SMPTE ST 2042-4:2018
Revision of ST 2042-4:2016

SMPTE STANDARD

Mapping a VC-2 Stream into the MXF Generic Container

Table of Contents

Approved November 20, 2018

Foreword

SMPTE (the Society of Motion Picture and Television Engineers) is an internationally-recognized standards developing organization. Headquartered and incorporated in the United States of America, SMPTE has members in over 80 countries on six continents. SMPTE's Engineering Documents, including Standards, Recommended Practices, and Engineering Guidelines, are prepared by SMPTE's Technology Committees. Participation in these Committees is open to all with a bona fide interest in their work. SMPTE cooperates closely with other standards-developing organizations, including ISO, IEC and ITU.

SMPTE Engineering Documents are drafted in accordance with the rules given in its Standards Operations Manual.

SMPTE ST 2042-4 was prepared by Technology Committee 31FS.

Intellectual Property

At the time of publication no notice had been received by SMPTE claiming patent rights essential to the implementation of this Engineering Document. However, attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. SMPTE shall not be held responsible for identifying any or all such patent rights.

1 Scope

This document specifies the mapping of VC-2 coded pictures into MXF. Specifically, it defines how to map a VC-2 stream (or part of a stream) into a Picture Element of a MXF generic container.

This standard specifies the Key, Length and Value fields of a VC-2 Picture Element. This standard also defines the Essence Container Label and Picture Essence Compression Label, a VC-2 Sub-Descriptor and the mapping of information from the VC-2 stream to the CDCI Picture Essence Descriptor.

2 Conformance Notation

Normative text is text that describes elements of the design that are indispensable or contains the conformance language keywords: "shall", "should", or "may". Informative text is text that is potentially helpful to the user, but not indispensable, and can be removed, changed, or added editorially without affecting interoperability. Informative text does not contain any conformance keywords.

All text in this document is, by default, normative, except: any section explicitly labeled as "Informative" or individual paragraphs that start with "Note: ”

The keywords "shall" and "shall not" indicate requirements strictly to be followed in order to conform to the document and from which no deviation is permitted.

The keywords, "should" and "should not" indicate that, among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others; or that a certain course of action is preferred but not necessarily required; or that (in the negative form) a certain possibility or course of action is deprecated but not prohibited.

The keywords "may" and "need not" indicate courses of action permissible within the limits of the document.

The keyword “reserved” indicates a provision that is not defined at this time, shall not be used, and may be defined in the future. The keyword “forbidden” indicates “reserved” and in addition indicates that the provision will never be defined in the future.

A conformant implementation according to this document is one that includes all mandatory provisions ("shall") and, if implemented, all recommended provisions ("should") as described. A conformant implementation need not implement optional provisions ("may") and need not implement them as described.

Unless otherwise specified, the order of precedence of the types of normative information in this document shall be as follows: Normative prose shall be the authoritative definition; Tables shall be next; then formal languages; then figures; and then any other language forms.

3 Normative References

The following standards contain provisions that, through reference in this text, constitute provisions of this standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this recommended practice are encouraged to investigate the possibility of applying the most recent edition of the standards indicated below.

SMPTE ST 377-1:2011, Material Exchange Format (MXF) – File Format Specification

SMPTE ST 379-1:2009, Material Exchange Format (MXF) — MXF Generic Container

SMPTE ST 379-2:2010, Material Exchange Format (MXF) – MXF Constrained Generic Container

SMPTE ST 2042-1:2017, VC-2 Video Compression

4 Glossary of Acronyms, Terms and Data Types

The general glossary of acronyms, terms and data types used in the MXF specification is given in SMPTE ST 377-1 and is supplemented in SMPTE ST 379-1 and SMPTE ST 379-2. These glossaries are not repeated here to avoid any divergence of meaning. In this document, terms are used as defined by SMPTE ST 377-1 in preference to any definition specific to SMPTE ST 379-1 or SMPTE ST 379-2.

4.1 VC-2 Terms and Notation

Definitions and representations of terms, abbreviations and symbols relating to VC-2 are given in SMPTE ST 2042-1.

