalEvents](#).

When [cudaSubmitTask](#) returns successfully, the fences present in [cudaSignalEvents](#) can be used to wait for the particular task to be completed. As cuDLA only supports syncpoints

at the present moment, the number of events specified in [`cudaTask::signalEvents`](#) cannot be more than 1. If more than 1 is specified, [`cudaErrorInvalidParam`](#) is returned.

This function can return [`cudaErrorUnsupportedOperation`](#) if

- ▶ stream being used in hybrid mode is in capturing state.
- ▶ application attempts to use NvSci functionalities in hybrid mode.
- ▶ loading of NvSci libraries failed for a particular platform.
- ▶ fence type other than [`CUDA\_NVSCISYNC\_FENCE`](#) is specified.

This function can return [`cudaErrorNvSci`](#) or [`cudaErrorInvalidAttribute`](#) in certain cases when the underlying NvSci operation fails.

**Note:**

This API can return task execution errors from previous DLA task submissions.

---

# Chapter 2. Data Structures

Here are the data structures with brief descriptions:

[cudaDevAttribute](#)

[cudaExternalMemoryHandleDesc](#)

[cudaExternalSemaphoreHandleDesc](#)

[CudaFence](#)

[cudaModuleAttribute](#)

[cudaModuleTensorDescriptor](#)

[cudaSignalEvents](#)

[cudaTask](#)

[cudaWaitEvents](#)

## 2.1. cudaDevAttribute Union Reference

Device attribute.

`uint32_t cudaDevAttribute::deviceVersion`

DLA device version. Xavier has 1.0 and Orin has 2.0.

`uint8_t`

`cudaDevAttribute::unifiedAddressingSupported`

Returns 0 if unified addressing is not supported.

## 2.2. cudaExternalMemoryHandleDesc\_t Struct Reference

External memory handle descriptor.

```
const void
*cudlaExternalMemoryHandleDesc_t::extBufObject
```

A handle representing an external memory object.

```
unsigned long long
cudlaExternalMemoryHandleDesc_t::size
```

Size of the memory allocation

## 2.3. cudlaExternalSemaphoreHandleDesc\_t Struct Reference

External semaphore handle descriptor.

```
const void
*cudlaExternalSemaphoreHandleDesc_t::extSyncObject
```

A handle representing an external synchronization object.

## 2.4. CudlaFence Struct Reference

Fence description.

```
void *CudlaFence::fence
```

Fence.

```
cudlaFenceType CudlaFence::type
```

Fence type.

## 2.5. cudlaModuleAttribute Union Reference

Module attribute.

`cudaModuleTensorDescriptor`  
`*cudaModuleAttribute::inputTensorDesc`

Returns an array of input tensor descriptors.

`uint32_t cudaModuleAttribute::numInputTensors`

Returns the number of input tensors.

`uint32_t cudaModuleAttribute::numOutputTensors`

Returns the number of output tensors.

`cudaModuleTensorDescriptor`  
`*cudaModuleAttribute::outputTensorDesc`

Returns an array of output tensor descriptors.

## 2.6. `cudaModuleTensorDescriptor` Struct Reference

Tensor descriptor.

## 2.7. `cudaSignalEvents` Struct Reference

Signal events for [cudaSubmitTask](#)

`const **cudaSignalEvents::devPtrs`

Array of registered synchronization objects (via [cudaImportExternalSemaphore](#)).

`CudaFence *cudaSignalEvents::eofFences`

Array of fences pointers for all the signal events corresponding to the synchronization objects.

`uint32_t cudaSignalEvents::numEvents`

Total number of signal events.

## 2.8. cudlaTask Struct Reference

Structure of Task.

`const **cudlaTask::inputTensor`

Array of input tensors.

`cudlaModule cudlaTask::moduleHandle`

cuDLA module handle.

`uint32_t cudlaTask::numInputTensors`

Number of input tensors.

`uint32_t cudlaTask::numOutputTensors`

Number of output tensors.

`const **cudlaTask::outputTensor`

Array of output tensors.

`cudlaSignalEvents *cudlaTask::signalEvents`

Signal events.

`const cudlaWaitEvents *cudlaTask::waitEvents`

Wait events.

## 2.9. cudlaWaitEvents Struct Reference

Wait events for [cudlaSubmitTask](#).

`uint32_t cudlaWaitEvents::numEvents`

Total number of wait events.



```
const CudlaFence *cudlaWaitEvents::preFences
```

Array of fence pointers for all the wait events.

---

# Chapter 3. Data Fields

Here is a list of all documented struct and union fields with links to the struct/union documentation for each field:

**deviceVersion**

[cudlaDevAttribute](#)

**devPtrs**

[cudlaSignalEvents](#)

**eofFences**

[cudlaSignalEvents](#)

**extBufObject**

[cudlaExternalMemoryHandleDesc](#)

**extSyncObject**

[cudlaExternalSemaphoreHandleDesc](#)

**fence**

[CudlaFence](#)

**inputTensor**

[cudlaTask](#)

**inputTensorDesc**

[cudlaModuleAttribute](#)

**moduleHandle**

[cudlaTask](#)

**numEvents**

[cudlaWaitEvents](#)

[cudlaSignalEvents](#)

**numInputTensors**

[cudlaTask](#)

[cudlaModuleAttribute](#)

**numOutputTensors**

[cudlaTask](#)

[cudlaModuleAttribute](#)

**outputTensor**

[cudlaTask](#)

**outputTensorDesc**

[cudlaModuleAttribute](#)

**preFences**[cudaWaitEvents](#)**signalEvents**[cudaTask](#)**size**[cudaExternalMemoryHandleDesc](#)**type**[CudlaFence](#)**unifiedAddressingSupported**[cudaDevAttribute](#)**waitEvents**[cudaTask](#)

## Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

## OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

## 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

© 2021-2022 NVIDIA Corporation & affiliates. All rights reserved.