



NVIDIA TensorRT

API Reference | NVIDIA Docs

Table of Contents

Chapter 1. Added, Deprecated, And Removed APIs.....	1
1.1. API Changes For TensorRT 8.6.0 EA.....	1
1.2. API Changes For TensorRT 8.5.1.....	6
1.3. API Changes For TensorRT 8.4.1.....	11
1.4. API Changes For TensorRT 8.4.0 EA.....	11
1.5. API Changes For TensorRT 8.2.0 EA.....	12
1.6. API Changes For TensorRT 8.0.1.....	13
Chapter 2. C++ API.....	21
Chapter 3. Python API.....	22

List of Tables

Table 1.	New C++ Classes.....	1
Table 2.	New C++ Macros.....	1
Table 3.	New C++ Functions.....	2
Table 4.	New C++ Enums.....	3
Table 5.	New C++ Enum Values.....	3
Table 6.	Deprecated C++ Enums.....	3
Table 7.	Deprecated C++ Macros.....	3
Table 8.	Deprecated C++ Functions.....	4
Table 9.	New Python Classes.....	4
Table 10.	New Python Functions.....	4
Table 11.	New Python Properties.....	4
Table 12.	New Python Enums.....	6
Table 13.	New Python Enum Values.....	6
Table 14.	Deprecated Python Enums.....	6
Table 15.	Deprecated Python Properties.....	6
Table 16.	New C++ API Functions.....	6
Table 17.	New C++ API Classes.....	8
Table 18.	New C++ API Enum Values.....	8
Table 19.	New C++ API Enum Classes.....	8
Table 20.	New Python API Functions and Properties.....	9
Table 21.	New Python API Classes.....	10
Table 22.	New Python API Enum Values.....	10
Table 23.	New Python API Enum Classes.....	11
Table 24.	New Python APIs.....	11

Table 25.	New C++ APIs.....	11
Table 26.	Deprecated C++ APIs.....	12
Table 27.	New Python APIs.....	12
Table 28.	Deprecated Python APIs.....	12
Table 29.	New C++ APIs.....	12
Table 30.	New Python APIs.....	13
Table 31.	New C++ APIs.....	13
Table 32.	Removed C++ APIs.....	14
Table 33.	Removed Plugins.....	16
Table 34.	Unsupported plugin methods removed in TensorRT 8.0.....	17
Table 35.	Updated versions for supported plugin methods.....	17
Table 36.	New Python APIs.....	17
Table 37.	Removed Python APIs.....	18
Table 38.	Deprecated APIs.....	20

Chapter 1. Added, Deprecated, And Removed APIs

1.1. API Changes For TensorRT 8.6.0 EA

The following tables show which APIs were added, deprecated, and removed for the NVIDIA® TensorRT™ 8.6.0 EA release.

C++ Changes

Table 1. New C++ Classes

New C++ Classes
<u>ICastLayer</u>
<u>IReverseSequenceLayer</u>
<u>INormalizationLayer</u>
<u>ILoggerFinder</u>

Table 2. New C++ Macros

New C++ Macros
<u>NV_TENSORRT_RELEASE_TYPE</u>
<u>NV_TENSORRT_RELEASE_TYPE_EARLY_ACCESS</u>
<u>NV_TENSORRT_RELEASE_TYPE_RELEASE_CANDIDATE</u>
<u>NV_TENSORRT_RELEASE_TYPE_GENERAL_AVAILABILITY</u>

