




# NVIDIA VIDEO ENCODER (NVENC) INTERFACE

NVENCODEAPI\_PG | Jan 2018

**Programming Guide**



## DOCUMENT CHANGE HISTORY

NVENCODAPI\_PG

Version	Date	Authors	Description of Change
1.0	2011/12/29	SD/CC	Initial release.
1.1	2012/05/04	SD	Update for Version 1.1
2.0	2012/10/12	SD	Update for Version 2.0
3.0	2013/07/25	AG	Update for Version 3.0
4.0	2014/07/01	SM	Update for Version 4.0
5.0	2014/11/30	MV	Update for Version 5.0
6.0	2015/10/15	VP	Update for Version 6.0
7.0	2016/06/10	SM	Update for Version 7.0
7.1	2016/11/15	SM	Update for Version 7.1
8.0	2017/02/15	SM	Update for Version 8.0
8.1	2018/1/10	SM	Update for Version 8.1

# TABLE OF CONTENTS

<b>Chapter 1. Introduction</b>	<b>1</b>
<b>Chapter 2. Basic Encoding Flow</b>	<b>2</b>
<b>Chapter 3. Setting Up Hardware for Encoding</b>	<b>3</b>
3.1 Opening an Encode Session	3
3.1.1 Initializing encode device	3
3.2 Selecting Encoder Codec GUID	4
3.3 Encoder Preset Configurations	5
3.3.1 Enumerating preset GUIDs	5
3.3.2 Selecting encoder preset configuration	5
3.4 Selecting an Encoder Profile	6
3.5 Getting Supported List of Input Formats	6
3.6 Querying encoder Capabilities	6
3.7 Initializing the Hardware Encoder Session	7
3.8 Encode Session Attributes	7
3.8.1 Configuring encode session attributes	7
3.8.2 Finalizing codec configuration for encoding	8
3.8.3 Rate control	9
3.8.4 Setting encode session attributes	10
3.9 Creating Resources Required to Hold Input/output Data	11
3.10 Retrieving Sequence Parameters	12
<b>Chapter 4. Encoding the Video Stream</b>	<b>13</b>
4.1 Preparing Input Buffers for Encoding	13
4.1.1 Input buffers allocated through NVIDIA Video Encoder Interface	13
4.1.2 Input buffers allocated externally	14
4.2 Configuring Per-Frame Encode Parameters	15
4.2.1 Forcing current frame to be encoded as intra frame	15
4.2.2 Forcing current frame to be used as a reference frame	15
4.2.3 Forcing current frame to be used as an IDR frame	15
4.2.4 Requesting generation of sequence parameters	15
4.3 Submitting Input Frame for Encoding	16
4.4 Retrieving Encoded Output	16
<b>Chapter 5. End of Encoding</b>	<b>17</b>
5.1 Notifying the End of Input Stream	17
5.2 Releasing Resources	17
5.3 Closing Encode Session	18
<b>Chapter 6. Modes of Operation</b>	<b>19</b>
6.1 Asynchronous Mode (Windows 7 and above)	19
6.2 Synchronous Mode	21
6.3 Threading Model	21
6.4 Encoder Features using CUDA	22
<b>Chapter 7. Motion Estimation Only Mode</b>	<b>23</b>

7.1	Query Motion-Estimation Only Mode Capability.....	23
7.2	Create Resources for Input/Output Data .....	24
7.3	Populate ME only mode settings .....	24
7.4	Run Motion Estimation .....	24
7.5	Enabling Motion estimation for stereo usecases .....	25
7.6	Release the Created Resources.....	26
<b>Chapter 8. Advanced Features and Settings .....</b>		<b>27</b>
8.1	Look-ahead .....	27
8.2	B-Frames As Reference .....	28
8.3	Reconfigure API.....	28
8.4	Adaptive Quantization (AQ) .....	29
8.4.1	Spatial AQ .....	29
8.4.2	Temporal AQ .....	29
8.5	High Bit Depth Encoding.....	30
8.6	Weighted Prediction .....	31
8.7	Long-Term Reference .....	31
8.8	Emphasis MAP .....	32
<b>Chapter 9. Recommended NVENC Settings.....</b>		<b>34</b>

# Chapter 1. INTRODUCTION

NVIDIA GPUs based on Kepler, Maxwell and the latest Pascal architectures contain a hardware-based H.264/HEVC video encoder (hereafter referred to as NVENC). The NVENC hardware takes YUV/RGB as input, and generates an H.264/HEVC compliant video bit stream. NVENC hardware's encoding capabilities can be accessed using the NVENCODE APIs, available in the NVIDIA Video Codec SDK.

This document provides information on how to program the NVENC using the NVENCODE APIs exposed in the SDK. The NVENCODE APIs expose encoding capabilities on Windows (Windows 7 and above) and Linux.

It is expected that the developers should understand H.264/HEVC video codecs and familiarity with Windows and/or Linux development environment.

NVENCODE API *guarantees* backward compatibility (and will make explicit reference whenever backward compatibility is broken). This means that applications compiled with older versions of released API will continue to work on future driver versions released by NVIDIA.

## Chapter 2. BASIC ENCODING FLOW

Developers can create a client application that calls NVENCODE API functions exposed by `nvEncodeAPI.dll` for Windows or `libnvidia-encode.so` for Linux. These libraries are installed as part of the NVIDIA display driver. The client application can either link to these libraries at run-time using `LoadLibrary()` on Windows or `dlopen()` on Linux.

The NVIDIA video encoder API is designed to accept raw video frames (in YUV or RGB format) and output the H.264 or HEVC bitstream. Broadly, the encoding flow consists of the following steps:

1. Initialize the encoder
2. Set up the desired encoding parameters
3. Allocate input/output buffers
4. Copy frames to input buffers and read bitstream from the output buffers. This can be done synchronously (Windows & Linux) or asynchronously (Windows 7 and above only).
5. Clean-up - release all allocated input/output buffers
6. Close the encoding session

These steps are explained in the rest of the document and demonstrated in the sample application included in the Video Codec SDK package.

# Chapter 3. SETTING UP HARDWARE FOR ENCODING

## 3.1 OPENING AN ENCODE SESSION

After loading the DLL or shared object library, the client's first interaction with the API is to call `NvEncodeAPICreateInstance`. This populates the input/output buffer passed to `NvEncodeAPICreateInstance` with pointers to functions which implement the functionality provided in the interface.

After loading the NVENC Interface, the client should first call `NvEncOpenEncodeSessionEx` to open an encoding session. This function returns an encode session handle which must be used for all subsequent calls to the API functions in the current session.

### 3.1.1 Initializing encode device

The NVIDIA Encoder supports use of the following types of devices:

#### 3.1.1.1 DirectX 9

The client should create a DirectX 9 device with behavior flags including:

```
D3DCREATE_FPU_PRESERVE  
D3DCREATE_MULTITHREADED  
D3DCREATE_HARDWARE_VERTEXPROCESSING
```

The client should pass a pointer to `IUnknown` interface of the created device (typecast to `void *`) as `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::device`, and set `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::deviceType` to `NV_ENC_DEVICE_TYPE_DIRECTX`. Use of DirectX devices is supported only on Windows 7 and later versions of the Windows OS.

### 3.1.1.2 DirectX 10

The client should pass a pointer to IUnknown interface of the created device (typecast to `void *`) as `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::device`, and set `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::deviceType` to `NV_ENC_DEVICE_TYPE_DIRECTX`. Use of DirectX devices is supported only on Windows 7 and later versions of Windows OS.

### 3.1.1.3 DirectX 11

The client should pass a pointer to IUnknown interface of the created device (typecast to `void *`) as `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::device`, and set `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::deviceType` to `NV_ENC_DEVICE_TYPE_DIRECTX`. Use of DirectX devices is supported only on Windows 7 and later versions of Windows OS.

### 3.1.1.4 CUDA

The client should create a floating CUDA context, and pass the CUDA context handle as `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::device`, and set `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::deviceType` to `NV_ENC_DEVICE_TYPE_CUDA`. Use of CUDA device for Encoding is supported on Linux and Windows 7 and later OS's.

### 3.1.1.5 OpenGL

The client should create an OpenGL context and make it current (in order to associate the context with the thread/process that is making calls to NVENC API) to the thread calling into NVENC API. `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::device` must be `NULL` and `NV_ENC_OPEN_ENCODE_SESSION_EX_PARAMS::deviceType` must be set to `NV_ENC_DEVICE_TYPE_OPENGL`. Use of the OpenGL device type for encoding is supported only on Linux.

## 3.2 SELECTING ENCODER CODEC GUID

The client should select an Encoding GUID that represents the desired codec for encoding the video sequence in the following manner:

1. The client should call `NvEncGetEncodeGUIDCount` to get the number of supported Encoder GUIDs from the NVIDIA Video Encoder Interface.
2. The client should use this count to allocate a buffer of sufficient size to hold the supported Encoder GUIDS.
3. The client should then call `NvEncGetEncodeGUIDs` to populate this list.

The client should select a GUID that matches its requirement from this list and use that as the `encodeGUID` for the remainder of the encoding session.

## 3.3 ENCODER PRESET CONFIGURATIONS

The NVIDIA Encoder Interface exposes various presets to cater to different video encoding use cases. Using these presets will automatically set all relevant encoding parameters. This is a coarse level of control exposed by the API. Specific attributes/parameters within the preset can be tuned, if required. This is explained in next two subsections.

### 3.3.1 Enumerating preset GUIDs

The client can enumerate supported Preset GUIDs for the selected `encodeGUID` as follows:

1. The client should call `NvEncGetEncodePresetCount` to get the number of supported Encoder GUIDs.
2. The client should use this count to allocate a buffer of sufficient size to hold the supported Preset GUIDs.
3. The client should then call `NvEncGetEncodePresetGUIDs` to populate this list.

### 3.3.2 Selecting encoder preset configuration

As mentioned above, the client can use the `presetGUID` for configuring the encode session directly. This will automatically set the hardware encoder with appropriate parameters for the use-case implied by the preset. If required, the client has the option to fine-tune the encoder configuration parameters in the preset, and override the preset defaults. This approach is often-times more convenient from programming point of view as the programmer only needs to change the configuration parameters which he/she is interested in, leaving everything else pre-configured as per the preset definition.

Here are the steps to fetch a preset encode configuration and optionally change select configuration parameters:

1. Enumerate the supported presets as described above, in Section 3.3.1.
2. Select the preset GUID for which the encode configuration is to be fetched.
3. The client should call `NvEncGetEncodePresetConfig` with the selected `encodeGUID` and `presetGUID` as inputs
4. The required preset encoder configuration can be retrieved through `NV_ENC_PRESET_CONFIG::presetCfg`.
5. Over-ride the default encoder parameters, if required, using the corresponding configuration APIs.

## 3.4 SELECTING AN ENCODER PROFILE

The client may specify a profile to encode for specific encoding scenario. For example, certain profiles are required for encoding video for playback on iPhone/iPod, encoding video for blue-ray disc authoring, etc.

The client should do the following to retrieve a list of supported encoder profiles:

1. The client should call `NvEncGetEncodeProfileGUIDCount` to get the number of supported Encoder GUIDs from the NVIDIA Video Encoder Interface.
2. The client should use this count to allocate a buffer of sufficient size to hold the supported Encode Profile GUIDS.
3. The client should then call `NvEncGetEncodeProfileGUIDs` to populate this list.

The client should select the profile GUID that best matches the requirement.

## 3.5 GETTING SUPPORTED LIST OF INPUT FORMATS

NVENC API accepts input frames in several different formats, such as YUV and RGB in specific formats, as enumerated in `NV_ENC_BUFFER_FORMAT`.

List of supported input formats can be retrieved as follows:

1. The client should call `NvEncGetInputFormatCount` to get the number of supported input formats.
2. The client should use this count to allocate a buffer to hold the list of supported input buffer formats (which are list elements of type `NV_ENC_BUFFER_FORMAT`).
3. Retrieve the supported input buffer formats by calling `NvEncGetInputFormats`.

The client should select a format enumerated in this list for creating input buffers.

## 3.6 QUERYING ENCODER CAPABILITIES

NVIDIA video encoder hardware has evolved over multiple generations, with many features being added in each new generation of the GPU. To facilitate application to dynamically figure out the capabilities of the underlying hardware encoder on the system, NVENC API provides a dedicated API to query these capabilities. It is a good programming practice to query for support of the desired encoder feature before making use of the feature.

Querying the encoder capabilities can be accomplished as follows:

1. Specify the capability attribute to be queried in `NV_ENC_CAPS_PARAM::capsToQuery` parameter. This should be a member of the `NV_ENC_CAPS` enum.
2. Call `NvEncGetEncodeCaps` to determine support for the required attribute.

Refer to the API reference `NV_ENC_CAPS` enum definition for interpretation of individual capability attributes.

## 3.7 INITIALIZING THE HARDWARE ENCODER SESSION

The client needs to call `NvEncInitializeEncoder` with a valid encoder configuration specified through `NV_ENC_INITIALIZE_PARAMS` and encoder handle (returned upon successful opening of encode session)

## 3.8 ENCODE SESSION ATTRIBUTES

### 3.8.1 Configuring encode session attributes

Encode session configuration is divided into three parts:

#### 3.8.1.1 Session parameters

Common parameters such as input format, output dimensions, display aspect ratio, frame rate, average bitrate, etc. are available in `NV_ENC_INITIALIZE_PARAMS` structure. The client should use an instance of this structure as input to `NvEncInitializeEncoder`.

The Client must populate the following members of the `NV_ENC_INITIALIZE_PARAMS` structure for the encode session to be successfully initialized:

- `NV_ENC_INITIALIZE_PARAMS::encodeGUID`: The client must select a suitable codec GUID as described in Section 3.2.
- `NV_ENC_INITIALIZE_PARAMS::encodeWidth`: The client must specify the desired width of the encoded video.
- `NV_ENC_INITIALIZE_PARAMS::encodeHeight`: The client must specify the desired height of the encoded video.

`NV_ENC_INITIALIZE_PARAMS::reportSliceOffsets` can be used to enable reporting of slice offsets. This feature requires `NV_ENC_INITIALIZE_PARAMS::enableEncodeAsync` to be set to 0, and does not work with MB-based and byte-based slicing on Kepler GPUs.

### 3.8.1.2 Advanced codec-level parameters

Parameters dealing with the encoded bit stream such as GOP length, encoder profile, rate control mode, etc. are exposed through the structure `NV_ENC_CONFIG`. The client can pass codec level parameters through `NV_ENC_INITIALIZE_PARAMS::encodeConfig` as explained below.