Table 3. New C++ Functions

New C++ Functions
<code>getBuilderSafePluginRegistry()</code>
<code>IAlgorithmIOInfo::getVectorizedDim()</code>
<code>IAlgorithmIOInfo::getComponentsPerElement()</code>
<code>IBuilderConfig::setBuilderOptimizationLevel()</code>
<code>IBuilderConfig::getBuilderOptimizationLevel()</code>
<code>IBuilderConfig::setHardwareCompatibilityLevel()</code>
<code>IBuilderConfig::getHardwareCompatibilityLevel()</code>
<code>IBuilderConfig::setPluginsToSerialize()</code>
<code>IBuilderConfig::getPluginToSerialize()</code>
<code>IBuilderConfig::getNbPluginsToSerialize()</code>
<code>IBuilderConfig::getMaxAuxStreams()</code>
<code>IBuilderConfig::setMaxAuxStreams()</code>
<code>IBuilder::getPluginRegistry()</code>
<code>ICudaEngine::getHardwareCompatibilityLevel()</code>
<code>ICudaEngine::getNbAuxStreams()</code>
<code>IExecutionContext::setAuxStreams()</code>
<code>ILayer::setMetadata()</code>
<code>ILayer::getMetadata()</code>
<code>INetworkDefinition::addCast()</code>
<code>INetworkDefinition::addNormalization()</code>
<code>INetworkDefinition::addReverseSequence()</code>
<code>INetworkDefinition::getBuilder()</code>
<code>IPluginRegistry::isParentSearchEnabled()</code>
<code>IPluginRegistry::setParentSearchEnabled()</code>
<code>IPluginRegistry::loadLibrary()</code>
<code>IPluginRegistry::deregisterLibrary()</code>
<code>IRuntime::setTemporaryDirectory()</code>
<code>IRuntime::getTemporaryDirectory()</code>
<code>IRuntime::setTempfileControlFlags()</code>
<code>IRuntime::getTempfileControlFlags()</code>
<code>IRuntime::getPluginRegistry()</code>

New C++ Functions[IRuntime::loadRuntime\(\)](#)[IRuntime::setEngineHostCodeAllowed\(\)](#)[IRuntime::getEngineHostCodeAllowed\(\)](#)[ITopKLayer::setInput\(\)](#)

Table 4. New C++ Enums

New C++ Enums[HardwareCompatibilityLevel](#)[TempfileControlFlag](#)

Table 5. New C++ Enum Values

New C++ Enum Values[BuilderFlag::kVERSION_COMPATIBLE](#)[BuilderFlag::kEXCLUDE_LEAN_RUNTIME](#)[DataType::kFP8](#)[LayerType::kREVERSE_SEQUENCE](#)[LayerType::kNORMALIZATION](#)[LayerType::kCAST](#)[MemoryPoolType::kTACTIC_DRAM](#)[PreviewFeature::kPROFILE_SHARING_0806](#)[TensorFormat::kDHW](#)[UnaryOperation::kISINF](#)

Table 6. Deprecated C++ Enums

Deprecated C++ Enums[PreviewFeature::kFASTER_DYNAMIC_SHAPES](#)

Table 7. Deprecated C++ Macros

Deprecated C++ Macros[NV_TENSORRT_SONAME_MAJOR](#)

Deprecated C++ Macros

NV_TENSORRT_SONAME_MINOR

NV_TENSORRT_SONAME_PATCH

Table 8. Deprecated C++ Functions

Deprecated C++ Functions

FieldMap::FieldMap()

IAlgorithmIOInfo::getTensorFormat()

Python Changes

Table 9. New Python Classes

New Python Classes[ICastLayer](#)[IReverseSequenceLayer](#)[INormalizationLayer](#)

Table 10. New Python Functions

New Python Functions[IExecutionContext.set_aux_streams](#)[INetworkDefinition.add_cast](#)[INetworkDefinition.add_normalization](#)[INetworkDefinition.add_reverse_sequence](#)[IPluginRegistry.load_library](#)[IPluginRegistry.deregister_library](#)[IRuntime.load_runtime](#)[ITopKLayer.set_input\(\)](#)

Table 11. New Python Properties

New Python Properties[BuilderFlag.EXCLUDE_LEAN_RUNTIME](#)[BuilderFlag.FP8](#)

New Python Properties

<code>BuilderFlag.VERSION_COMPATIBLE</code>
<code>DataType.FP8</code>
<code>HardwareCompatibilityLevel.AMPERE_PLUS</code>
<code>HardwareCompatibilityLevel.NONE</code>
<code>IAlgorithmIOInfo.components_per_element</code>
<code>IAlgorithmIOInfo.vectorized_dim</code>
<code>IBuilder.get_plugin_registry</code>
<code>IBuilderConfig.builder_optimization_level</code>
<code>IBuilderConfig.hardware_compatibility_level</code>
<code>IBuilderConfig.max_aux_streams</code>
<code>IBuilderConfig.plugins_to_serialize</code>
<code>ICudaEngine.hardware_compatibility_level</code>
<code>ICudaEngine.num_aux_streams</code>
<code>ILayer.metadata</code>
<code>INetworkDefinition.builder</code>
<code>INormalizationLayer.axes</code>
<code>INormalizationLayer.compute_precision</code>
<code>INormalizationLayer.epsilon</code>
<code>INormalizationLayer.num_groups</code>
<code>IPluginRegistry.parent_search_enabled</code>
<code>IReverseSequenceLayer.batch_axis</code>
<code>IReverseSequenceLayer.sequence_axis</code>
<code>IRuntime.engine_host_code_allowed</code>
<code>IRuntime.load_runtime</code>
<code>IRuntime.tempfile_control_flags</code>
<code>IRuntime.temporary_directory</code>
<code>LayerType.CAST</code>
<code>LayerType.NORMALIZATION</code>
<code>LayerType.REVERSE_SEQUENCE</code>
<code>MemoryPoolType.TACTIC_DRAM</code>
<code>PreviewFeature.PROFILE_SHARING_0806</code>
<code>TempfileControlFlag.ALLOW_IN_MEMORY_FILES</code>

New Python Properties

TempfileControlFlag.ALLOW_TEMPORARY_FILES

[TensorFormat.DHWC](#)

Table 12. New Python Enums

New Python Enums

HardwareCompatibilityLevel

TempfileControlFlag

Table 13. New Python Enum Values

New Python Enum Values[UnaryOperation.ISINF](#)

Table 14. Deprecated Python Enums

Deprecated Python Enums

PreviewFeature.FASTER_DYNAMIC_SHAPES

Table 15. Deprecated Python Properties

Deprecated Python Properties

IAlgorithmIOInfo.tensor_format

1.2. API Changes For TensorRT 8.5.1

The following tables show which APIs were added, deprecated, and removed for the NVIDIA® TensorRT™ 8.5.1 release.

C++ changes

Table 16. New C++ API Functions

New C++ API Functions[ITensor::setDimensionName\(\)](#)[ITensor::getDimensionName\(\)](#)

New C++ API Functions

<u>IResizeLayer::setCubicCoeff()</u>
<u>IResizeLayer::getCubicCoeff()</u>
<u>IResizeLayer::setExcludeOutside()</u>
<u>IResizeLayer::getExcludeOutside()</u>
<u>IBuilderConfig::setPreviewFeature()</u>
<u>IBuilderConfig::getPreviewFeature()</u>
<u>ICudaEngine::getTensorShape()</u>
<u>ICudaEngine::getTensorDataType()</u>
<u>ICudaEngine::getTensorLocation()</u>
<u>ICudaEngine::isShapeInferenceIO()</u>
<u>ICudaEngine::getTensorIOMode()</u>
<u>ICudaEngine::getTensorBytesPerComponent()</u>
<u>ICudaEngine::getTensorComponentsPerElement()</u>
<u>ICudaEngine::getTensorFormat()</u>
<u>ICudaEngine::getTensorFormatDesc()</u>
<u>ICudaEngine::getProfileShape()</u>
<u>ICudaEngine::getNbIOTensors()</u>
<u>ICudaEngine::getIOTensorName()</u>
<u>IExecutionContext::getTensorStrides()</u>
<u>IExecutionContext::setInputShape()</u>
<u>IExecutionContext::getTensorShape()</u>
<u>IExecutionContext::setTensorAddress()</u>
<u>IExecutionContext::getTensorAddress()</u>
<u>IExecutionContext::setInputTensorAddress()</u>
<u>IExecutionContext::getOutputTensorAddress()</u>
<u>IExecutionContext::inferShapes()</u>
<u>IExecutionContext::setInputConsumedEvent()</u>
<u>IExecutionContext::getInputConsumedEvent()</u>
<u>IExecutionContext::setOutputAllocator()</u>
<u>IExecutionContext::getOutputAllocator()</u>
<u>IExecutionContext::getMaxOutputSize()</u>
<u>IExecutionContext::setTemporaryStorageAllocator()</u>

New C++ API Functions

[ExecutionContext::getTemporaryStorageAllocator\(\)](#)

[ExecutionContext::enqueueV3\(\)](#)