### 3.8.1.3 Advanced codec-specific parameters

Advanced H.264 and HEVC specific parameters are available in structures `NV_ENC_CONFIG_H264` and `NV_ENC_CONFIG_HEVC` respectively.

The client can pass codec-specific parameters through the structure `NV_ENC_CONFIG::encodeCodecConfig`.

## 3.8.2 Finalizing codec configuration for encoding

### 3.8.2.1 High-level control using presets

This is the simplest method of configuring the NVIDIA Video Encoder Interface, and involves minimal setup steps to be performed by the client. This is intended for use cases where the client does not need to fine-tune any codec level parameters.

In this case, the client should follow these steps:

1. The client should specify the session parameters as described in Section 3.8.1.1.
2. Optionally, the client can enumerate and select preset GUID that best suits the current use case, as described in Section 3.3.1. The client should then pass the selected preset GUID using `NV_ENC_INITIALIZE_PARAMS::presetGUID`. This helps the NVIDIA Video Encoder interface to correctly configure the encoder session based on the `encodeGUID` and `presetGUID` provided.
3. The client should set the advanced codec-level parameter pointer `NV_ENC_INITIALIZE_PARAMS::encodeConfig::encodeCodecConfig` to `NULL`.

### 3.8.2.2 Finer control by overriding preset parameters

The client can choose to edit some encoding parameters on top of the parameters set by the individual preset, as follows:

1. The client should specify the session parameters as described in Section 3.8.1.1.
2. The client should enumerate and select a preset GUID that best suites the current use case, as described in Section 3.3.1. The client should retrieve a preset encode configuration as described in Section 3.3.2.

3. The client may need to explicitly query the capability of the encoder to support certain features or certain encoding configuration parameters. For this, the client should do the following:
  4. Specify the capability desired attribute through `NV_ENC_CAPS_PARAM::capsToQuery` parameter. This should be a member of the `NV_ENC_CAPS` enum.
  5. Call `NvEncGetEncodeCaps` to determine support for the required attribute. Refer to `NV_ENC_CAPS` enum definition in the API reference for interpretation of individual capability attributes.
  6. Select a desired preset GUID and fetch the corresponding Preset Encode Configuration as described in Section 3.3.
  7. The client can then override any parameters from the preset `NV_ENC_CONFIG` according to its requirements. The client should pass the fine-tuned `NV_ENC_CONFIG` structure using `NV_ENC_INITIALIZE_PARAMS::encodeConfig::encodeCodecConfig` pointer.
  8. Additionally, the client should also pass the selected preset GUID through `NV_ENC_INITIALIZE_PARAMS::presetGUID`. This is to allow the NVIDIA Video Encoder interface to program internal parameters associated with the encoding session to ensure that the encoded output conforms to the client's request. Note that passing the preset GUID here will not override the fine-tuned parameters.

### 3.8.3 Rate control

NVENC supports several rate control modes and provides control over various parameters related to the rate control algorithm via structure `NV_ENC_INITIALIZE_PARAMS::encodeConfig::rcParams`. The rate control algorithm is implemented in NVENC firmware.

NVENC supports 1-pass (`NV_ENC_PARAMS_RC_VBR` and `NV_ENC_PARAMS_RC_CBR`), 2-pass (`NV_ENC_PARAMS_RC_CBR_LOWDELAY_HQ`, `NV_ENC_PARAMS_RC_CBR_HQ` and `NV_ENC_PARAMS_RC_VBR_HQ`), constant QP and Constant quality rate control modes. In 1-pass rate control mode, NVENC estimates the required QP for the macroblock and immediately encodes the macroblock. In 2-pass rate control modes, NVENC estimates the complexity of the frame to be encoded and determines bit distribution across the frame in the first pass. In the second pass, NVENC encodes macroblocks in the frame using the distribution determined in the first pass. As a result, quality of the encoded bitstream is generally better with 2-pass encoding, at the cost of some performance.

At a high level, NVENC supports following rate control modes:

- **Constant bitrate (CBR):** Constant bitrate is specified by setting `rateControlMode` to `NV_ENC_PARAMS_RC_CBR`, `NV_ENC_PARAMS_RC_CBR_LOWDELAY_HQ` or

`NV_ENC_PARAMS_RC_CBR_HQ`. In this mode, only `averageBitRate` is required and used as the target output bitrate by the rate control algorithm.

- **Variable bitrate (VBR):** Variable bitrate is specified by setting `rateControlMode` to `NV_ENC_PARAMS_RC_VBR` or `NV_ENC_PARAMS_RC_VBR_HQ`. The encoder tries to conform to average bitrate of `averageBitRate` over the long term while not exceeding `maxBitRate` any time during the encoding. In this mode, `averageBitRate` must be specified. If `maxBitRate` isn't specified, NVENC will set it to an internally determined default value. It is recommended that the client specify both parameters `maxBitRate` and `averageBitRate` for better control.
- **Constant QP:** This mode is specified by setting `rateControlMode` to `NV_ENC_PARAMS_RC_CONSTQP`. In this mode, the entire frame is encoded using QP specified in `NV_ENC_RC_PARAMS::constQP`.
- **Target quality:** This mode is specified by setting `rateControlMode` to one of the VBR modes and desired target quality in `targetQuality`. The range of this target quality is 0 to 51 (fractional values are also supported in Video Codec SDK 8.0 and above). In this mode, the encoder tries to maintain constant quality for each frame, by allowing the bitrate to vary subject to the bitrate parameter specified in `maxBitRate`. The resulting average bitrate can, therefore, vary significantly depending on the video content being encoded. If `maxBitRate` is not specified, the encoder will use as many bits as needed to achieve the target quality. However, if `maxBitRate` is set, it will form an upper bound on the actual bitrate. If `maxBitRate` is set too low in this mode, the bitrate may become constrained, resulting in the desired target quality possibly not being achieved.

## 3.8.4 Setting encode session attributes

Once all Encoder settings have been finalized, the client should populate a `NV_ENC_CONFIG` structure, and use it as an input to `NvEncInitializeEncoder` to freeze the Encode settings for the current encodes session. Some settings such as rate control mode, average bitrate, resolution etc. can be changed on-the-fly.

The client is required to explicitly specify the following while initializing the Encode Session:

### 3.8.4.1 Mode of operation

The client should set `NV_ENC_INITIALIZE_PARAMS::enableEncodeAsync` to 1 if it wants to operate in asynchronous mode and 0 for operating in synchronous mode.

*Asynchronous mode encoding is supported only on Windows 7 and later. Refer to Chapter 6 for more detailed explanation.*

### 3.8.4.2 Picture-type decision

If the client wants to send the input buffers in display order, it must set `enablePTD = 1`. If `enablePTD` is set to 1 the decision of determining the picture type will be taken by NVENCODE API.

If the client wants to send the input buffers in encode order, it must set `enablePTD = 0`, and must specify

- `NV_ENC_PIC_PARAMS::pictureType`
- `NV_ENC_PIC_PARAMS_H264/NV_ENC_PIC_PARAMS_HEVC::displayPOCSyntax`
- `NV_ENC_PIC_PARAMS_H264/NV_ENC_PIC_PARAMS_HEVC::refPicFlag`

## 3.9 CREATING RESOURCES REQUIRED TO HOLD INPUT/OUTPUT DATA

Once the encode session is initialized, the client should allocate buffers to hold the input/output data.