[ExecutionContext::setPersistentCacheLimit\(\)](#)

[ExecutionContext::getPersistentCacheLimit\(\)](#)

[ExecutionContext::setNvtxVerbosity\(\)](#)

[ExecutionContext::getNvtxVerbosity\(\)](#)

[INetworkDefinition::addOneHot\(\)](#)

[INetworkDefinition::addNonZero\(\)](#)

[INetworkDefinition::addGridSample\(\)](#)

[INetworkDefinition::addNMS\(\)](#)

Table 17. New C++ API Classes

New C++ API Classes

[IOneHotLayer](#)

[IGridSampleLayer](#)

[INonZeroLayer](#)

[INMSLayer](#)

[IOutputAllocator](#)

Table 18. New C++ API Enum Values

New C++ API Enum Values

[InterpolationMode::kCUBIC](#)

[FillOperation::kRANDOM_NORMAL](#)

[BuilderFlag::kREJECT_EMPTY_ALGORITHMS](#)

[BuilderFlag::kENABLE_TACTIC_HEURISTIC](#)

[TacticSource::kJIT_CONVOLUTIONS](#)

[DataType::kUINT8](#)

Table 19. New C++ API Enum Classes

New C++ API Enum Classes

[TensorIOMode](#)

New C++ API Enum Classes

[PreviewFeature](#)

Python changes

Table 20. New Python API Functions and Properties

New Python API Functions and Properties
<u>ITensor.set_dimension_name()</u>
<u>ITensor.get_dimension_name()</u>
<u>IResizeLayer.cubic_coeff</u>
<u>IResizeLayer.exclude_outside</u>
<u>IBuilderConfig.set_preview_feature()</u>
<u>IBuilderConfig.get_preview_feature()</u>
<u>ICudaEngine.get_tensor_shape()</u>
<u>ICudaEngine.get_tensor_dtype()</u>
<u>ICudaEngine.get_tensor_location()</u>
<u>ICudaEngine.is_shape_inference_io()</u>
<u>ICudaEngine.get_tensor_mode()</u>
<u>ICudaEngine.get_tensor_bytes_per_component()</u>
<u>ICudaEngine.get_tensor_components_per_element()</u>
<u>ICudaEngine.get_tensor_format()</u>
<u>ICudaEngine.get_tensor_format_desc()</u>
<u>ICudaEngine.get_tensor_profile_shape()</u>
<u>ICudaEngine.num_io_tensors</u>
<u>ICudaEngine.get_tensor_name()</u>
<u>IExecutionContext.get_tensor_strides()</u>
<u>IExecutionContext.set_input_shape()</u>
<u>IExecutionContext.get_tensor_shape()</u>
<u>IExecutionContext.set_tensor_address()</u>
<u>IExecutionContext.get_tensor_address()</u>
<u>IExecutionContext.infer_shapes()</u>
<u>IExecutionContext.set_input_consumed_event()</u>
<u>IExecutionContext.get_input_consumed_event()</u>

New Python API Functions and Properties
<code><u>ExecutionContext.set_output_allocator()</u></code>
<code><u>ExecutionContext.get_output_allocator()</u></code>
<code><u>ExecutionContext.get_max_output_size()</u></code>
<code><u>ExecutionContext.temporary_allocator</u></code>
<code><u>ExecutionContext.execute_async_v3()</u></code>
<code><u>ExecutionContext.persistent_cache_limit</u></code>
<code><u>ExecutionContext.nvtx_verbosity</u></code>
<code><u>INetworkDefinition.add_one_hot()</u></code>
<code><u>INetworkDefinition.add_non_zero()</u></code>
<code><u>INetworkDefinition.add_grid_sample()</u></code>
<code><u>INetworkDefinition.add_nms()</u></code>

Table 21. New Python API Classes

New Python API Classes
<code><u>IOneHotLayer</u></code>
<code><u>IGridSampleLayer</u></code>
<code><u>INonZeroLayer</u></code>
<code><u>INMSLayer</u></code>
<code><u>IOutputAllocator</u></code>

Table 22. New Python API Enum Values

New Python API Enum Values
<code><u>InterpolationMode.CUBIC</u></code>
<code><u>FillOperation.RANDOM_NORMAL</u></code>
<code><u>BuilderFlag.REJECT_EMPTY_ALGORITHMS</u></code>
<code><u>BuilderFlag.ENABLE_TACTIC_HEURISTIC</u></code>
<code><u>TacticSource.JIT_CONVOLUTIONS</u></code>
<code><u>DataType.UINT8</u></code>

Table 23. New Python API Enum Classes

New Python API Enum Classes
<u>TensorIOMode</u>
<u>PreviewFeature</u>

1.3. API Changes For TensorRT 8.4.1

The following tables show which APIs were added, deprecated, and removed for the NVIDIA® TensorRT™ 8.4.1 release.

Python changes

Table 24. New Python APIs

New Python APIs
<u>TacticSource::kEDGE_MASK_CONVOLUTIONS</u>

1.4. API Changes For TensorRT 8.4.0 EA

The following tables show which APIs were added, deprecated, and removed for the NVIDIA® TensorRT™ 8.4.0 EA release.

C++ changes

Table 25. New C++ APIs

New C++ APIs
<u>setMemoryPoolLimit (IBuilderConfig::setMemoryPoolLimit)</u>
<u>getMemoryPoolLimit (IBuilderConfig::getMemoryPoolLimit)</u>
<u>MemoryPoolType</u>
<u>setMaxThreads (IBuilder::setMaxThreads, IRefitter::setMaxThreads, IRuntime::setMaxThreads)</u>
<u>getMaxThreads (IBuilder::getMaxThreads, IRefitter::getMaxThreads, IRuntime::getMaxThreads)</u>
<u>getBuilderPluginRegistry</u>

Table 26. Deprecate C++ APIs

Deprecated C++ APIs
<code>IFullyConnectedLayer</code>
<code>getMaxWorkspaceSize</code>
<code>setMaxWorkspaceSize</code>

Python changes

Table 27. New Python APIs

New Python APIs
<code>set_memory_pool_limit</code>
<code>get_memory_pool_limit</code>
<code>MemoryPoolType</code>
<code>max_threads</code> property (<code>Builder.max_threads</code> , <code>Refitter.max_threads</code> , <code>Runtime.max_threads</code>)
<code>get_builder_plugin_registry</code>

Table 28. Deprecate Python APIs

Deprecated Python APIs
<code>IFullyConnectedLayer</code>
<code>get_max_workspace_size</code>
<code>set_max_workspace_size</code>

1.5. API Changes For TensorRT 8.2.0 EA

The following tables show which APIs were added, deprecated, and removed for the NVIDIA® TensorRT™ 8.2.0 EA release.

C++ changes

Table 29. New C++ APIs

New C++ APIs
<code>IAssertionLayer</code>
<code>IConditionLayer</code>

New C++ APIs[IEinsumLayer](#)[IScatterLayer](#)**Python changes**

Table 30. New Python APIs

New Python APIs[IAssertionLayer](#)[IConditionLayer](#)[IEinsumLayer](#)[IScatterLayer](#)

1.6. API Changes For TensorRT 8.0.1

The following tables show which APIs were added, deprecated, and removed for the TensorRT 8.0.1 release.

C++ changes

Table 31. New C++ APIs

New C++ APIs[class IDequantizeLayer](#)[class IQuantizeLayer](#)[class ITimingCache](#)[IBuilder::buildSerializedNetwork\(\)](#)[IBuilderConfig::getTimingCache\(\)](#)[IBuilderConfig::setTimingCache\(\)](#)[IGpuAllocator::reallocate\(\)](#)[INetworkDefinition::addDequantize\(\)](#)[INetworkDefinition::addQuantize\(\)](#)[INetworkDefinition::setWeightsName\(\)](#)[IPluginRegistry::deregisterCreator\(\)](#)[IRefitter::getMissingWeights\(\)](#)