The client may choose to allocate input buffers through NVIDIA Video Encoder Interface by calling `NvEncCreateInputBuffer` API. In this case, the client is responsible for destroying the allocated input buffers before closing the encode session. It is also the client's responsibility to fill the input buffer with valid input data according to the chosen input buffer format.

The client should allocate buffers to hold the output encoded bit stream using the `NvEncCreateBitstreamBuffer` API. It is the client's responsibility to destroy these buffers before closing the encode session.

Alternatively, in scenarios where the client cannot or does not want to allocate input buffers through the NVIDIA Video Encoder Interface, it can use any externally allocated DirectX resource as an input buffer. However, the client must perform some simple processing to map these resources to resource handles that are recognized by the NVIDIA Video Encoder Interface before use. The translation procedure is explained in Section 4.1.2.

If the client has used a CUDA device to initialize the encoder session, and wishes to use input buffers NOT allocated through the NVIDIA Video Encoder Interface, the client is required to use buffers allocated using the `cuMemAlloc` family of APIs. The NVIDIA Video Encoder Interface version 8.0 only supports `CUdeviceptr` as a supported input format. Support for `CUarray` inputs is not supported.

If the client has used the OpenGL device type to initialize the encoder session and wishes to use input buffers NOT allocated through the NVIDIA Video Encoder Interface, the client is required to provide the textures allocated earlier.

The client may generate textures using `glGenTextures()`, bind it to either the `NV_ENC_INPUT_RESOURCE_OPENGL_TEX::GL_TEXTURE_RECTANGLE` or `NV_ENC_INPUT_RESOURCE_OPENGL_TEX::GL_TEXTURE_2D` target, allocate storage for it using `glTexImage2D()` and copy data to it.

Note that the OpenGL interface for NVENCODE API is only supported on Linux.

**Note:** The client should allocate at least  $(1 + N_B)$  input and output buffers, where  $N_B$  is the number of B frames between successive P frames.

## 3.10 RETRIEVING SEQUENCE PARAMETERS

After configuring the encode session, the client can retrieve the sequence parameter information (SPS) at any time by calling `NvEncGetSequenceParams`. It is the client's responsibility to allocate and eventually de-allocate a buffer of size `MAX_SEQ_HDR_LEN` to hold the sequence parameter information.

By default, SPS/PPS data will be attached to every IDR frame. However, the client can request the encoder to generate SPS/PPS data on demand as well. To accomplish this, set `NV_ENC_PIC_PARAMS::encodePicFlags = NV_ENC_PIC_FLAG_OUTPUT_SPSPPS`. The output bitstream generated for the current input will then include SPS/PPS.

The client can call `NvEncGetSequenceParams` at any time, after the encoder has been initialized (`NvEncInitializeEncoder`) and the session is active.

# Chapter 4. ENCODING THE VIDEO STREAM

Once the encode session is configured and input/output buffers are allocated, the client can start streaming the input data for encoding. The client is required to pass a handle to a valid input buffer and a valid bit stream (output) buffer to the NVIDIA Video Encoder Interface for encoding an input picture.

## 4.1 PREPARING INPUT BUFFERS FOR ENCODING

There are two methods to allocate and pass input buffers to the video encoder.

### 4.1.1 Input buffers allocated through NVIDIA Video Encoder Interface

If the client has allocated input buffers through `NvEncCreateInputBuffer`, the client needs to fill valid input data before using the buffer as input for encoding. For this, the client should call `NvEncLockInputBuffer` to get a CPU pointer to the input buffer. Once the client has filled input data, it should call `NvEncUnlockInputBuffer`. The input buffer should be passed to the encoder only after unlocking it. Any input buffers should be unlocked by calling `NvEncUnlockInputBuffer` before destroying/reallocating them.

## 4.1.2 Input buffers allocated externally

To pass externally allocated buffers to the encoder, follow these steps:

1. Populate `NV_ENC_REGISTER_RESOURCE` with attributes of the externally allocated buffer.
2. Call `NvEncRegisterResource` with the `NV_ENC_REGISTER_RESOURCE` populated in the above step.
3. `NvEncRegisterResource` returns an opaque handle in `NV_ENC_REGISTER_RESOURCE::registeredResource` which should be saved.
4. Call `NvEncMapInputResource` with the handle returned above.
5. The mapped handle will then be available in `NV_ENC_MAP_INPUT_RESOURCE::mappedResource`.
6. The client should use this mapped handle (`NV_ENC_MAP_INPUT_RESOURCE::mappedResource`) as the input buffer handle in `NV_ENC_PIC_PARAMS`.
7. After the client has finished using the resource `NvEncUnmapInputResource` must be called.
8. The client must also call `NvEncUnregisterResource` with the handle returned by `NvEncRegisterResource` before destroying the registered resource.

The mapped resource handle (`NV_ENC_MAP_INPUT_RESOURCE::mappedResource`) should not be used for any other purpose outside the NVIDIA Video Encoder Interface while it is in mapped state. Such usage is not supported and may lead to undefined behavior.

## 4.2 CONFIGURING PER-FRAME ENCODE PARAMETERS

The client should populate `NV_ENC_PIC_PARAMS` with the parameters to be applied to the current input picture. The client can do the following on a per-frame basis.

### 4.2.1 Forcing current frame to be encoded as intra frame

To force the current frame as intra (I) frame, set

```
NV_ENC_PIC_PARAMS::encodePicFlags = NV_ENC_PIC_FLAG_FORCEINTRA
```

### 4.2.2 Forcing current frame to be used as a reference frame

To force the current frame to be used as a reference frame, set

```
NV_ENC_PIC_PARAMS_H264/NV_ENC_PIC_PARAMS_HEVC::refPicFlag = 1
```

### 4.2.3 Forcing current frame to be used as an IDR frame

To force the current frame to be encoded as IDR frame, set

```
NV_ENC_PIC_PARAMS::encodePicFlags = NV_ENC_PIC_FLAG_FORCEIDR
```

### 4.2.4 Requesting generation of sequence parameters

To include SPS/PPS along with the currently encoded frame, set

```
NV_ENC_PIC_PARAMS::encodePicFlags = NV_ENC_PIC_FLAG_OUTPUT_SPSPPS
```

## 4.3 SUBMITTING INPUT FRAME FOR ENCODING

The client should call `NvEncEncodePicture` to perform encoding.

The input picture data will be taken from the specified input buffer, and the encoded bit stream will be available in the specified bit stream (output) buffer once the encoding process completes.

Codec-agnostic parameters such as timestamp, duration, input buffer pointer, etc. are passed via the structure `NV_ENC_PIC_PARAMS` while codec-specific parameters are passed via the structure `NV_ENC_PIC_PARAMS_H264/NV_ENC_PIC_PARAMS_HEVC` depending upon the codec in use.

The client should specify the codec-specific structure in `NV_ENC_PIC_PARAMS` using the `NV_ENC_PIC_PARAMS::codecPicParams` member.

## 4.4 RETRIEVING ENCODED OUTPUT

Upon completion of the encoding process for an input picture, the client is required to call `NvEncLockBitstream` to get a CPU pointer to the encoded bit stream. The client can make a local copy of the encoded data or pass the CPU pointer for further processing (e.g. to a media file writer).

The CPU pointer will remain valid until the client calls `NvEncUnlockBitstream`. The client should call `NvEncUnlockBitstream` after it completes processing the output data.

The client must ensure that all bit stream buffers are unlocked before destroying/de-allocating them (e.g. while closing an encode session) or even before reusing them as output buffers for subsequent frames.

# Chapter 5. END OF ENCODING

## 5.1 NOTIFYING THE END OF INPUT STREAM

To notify the end of input stream, the client must call `NvEncEncodePicture` with the flag `NV_ENC_PIC_PARAMS:: encodePicFlags` set to `NV_ENC_FLAGS_EOS` and all other members of `NV_ENC_PIC_PARAMS` set to 0. No input buffer is required while calling `NvEncEncodePicture` for EOS notification.

EOS notification effectively flushes the encoder. This can be called multiple times in a single encode session. This operation however must be done before closing the encode session.

## 5.2 RELEASING RESOURCES

Once encoding completes, the client should destroy all allocated resources.

The client should call `NvEncDestroyInputBuffer` if it had allocated input buffers through the NVIDIA Video Encoder Interface. The client must ensure that input buffer is first *unlocked* by calling `NvEncUnlockInputBuffer` before destroying it.

The client should call `NvEncDestroyBitstreamBuffer` to destroy each bitstream buffer it had allocated. The client must ensure that the bitstream buffer is first *unlocked* by calling `NvEncUnlockBitstream` before destroying it.

## 5.3 CLOSING ENCODE SESSION

The client should call `NvEncDestroyEncoder` to close the encoding session. The client should ensure that all resources tied to the encode session being closed have been destroyed before calling `NvEncDestroyEncoder`. These include input buffers, bit stream buffers, SPS/PPS buffer, etc.

It must also ensure that all registered events are unregistered, and all mapped input buffer handles are unmapped.

# Chapter 6. MODES OF OPERATION

The NVIDIA Video Encoder Interface supports the following two modes of operation.

## 6.1 ASYNCHRONOUS MODE (WINDOWS 7 AND ABOVE)

This mode of operation is used for asynchronous output buffer processing. For this mode, the client allocates an event object and associates the event with an allocated output buffer. This event object is passed to the NVIDIA Encoder Interface as part of the `NvEncEncodePicture` API. The client can wait on the event in a separate thread. When the event is signaled, the client calls the NVIDIA Video Encoder Interface to copy output bitstream produced by the encoder. Note that the encoder supports asynchronous mode of operation only for Windows 7 and above, with driver running in WDDM mode. In Linux and Windows with TCC mode (TCC mode is available on Tesla boards<sup>1</sup>), ONLY synchronous mode is supported (refer to Section 6.2.)

The client should set the flag `NV_ENC_INITIALIZE_PARAMS::enableEncodeAsync` to 1 to indicate that it wants to operate in asynchronous mode. After creating the event objects (one object for each output bitstream buffer allocated), the client needs to register them with the NVIDIA Video Encoder Interface using the `NvEncRegisterAsyncEvent`. The client is required to pass a bitstream buffer handle and the corresponding event handle as input to `NvEncEncodePicture`. The NVIDIA Video Encoder Interface will signal this event when the hardware encoder finishes encoding the current input data. The client can

---

<sup>1</sup> To check the mode in which your board is running, run the command-line utility `nvidia-smi` (`nvidia-smi.exe` on Windows) included with the driver.

then call `NvEncLockBitstream` in non-blocking mode with `NV_ENC_LOCK_BITSTREAM::doNotWait` flag set to 1 to fetch the output data.

The client should call `NvEncUnregisterAsyncEvent` to unregister the Event handles before destroying the event objects. Whenever possible, NVIDIA recommends using the asynchronous mode of operation instead of synchronous mode.

A step-by-step control flow for asynchronous mode is as follows:

1. When working in asynchronous mode, the output sample must consist of an event + output buffer and clients must work in multi-threaded manner (D3D9 device should be created with `MULTITHREADED` flag).
2. The output buffers are allocated using `NvEncCreateBitstreamBuffer` API. The NVIDIA Video Encoder Interface will return an opaque pointer to the output memory in `NV_ENC_CREATE_BITSTREAM_BUFFER::bitstreambuffer`. This opaque output pointer should be used in `NvEncEncodePicture` and `NvEncLockBitstream/NvEncUnlockBitstream` calls. For accessing the output memory using CPU, client must call `NvEncLockBitstream` API. The number of IO buffers should be at least 4 + number of B frames.
3. The events are windows event handles allocated using Windows' `CreateEvent` API and registered using the function `NvEncRegisterAsyncEvent` before encoding. The registering of events is required only once per encoding session. Clients must unregister the events using `NvEncUnregisterAsyncEvent` before destroying the event handles. The number of event handles must be same as number of output buffers as each output buffer is associated with an event.
4. Client must create a secondary thread in which it can wait on the completion event and copy the bitstream data from the output sample. Client will have two threads: one is the main application thread which submits encoding work to NVIDIA Encoder while secondary thread waits on the completion events and copies the compressed bitstream data from the output buffer.
5. Client must send the output buffer and event in `NV_ENC_PIC_PARAMS::outputBitstream` and `NV_ENC_PIC_PARAMS::completionEvent` fields respectively as part of `NvEncEncodePicture` API call.
6. Client should then wait on the event on the secondary thread in the same order in which it has called `NvEncEncodePicture` calls irrespective of input buffer re-ordering (encode order! = display order). NVIDIA Encoder takes care of the reordering in case of B frames and should be transparent to the encoder clients.
7. When the event gets signalled client must send down the output buffer of sample event it was waiting on in `NV_ENC_LOCK_BITSTREAM::outputBitstream` field as part of `NvEncLockBitstream` call.
8. The NVIDIA Encoder Interface returns a CPU pointer and bitstream size in bytes as part of the `NV_ENC_LOCK_BITSTREAM`.

9. After copying the bitstream data, client must call `NvEncUnlockBitstream` for the locked output bitstream buffer.

**Note:**

1. The client will receive the event's signal and output buffer in the same order in which they were queued.
2. The `NV_ENC_LOCK_BITSTREAM::pictureType` notifies the output picture type to the clients.
3. Both, the input and output sample (output buffer and the output completion event) are free to be reused once the NVIDIA Video Encoder Interface has signalled the event and the client has copied the data from the output buffer.

## 6.2 SYNCHRONOUS MODE

This mode of operation is used for synchronous output buffer processing. In this mode the client makes a blocking call to the NVIDIA Video Encoder Interface to retrieve the output bitstream data from the encoder. The client sets the flag `NV_ENC_INITIALIZE_PARAMS::enableEncodeAsync` to 0 for operation in synchronous mode. The client then must call `NvEncEncodePicture` without setting a completion event handle. The client must call `NvEncLockBitstream` with flag `NV_ENC_LOCK_BITSTREAM::doNotWait` set to 0, so that the lock call blocks until the hardware encoder finishes writing the output bitstream. The client can then operate on the generated bitstream data and call `NvEncUnlockBitstream`. This is the only mode supported on Linux.

## 6.3 THREADING MODEL

To get maximum performance for encoding, the encoder client should create a separate thread to wait on events or when making any blocking calls to the encoder interface.

The client should avoid making any blocking calls from the main encoder processing thread. The main encoder thread should be used only for encoder initialization and to submit work to the HW Encoder using `NvEncEncodePicture` API, which is non-blocking.