New C++ APIs
<code>IRefitter::getAllWeights()</code>
<code>IRefitter::setNamedWeights()</code>
<code>IResizeLayer::getCoordinateTransformation()</code>
<code>IResizeLayer::getNearestRounding()</code>
<code>IResizeLayer::getSelectorForSinglePixel()</code>
<code>IResizeLayer::setCoordinateTransformation()</code>
<code>IResizeLayer::setNearestRounding()</code>
<code>IResizeLayer::setSelectorForSinglePixel()</code>
<code>IScaleLayer::setChannelAxis()</code>
<code>enum ResizeCoordinateTransformation</code>
<code>enum ResizeMode</code>
<code>BuilderFlag::kSPARSE_WEIGHTS</code>
<code>TacticSource::kCUDNN</code>
<code>TensorFormat::kDLA_HWC4</code>
<code>TensorFormat::kDLA_LINEAR</code>
<code>TensorFormat::kHWC16</code>

Table 32. Removed C++ APIs

Removed C++ APIs
Core Library
<code>DimensionType</code>
<code>Dims::Type</code>
<code>class DimsCHW</code>
<code>class DimsNCHW</code>
<code>class IOutputDimensionFormula</code>
<code>class IPlugin</code>
<code>class IPluginFactory</code>
<code>class IPluginLayer</code>
<code>class IRNNLayer</code>
<code>IBuilder::getEngineCapability()</code>
<code>IBuilder::allowGPUPFallback()</code>
<code>IBuilder::buildCudaEngine()</code>

Removed C++ APIs

<code>IBuilder::canRunOnDLA()</code>
<code>IBuilder::createNetwork()</code>
<code>IBuilder::getAverageFindIterations()</code>
<code>IBuilder::getDebugSync()</code>
<code>IBuilder::getDefaultDeviceType()</code>
<code>IBuilder::getDeviceType()</code>
<code>IBuilder::getDLACore()</code>
<code>IBuilder::getFp16Mode()</code>
<code>IBuilder::getHalf2Mode()</code>
<code>IBuilder::getInt8Mode()</code>
<code>IBuilder::getMaxWorkspaceSize()</code>
<code>IBuilder::getMinFindIterations()</code>
<code>IBuilder::getRefittable()</code>
<code>IBuilder::getStrictTypeConstraints()</code>
<code>IBuilder::isDeviceTypeSet()</code>
<code>IBuilder::reset()</code>
<code>IBuilder::resetDeviceType()</code>
<code>IBuilder::setAverageFindIterations()</code>
<code>IBuilder::setDebugSync()</code>
<code>IBuilder::setDefaultDeviceType()</code>
<code>IBuilder::setDeviceType()</code>
<code>IBuilder::setDLACore()</code>
<code>IBuilder::setEngineCapability()</code>
<code>IBuilder::setFp16Mode()</code>
<code>IBuilder::setHalf2Mode()</code>
<code>IBuilder::setInt8Calibrator()</code>
<code>IBuilder::setInt8Mode()</code>
<code>IBuilder::setMaxWorkspaceSize()</code>
<code>IBuilder::setMinFindIterations()</code>
<code>IBuilder::setRefittable()</code>
<code>IBuilder::setStrictTypeConstraints()</code>
<code>ICudaEngine::getWorkspaceSize()</code>

Removed C++ APIs
<code>IMatrixMultiplyLayer::getTranspose()</code>
<code>IMatrixMultiplyLayer::setTranspose()</code>
<code>INetworkDefinition::addMatrixMultiply()</code>
<code>INetworkDefinition::addPlugin()</code>
<code>INetworkDefinition::addPluginExt()</code>
<code>INetworkDefinition::addRNN()</code>
<code>INetworkDefinition::getConvolutionOutputDimensionsFormula()</code>
<code>INetworkDefinition::getDeconvolutionOutputDimensionsFormula()</code>
<code>INetworkDefinition::getPoolingOutputDimensionsFormula()</code>
<code>INetworkDefinition::setConvolutionOutputDimensionsFormula()</code>
<code>INetworkDefinition::setDeconvolutionOutputDimensionsFormula()</code>
<code>INetworkDefinition::setPoolingOutputDimensionsFormula()</code>
<code>ITensor::getDynamicRange()</code>
<code>TensorFormat::kNHWC8</code>
<code>TensorFormat::NCHW</code>
<code>TensorFormat::kNC2HW2</code>
Caffe Parser
<code>class IPluginFactory</code>
<code>class IPluginFactoryExt</code>
<code>setPluginFactory()</code>
<code>setPluginFactoryExt()</code>
UFF Parser
<code>class IPluginFactory</code>
<code>class IPluginFactoryExt</code>
<code>setPluginFactory()</code>
<code>setPluginFactoryExt()</code>