Output buffer processing, such as waiting on the completion event in asynchronous mode or calling the blocking API's such as `NvEncLockBitstream/NvEncUnlockBitstream` in

synchronous mode, should be done in the secondary thread. This ensures that the main encoder thread is never blocked except when the encoder client runs out of resources.

It is also recommended to allocate many input and output buffers in order to avoid resource hazards and improve overall encoder throughput.

## 6.4 ENCODER FEATURES USING CUDA

Although the core video encoder hardware on GPU is completely independent of CUDA cores or graphics engine on the GPU, following encoder features internally use CUDA for hardware acceleration. Note that the impact of enabling these features on overall CUDA or graphics performance is minimal, and this list is provided purely for information purposes.

- Two-pass rate control modes for high quality presets
- Look-ahead
- All adaptive quantization modes
- Weighted prediction
- Encoding of RGB contents

# Chapter 7. MOTION ESTIMATION ONLY MODE

NVENC can be used as a hardware accelerator to perform motion search and generate motion vectors and mode information. The resulting motion vectors or mode decisions can be used, for example, in motion compensated filtering or for supporting other codecs not fully supported by NVENC or simply as motion vector hints for a custom encoder. The procedure to use the feature is explained below.

## 7.1 QUERY MOTION-ESTIMATION ONLY MODE CAPABILITY

Before using the motion-estimation (ME) only mode, the client should explicitly query the capability of the encoder to support ME only mode. For this, the client should do the following:

1. Specify the capability attribute as `NV_ENC_CAPS_SUPPORT_MEONLY_MODE` to query through the `NV_ENC_CAPS_PARAM::capsToQuery` parameter.
2. The client should call `NvEncGetEncoderCaps` to determine support for the required attribute.

`NV_ENC_CAPS_SUPPORT_MEONLY_MODE` indicates support of ME only mode in hardware.

0: ME only mode not supported.

1: ME only mode supported.

## 7.2 CREATE RESOURCES FOR INPUT/OUTPUT DATA

The client should allocate at least one buffer for the input picture by calling `NvEncCreateInputBuffer` API, and should also allocate one buffer for the reference frame by using `NvEncCreateInputBuffer` API. The client is responsible for filling in valid input data.

After input resources are created, client needs to allocate resources for the output data by using `NvEncCreateMVBuffer` API.

## 7.3 POPULATE ME ONLY MODE SETTINGS

The structure `NV_ENC_CODEC_CONFIG::NV_ENC_CONFIG_H264_MEONLY` provides the ability to control the partition types of motion vectors and modes returned by NVENC hardware. Specifically, the client can disable intra mode and/or specific MV partition sizes by setting the following flags:

```
NV_ENC_CONFIG_H264_MEONLY::disableIntraSearch
NV_ENC_CONFIG_H264_MEONLY::disablePartition16x16
NV_ENC_CONFIG_H264_MEONLY::disablePartition8x16
NV_ENC_CONFIG_H264_MEONLY::disablePartition16x8
NV_ENC_CONFIG_H264_MEONLY::disablePartition8x8
```

The API also exposes a parameter `NV_ENC_CONFIG::NV_ENC_MV_PRECISION` to control the precision of motion vectors returned by the hardware. For full-pel precision, client must ignore two LSBs of the motion vector. For sub-pel precision, the two LSBs of the motion vector represent fractional part of the motion vector. To get motion vectors for each macroblock, it is recommended to disable intra modes by setting `NV_ENC_CONFIG_H264_MEONLY::disableIntraSearch = 1` and let NVENC decide the optimal partition sizes for motion vectors.

## 7.4 RUN MOTION ESTIMATION

The client should create an instance of `NV_ENC_MEONLY_PARAMS`.

The pointers of the input picture buffer and the reference frame buffer need to be fed to `NV_ENC_MEONLY_PARAMS::inputBuffer` and `NV_ENC_MEONLY_PARAMS::referenceFrame` respectively.

The pointer returned by `NvEncCreateMVBuffer` API in the `NV_ENC_CREATE_MV_BUFFER::mvBuffer` field needs to be fed to `NV_ENC_MEONLY_PARAMS::mvBuffer`.

In order to operate in asynchronous mode, the client should create an event and pass this event in `NV_ENC_MEONLY_PARAMS::completionEvent`. This event will be signaled upon completion of motion estimation. Each output buffer should be associated with a distinct event pointer.

Client should call `NvEncRunMotionEstimationOnly` to run the motion estimation on hardware encoder.

For asynchronous mode client should wait for motion estimation completion signal before reusing output buffer and application termination.

Client must lock `NV_ENC_CREATE_MV_BUFFER::mvBuffer` using `NvEncLockBitstream` to get the motion vector data.

Finally, `NV_ENC_LOCK_BITSTREAM::bitstreamBufferPtr` which contains the output motion vectors should be typecast to `NV_ENC_H264_MV_DATA*/NV_ENC_HEVC_MV_DATA*` for H.264/HEVC respectively. Client should then unlock `NV_ENC_CREATE_MV_BUFFER::mvBuffer` by calling `NvEncUnlockBitstream`.

## 7.5 ENABLING MOTION ESTIMATION FOR STEREO USECASES

For stereo use cases where in two views need to be processed we suggest the following approach for better performance and quality of motion vectors:

1. Client should create single encode session.
2. The client should kick-off the processing of left and right views on separate threads.
3. The client should set `NV_ENC_MEONLY_PARAMS::viewID` to 0 and 1 for left and right views.
4. The main thread should wait for completion of the threads which have been kicked off NVENC for left and right views.

## 7.6 RELEASE THE CREATED RESOURCES

Once the usage of motion estimation is done, the client should call `NvEncDestroyInputBuffer` to destroy the input picture buffer and the reference frame buffer and should call `NvEncDestroyMVBuffer` to destroy the motion vector data buffer.

# Chapter 8. ADVANCED FEATURES AND SETTINGS

## 8.1 LOOK-AHEAD

Look-ahead improves the video encoder's rate control accuracy by enabling the encoder to buffer the specified number of frames, estimate their complexity and allocate the bits appropriately among these frames proportional to their complexity. This also dynamically allocates B and P frames.

To use this feature, the client must follow these steps:

1. The availability of the feature in the current hardware can be queried using `NvEncGetEncodeCaps` and checking for `NV_ENC_CAPS_SUPPORT_LOOKAHEAD`.
2. Look-ahead needs to be enabled during initialization by setting `NV_ENC_INITIALIZE_PARAMS::encodeconfig->rcParams.enableLookahead = 1`.
3. The number of frames to be looked ahead should be set in `NV_ENC_INITIALIZE_PARAMS::encodeconfig->rcParams.lookaheadDepth` which can be up to 32.
4. By default, look-ahead enables adaptive insertion of intra frames and B frames. They can however be disabled by setting `NV_ENC_INITIALIZE_PARAMS::encodeconfig->rcParams.disableIadapt` and/or `NV_ENC_INITIALIZE_PARAMS::encodeconfig->rcParams.disableBadapt` to 1.
5. When the feature is enabled, frames are queued up in the encoder and hence `NvEncEncodePicture` will return `NV_ENC_ERR_NEED_MORE_INPUT` until the encoder has sufficient number of input frames to satisfy the look-ahead requirement. Frames should be continuously fed in until `NvEncEncodePicture` returns `NV_ENC_SUCCESS`.

## 8.2 B-FRAMES AS REFERENCE

Using B frame as a reference improves subjective and objective encoded quality with no performance impact. Hence the users enabling multiple B frames are strongly recommended to enable this feature.

To use the feature, follow these steps:

1. Query availability of the feature using `NvEncGetEncodeCaps` API and checking for `NV_ENC_CAPS_SUPPORT_BFRAME_REF_MODE` in the return value.
2. During encoder initialization, set `NV_ENC_CONFIG_H264:: useBFramesAsRef = NV_ENC_BFRAME_REF_MODE_MIDDLE`. This will set the  $(N/2)^{\text{th}}$  B frame as reference where  $N$  = number of B frames. In case  $N$  is odd, then  $(N-1)/2^{\text{th}}$  frame will be picked up as reference.

## 8.3 RECONFIGURE API

`NvEncReconfigureEncoder` allows clients to change the encoder initialization parameters in `NV_ENC_INITIALIZE_PARAMS` without closing existing encoder session and re-creating a new encoding session. This helps clients avoid the latency introduced due to destruction and re-creation of the encoding session. This API is useful in scenarios which are prone to instabilities in transmission mediums during video conferencing, game streaming etc.

Using this API clients can change parameters like bit-rate, frame-rate, resolution dynamically using the same encode session. The reconfigured parameters are passed via `NV_ENC_RECONFIGURE_PARAMS::reInitEncodeParams`.

However, The API currently doesn't support reconfiguration of all parameters, some of which are listed below:

- Changing the GOP structure (`NV_ENC_CONFIG_H264::idrPeriod`, `NV_ENC_CONFIG::gopLength`, `NV_ENC_CONFIG:: frameIntervalP`)
- Changing from synchronous mode of encoding to asynchronous mode and vice-versa.
- Changing `NV_ENC_INITIALIZE_PARAMS:: maxEncodeWidth` and `NV_ENC_INITIALIZE_PARAMS::maxEncodeHeight`.
- Changing picture type decision in `NV_ENC_INITIALIZE_PARAMS::enablePTD`.
- Changing bit-depth.

The API would fail if any attempt is made to reconfigure the parameters which is not supported.

Resolution change is possible only if `NV_ENC_INITIALIZE_PARAMS::maxEncodeWidth` and `NV_ENC_INITIALIZE_PARAMS::maxEncodeHeight` of is set while creating encoder session.

If the client wishes to change the resolution using this API, it is advisable to force the next frame following the reconfiguration as an IDR frame by setting `NV_ENC_RECONFIGURE_PARAMS::forceIDR` to 1.

If the client wishes to reset the internal rate control states, set `NV_ENC_RECONFIGURE_PARAMS::resetEncoder` to 1.

## 8.4 ADAPTIVE QUANTIZATION (AQ)

This feature improves visual quality by adjusting encoding QP (on top of QP evaluated by the Rate Control Algorithm) based on spatial and temporal characteristics of the sequence. The current SDK support two flavors of AQ which are explained as follows:

### 8.4.1 Spatial AQ

Spatial AQ mode adjusts the QP values based on spatial characteristics of the frame. Since the low complexity flat regions are visually more perceptible to quality differences than high complexity detailed regions, extra bits are allocated to flat regions of the frame at the cost of the regions having high spatial detail. Although spatial AQ improves the perceptible visual quality of the encoded video, the required bit redistribution results in PSNR drop in most of the cases. Therefore, during PSNR-based evaluation, this feature should be turned off.

To use spatial AQ, follow these steps in your application.

- Spatial AQ can be enabled during initialization by setting `NV_ENC_INITIALIZE_PARAMS::encodeconfig->rcParams.enableAQ = 1`.
- The intensity of QP adjustment can be controlled by setting `NV_ENC_INITIALIZE_PARAMS::encodeconfig->rcParams.aqStrength` which ranges from 1 (least aggressive) to 15 (most aggressive). If not set, strength is auto selected by driver.

### 8.4.2 Temporal AQ

Temporal AQ tries to adjust encoding QP (on top of QP evaluated by the rate control algorithm) based on temporal characteristics of the sequence. Temporal AQ improves the quality of encoded frames by adjusting QP for regions which are constant or have low motion across frames but have high spatial detail, such that they become better reference for future frames. Allocating extra bits to such regions in reference frames is better than allocating them to the residuals in referred frames because it helps improve the overall

encoded video quality. If majority of the region within a frame has little or no motion, but has high spatial details (e.g. high-detail non-moving background) enabling temporal AQ will benefit the most.

One of the potential disadvantages of temporal AQ is that enabling temporal AQ *may* result in high fluctuation of bits consumed per frame within a GOP. I/P-frames will consume more bits than average P-frame size and B-frames will consume lesser bits. Although target bitrate will be maintained at the GOP level, the frame size will fluctuate from one frame to next within a GOP more than it would without temporal AQ. If a strict CBR profile is required for every frame size within a GOP, it is not recommended to enable temporal AQ. Additionally, since some of the complexity estimation is performed in CUDA, there may be some performance impact when temporal AQ is enabled.

To use temporal AQ, follow these steps in your application.

1. Query the availability of temporal AQ for the current hardware by calling the API `NvEncGetEncodeCaps` and checking for `NV_ENC_CAPS_SUPPORT_TEMPORAL_AQ`.
2. If supported, temporal AQ can be enabled during initialization by setting `NV_ENC_INITIALIZE_PARAMS::encodeconfig->rcParams.enableTemporalAQ = 1`.

Temporal AQ is supported for H.264 only.

Temporal AQ uses CUDA pre-processing and hence requires CUDA processing power, depending upon resolution and content.

Enabling temporal AQ may result in very minor degradation in encoder performance.

## 8.5 HIGH BIT DEPTH ENCODING

All NVIDIA GPUs support 8-bit encoding (RGB/YUV input with 8-bit precision). NVIDIA GPUs starting from Pascal generation support high-bit-depth HEVC encoding (HEVC main-10 profile with 10-bit input precision). To encode 10-bit content the following steps are to be followed.

1. The availability of the feature can be queried using `NvEncGetEncodeCaps` and checking for `NV_ENC_CAPS_SUPPORT_10BIT_ENCODE`.
2. Create the encoder session with `NV_ENC_HEVC_PROFILE_MAIN10_GUID`.
3. During encoder initialization, set `encodeConfig->encodeCodecConfig.hevcConfig.pixelBitDepthMinus8 = 2`.
4. The input surface format needs be set to `NV_ENC_BUFFER_FORMAT_YUV420_10BIT` OR `NV_ENC_BUFFER_FORMAT_ABGR_10` or `NV_ENC_BUFFER_FORMAT_ARGB10` or `NV_ENC_BUFFER_FORMAT_YUV444_10BIT`, depending upon nature of input.
5. Other encoding parameters such as preset, rate control mode etc. can be set as desired.

## 8.6 WEIGHTED PREDICTION

Weighted prediction involves calculation of a multiplicative weighting factor and an additive offset to the motion compensated prediction. Weighted prediction provides significant quality gain for contents having illumination changes. NVENC API supports weighed prediction for HEVC and H.264 starting from Pascal generation GPUs.

The following steps need to be followed for enabling weighted prediction.

1. The availability of the feature can be queried using `NvEncGetEncodeCaps` and checking for `NV_ENC_CAPS_SUPPORT_WEIGHTED_PREDICTION`.
2. During encoder initialization, set `NV_ENC_INITIALIZE_PARAMS::enableWeightedPrediction = 1`.