Table 33. Removed Plugins

Removed Plugins
<code>class INvPlugin</code>
<code>createLReLUPlugin()</code>
<code>createClipPlugin()</code>

Removed Plugins
<code>PluginType</code>
<code>struct SoftmaxTree</code>

For plugins based on `IPluginV2DynamicExt` and `IPluginV2IOExt`, certain methods with legacy function signatures (derived from `IPluginV2` and `IPluginV2Ext` base classes) which were deprecated and marked for removal in TensorRT 8.0 will no longer be available. Plugins using these interface methods must stop using them or implement the versions with updated signatures, as applicable.

Table 34. Unsupported plugin methods removed in TensorRT 8.0

Removed Plugins
<code>IPluginV2DynamicExt::canBroadcastInputAcrossBatch()</code>
<code>IPluginV2DynamicExt::isOutputBroadcastAcrossBatch()</code>
<code>IPluginV2DynamicExt::getTensorRTVersion()</code>
<code>IPluginV2IOExt::configureWithFormat()</code>
<code>IPluginV2IOExt::getTensorRTVersion()</code>

Table 35. Updated versions for supported plugin methods

Removed Plugin	Replaced with
	<code>IPluginV2DynamicExt::configurePlugin()</code>
	<code>IPluginV2DynamicExt::enqueue()</code>
	<code>IPluginV2DynamicExt::getOutputDimensions()</code>
	<code>IPluginV2DynamicExt::getWorkspaceSize()</code>
	<code>IPluginV2IOExt::configurePlugin()</code>
<code>IPluginV2DynamicExt::supportsFormat()</code>	<code>IPluginV2DynamicExt::supportsFormatCombination()</code>
<code>IPluginV2IOExt::supportsFormat()</code>	<code>IPluginV2IOExt::supportsFormatCombination()</code>

Python changes

Table 36. New Python APIs

New Python APIs
<code>class IDequantizeLayer</code>
<code>class IQuantizeLayer</code>

New Python APIs
<code>class ITimingCache</code>
<code>Builder.build_serialized_network()</code>
<code>IBuilderConfig.get_timing_cache()</code>
<code>IBuilderConfig.set_timing_cache()</code>
<code>IGpuAllocator.reallocate()</code>
<code>INetworkDefinition.add_dequantize()</code>
<code>INetworkDefinition.add_quantize()</code>
<code>INetworkDefinition.set_weights_name()</code>
<code>IPluginRegistry.deregister_creator()</code>
<code>Refitter.get_all_weights()</code>
<code>Refitter.get_missing_weights()</code>
<code>Refitter::set_named_weights()</code>
<code>IResizeLayer.coordinate_transformation</code>
<code>IResizeLayer.nearest_rounding</code>
<code>IResizeLayer.selector_for_single_pixel</code>
<code>IScaleLayer.channel_axis</code>
<code>enum ResizeCoordinateTransformationDoc</code>
<code>enum ResizeMode</code>
<code>BuilderFlag.SPARSE_WEIGHTS</code>
<code>TacticSource.CUDNN</code>
<code>TensorFormat.DLA_HWC4</code>
<code>TensorFormat.DLA_LINEAR</code>
<code>TensorFormat.HWC16</code>

Table 37. Removed Python APIs

Removed Python APIs
Core Library
<code>class DimsCHW</code>
<code>class DimsNCHW</code>
<code>class IPlugin</code>
<code>class IPluginFactory</code>
<code>class IPluginLayer</code>