Weighted prediction is not supported if the encode session is configured with B frames.

Weighted prediction uses CUDA pre-processing and hence requires CUDA processing power, depending upon resolution and content.

Enabling weighted prediction may also result in very minor degradation in encoder performance.

## 8.7 LONG-TERM REFERENCE

NVENC API provides the functionality to mark and use specific frames as long-term reference (LTR) frames, which can later be used as reference for encoding the current picture. This helps in error concealment where in the client decoders can predict from the long-term reference frame in case of an intermediate frame's data loss. The feature is useful in video streaming applications to recover from frame losses at the receiver.

Following steps are to be followed to enable the feature.

1. During encoder initialization,
  - a. For H.264, set `NV_ENC_CONFIG_H264:enableLTR = 1`
  - b. For HEVC, set `NV_ENC_CONFIG_HEVC:enableLTR = 1`
2. The maximum number of long term reference pictures supported in the current hardware can be queried using `NvEncGetEncoderCaps` and checking for `NV_ENC_CAPS_NUM_MAX_LTR_FRAMES`.

During normal encoding operation, following steps need to be followed to mark specific frame(s) as LTR frame(s).

1. Configure the number of LTR frames:

- a. For H.264, set `NV_ENC_CONFIG_H264:ltrNumFrames`
  - b. For HEVC, set `NV_ENC_CONFIG_HEVC:ltrNumFrames`
2. The client can mark any frame as LTR by setting `NV_ENC_PIC_PARAMS_H264::ltrMarkFrame = 1` OR `NV_ENC_PIC_PARAMS_HEVC::ltrMarkFrame = 1` for H.264 and HEVC respectively. Each LTR frame needs to be assigned an LTR frame index. This value should be between 0 and `ltrNumFrames - 1`.
3. The LTR frame index can be assigned by setting `NV_ENC_PIC_PARAMS_H264::ltrMarkFrameIdx` OR `NV_ENC_PIC_PARAMS_HEVC::ltrMarkFrameIdx` for H264 and HEVC respectively.

The frames earlier marked as long term reference frames can be used for prediction of the current frame in the following manner:

1. The LTR frames that are to be used for reference have to be specified using `NV_ENC_PIC_PARAMS_H264::ltrUseFrameBitmap` OR `NV_ENC_PIC_PARAMS_HEVC::ltrUseFrameBitmap` for H.264 and HEVC respectively. The bit location specifies the LTR frame index of the frame that will be used as reference. Note that NVENC can only use one frame as a reference frame for prediction. Therefore, only one bit should be set in `ltrUseFrameBitmap` at any time.
2. Setting `NV_ENC_PIC_PARAMS_H264::ltrUsageMode = 1` and `NV_ENC_PIC_PARAMS_HEVC::ltrUsageMode = 1` for H.264 and HEVC respectively will clear out all the short term (other than current) reference pictures.

The current SDK does not support LTR when the encoding session is configured with B frames.

## 8.8 EMPHASIS MAP

The emphasis map feature in NVENCODE API provides a way to specify regions in the frame to be encoded at varying levels of quality, at macroblock-level granularity. Depending upon the actual *emphasis level* for each macroblock, the encoder applies an adjustment to the quantization parameter used to encode that macroblock. The value of this adjustment depends on the following factors:

1. Absolute value of the QP as decided by the rate control algorithm, depending upon the rate control constraints. In general, for a given emphasis level, higher the QP determined by the rate control, higher the (negative) adjustment.
2. Emphasis level value for the macroblock.

Note that the QP adjustment is performed after the rate control algorithm has run. So there is a possibility of VBV/rate violations when using this feature.

Emphasis level map is useful when the client has prior knowledge of the image complexity (e.g. NVFBC's Classification Map feature) and encoding those high-complexity areas at higher quality (lower QP) is important, even at the possible cost of violating bitrate/VBV buffer size constraints. This feature is not supported when AQ (Spatial/Temporal) is enabled.

Follow these steps to enable the feature.

1. Query availability of the feature using `NvEncGetEncodeCaps` API and checking for `NV_ENC_CAPS_SUPPORT_EMPHASIS_LEVEL_MAP`.
2. Set `NV_ENC_RC_PARAMS::qpMapMode = NV_ENC_QP_MAP_EMPHASIS`.
3. Fill up the `NV_ENC_PIC_PARAMS::qpDeltaMap` (which is a signed byte array containing value per macroblock in raster scan order for the current picture) with a value from `enum NV_ENC_EMPHASIS_MAP_LEVEL`.

As explained above, higher values of `NV_ENC_EMPHASIS_MAP_LEVEL` imply higher (negative) adjustment made to the QP to *emphasize* quality of that macroblock. The user can choose higher emphasis level for the regions (s)he wants to encode with a higher quality.

# Chapter 9. RECOMMENDED NVENC SETTINGS

NVIDIA hardware video encoder is used for several purposes in various applications. Some of the common applications include: Video-recording (archiving), game-casting (broadcasting/multicasting video gameplay online), transcoding (live and video-on-demand) and streaming (games or live content). Each of these use-cases has its unique requirements for quality, bitrate, latency tolerance, performance constraints etc. Although NVIDIA encoder interface provides flexibility to control the settings with many API's, below table can be used as a general guideline for recommended settings for some of the popular use-cases to deliver the best encoded bit stream quality. These recommendations are particularly applicable to GPUs based on second generation Maxwell architecture beyond. For earlier GPUs (Kepler and first-generation Maxwell), it is recommended that clients use the Table 1 as a starting point and adjust the settings to achieve appropriate performance-quality tradeoff.

Table 1. Recommended NVENC settings for various use-cases

Use-case	Recommended settings for optimal quality and performance
Recording/Archiving	<ul style="list-style-type: none"> <li>➤ High quality preset</li> <li>➤ Rate control mode = VBR</li> <li>➤ Very large VBV buffer size (4 seconds)</li> <li>➤ B Frames*</li> <li>➤ Look-ahead</li> <li>➤ B frame as reference</li> <li>➤ Finite GOP length (2 seconds)</li> <li>➤ Adaptive quantization (AQ) enabled**</li> </ul>
Game-casting & cloud transcoding	<ul style="list-style-type: none"> <li>➤ High quality preset</li> <li>➤ Rate control mode = CBR</li> <li>➤ Medium VBV buffer size (1 second)</li> <li>➤ B Frames*</li> <li>➤ Look-ahead</li> <li>➤ B frame as reference</li> <li>➤ Finite GOP length (2 seconds)</li> <li>➤ Adaptive quantization (AQ) enabled**</li> </ul>
Low-latency use cases like game-streaming, video conferencing etc.	<ul style="list-style-type: none"> <li>➤ Low-latency high quality preset</li> <li>➤ Rate control mode = CBR</li> <li>➤ Very low VBV buffer size (single frame)</li> <li>➤ No B Frames</li> <li>➤ Infinite GOP length</li> <li>➤ Long term reference pictures</li> <li>➤ Intra refresh</li> <li>➤ Adaptive quantization (AQ) enabled**</li> </ul>

\*: Recommended for low motion games and natural video.

\*\* : Recommended on second generation Maxwell GPUs and above.

## 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.

## HDMI

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

## 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

© 2011-2018 NVIDIA Corporation. All rights reserved.