Removed Python APIs
<code>class IRNNLayer</code>
<code>Builder.build_cuda_engine()</code>
<code>Builder.average_find_iterations</code>
<code>Builder.debug_sync</code>
<code>Builder.fpl6_mode</code>
<code>IBuilder.int8_mode</code>
<code>Builder.max_workspace_size</code>
<code>Builder.min_find_iterations</code>
<code>Builder.refittable</code>
<code>Builder.strict_type_constraints</code>
<code>ICudaEngine.max_workspace_size</code>
<code>IMatrixMultiplyLayer.transpose0</code>
<code>INetworkDefinition.add_matrix_multiply_deprecated()</code>
<code>INetworkDefinition.add_plugin()</code>
<code>INetworkDefinition.add_plugin_ext()</code>
<code>INetworkDefinition.add_rnn()</code>
<code>INetworkDefinition.convolution_output_dimensions_formula</code>
<code>INetworkDefinition.deconvolution_output_dimensions_formula</code>
<code>INetworkDefinition.pooling_output_dimensions_formula</code>
<code>ITensor.get_dynamic_range()</code>
<code>Dims.get_type()</code>
<code>TensorFormat.HWC8</code>
<code>TensorFormat.NCHW</code>
<code>TensorFormat.NCHW2</code>
Caffe Parser
<code>class IPluginFactory</code>
<code>class IPluginFactoryExt</code>
<code>CaffeParser.plugin_factory</code>
<code>CaffeParser.plugin_factory_ext</code>
UFF Parser
<code>class IPluginFactory</code>
<code>class IPluginFactoryExt</code>

Removed Python APIs
<code>UffParser.plugin_factory</code>
<code>UffParser.plugin_factory_ext</code>

Deprecated

For our deprecation policy, refer to the [TensorRT Deprecation Policy](#).

Table 38. Deprecated APIs

Deprecated APIs	Replaced with
<code>nvinfer1::IResizeLayer::setAlignCorners</code>	<code>nvinfer1::IResizeLayer::setCoordinateTransformation</code>
<code>nvinfer1::IResizeLayer::getAlignCorners</code>	<code>nvinfer1::IResizeLayer::setSelectorForSinglePixel</code>
	<code>nvinfer1::IResizeLayer::setNearestRounding</code>

Chapter 2. C++ API

To view this API, refer to the [TensorRT C++ API](#).

The NVIDIA TensorRT C++ API allows developers to import, calibrate, generate and deploy networks using C++. Networks can be imported directly from ONNX. They may also be created programmatically by instantiating individual layers and setting parameters and weights directly.

Within the core C++ API in `NvInfer.h`, the following APIs are included:

- ▶ [Builder API](#)
- ▶ [Execution API](#)
- ▶ [Network Definition API](#)
- ▶ [ONNX Parser API](#)
- ▶ [Plugin API](#)

For more information about the C++ API, including sample code, refer to the [NVIDIA TensorRT Developer Guide](#).

Chapter 3. Python API

To view this API, refer to the [TensorRT Python API](#).

The NVIDIA TensorRT Python API enables developers in Python based development environments and those looking to experiment with TensorRT to easily parse models (for example, from ONNX) and generate and run PLAN files.

For more information about the Python API, including sample code, refer to the [NVIDIA TensorRT Developer Guide](#).

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.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND 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. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Arm

Arm, AMBA and Arm Powered are registered trademarks of Arm Limited. Cortex, MPCore and Mali are trademarks of Arm Limited. "Arm" is used to represent Arm Holdings plc; its operating company Arm Limited; and the regional subsidiaries Arm Inc.; Arm KK; Arm Korea Limited.; Arm Taiwan Limited; Arm France SAS; Arm Consulting (Shanghai) Co. Ltd.; Arm Germany GmbH; Arm Embedded Technologies Pvt. Ltd.; Arm Norway, AS and Arm Sweden AB.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

BlackBerry/QNX

Copyright © 2020 BlackBerry Limited. All rights reserved.

Trademarks, including but not limited to BLACKBERRY, EMBLEM Design, QNX, AVIAGE, MOMENTICS, NEUTRINO and QNX CAR are the trademarks or registered trademarks of BlackBerry Limited, used under license, and the exclusive rights to such trademarks are expressly reserved.

Google

Android, Android TV, Google Play and the Google Play logo are trademarks of Google, Inc.

Trademarks

NVIDIA, the NVIDIA logo, and BlueField, CUDA, DALI, DRIVE, Hopper, JetPack, Jetson AGX Xavier, Jetson Nano, Maxwell, NGC, Nsight, Orin, Pascal, Quadro, Tegra, TensorRT, Triton, Turing and Volta are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2017-2024 NVIDIA Corporation & affiliates. All rights reserved.