Note: For example: `video_parameters[frame_height]` refers to the `frame_height` property of the `video_parameters` map, after the Sequence Header has been parsed. VC-2 defines a process for populating the `video_parameters` map, which involves first initializing the properties to default values that are then overridden if a value for a property appears in the Sequence Header. Therefore the VC-2 properties referenced in this standard will always have values defined.

5 Overview (Informative)

A summary is given below of some of the features of this standard that could be relevant when, for example, testing or deploying an implementation.

- A recommended Operating Mode is defined which places constraints on the wrapped VC-2 stream in order to simplify handling of the content. The use of this Operating Mode is optional.
- Even in the recommended Operating Mode the structure of the wrapped VC-2 stream is only partially constrained. For example: Sequence Headers can be repeated within an Edit Unit; Padding and Auxiliary data units can be included.
- This standard places no restrictions on the video pictures that can be encoded by the wrapped VC-2 stream. For example: no restrictions are placed on the video frame size, frame rate or color specification.
- This standard places no restrictions on the wavelet transform (coding) parameters used in the wrapped VC-2 stream. The values permitted are not restricted by this standard nor does this standard require that any of the wavelet transform parameters remain constant throughout the wrapped VC-2 stream.
- The VC-2 stream can be clip-wrapped or frame-wrapped.
- The size of the Edit Units can be constant throughout or variable.

6 VC-2 Stream Syntax (Informative)

SMPTE ST 2042-1 specifies the compressed stream syntax and reference decoder operations for a video compression system known as VC-2. A VC-2 stream can contain multiple, concatenated VC-2 sequences. The data contained in a VC-2 sequence corresponds to a single video sequence with constant video parameters (but not necessarily constant wavelet transform (coding) parameters). A VC-2 sequence can be extracted from a VC-2 stream and decoded as an independent entity.

The VC-2 sequence structure is illustrated in Figure 1 and is identical to that illustrated in SMPTE ST 2042-1 Figure 10.2. The VC-2 sequence is composed of:

- Parse Info headers: components of the VC-2 stream that provide information on how to parse the stream. They support navigating through the stream without having to decode every frame.
- Data Units: Sequence Header; Picture; Auxiliary; Padding.

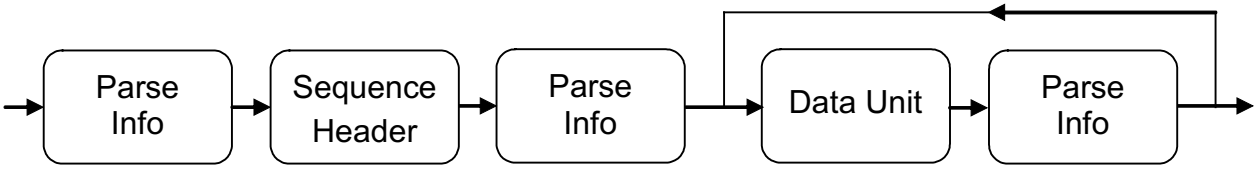

Figure 1 – Illustration of the VC-2 Sequence Structure

7 The Wrapped VC-2 Stream

7.1 Edit Units of the Wrapped VC-2 Stream

When concatenated consecutively the Edit Units shall form a valid and complete VC-2 stream except that:

- The first Parse Info header in the wrapped VC-2 stream need not have a `previous_parse_offset` value of zero.
- An end of sequence Parse Info header need not be present at the end of the wrapped VC-2 stream.

Each Edit Unit shall be a portion, or the entirety, of a single valid VC-2 sequence and shall begin with a Parse Info header, followed by a Sequence Header.

Within an Edit Unit there shall be all of the VC-2 Picture data units for a single video frame and no other VC-2 Picture data units.

7.2 Understanding the Wrapped VC-2 Stream (Informative)

The wrapped VC-2 stream (which is the consecutive concatenation of all of the Edit Units) could contain data from more than one VC-2 sequence. In Operating Mode A (see Section 13) there is exactly one VC-2 sequence for each Edit Unit.

In Operating Mode A the wrapped VC-2 stream will always be a valid and complete VC-2 stream.

If video is VC-2 coded as frames there will be one VC-2 Picture data unit per Edit Unit.

If video is VC-2 coded as fields there will be two VC-2 Picture data units per Edit Unit and the earlier of the two will have an even `picture_number`. This precludes the case of wrapping individual fields described by SMPTE ST 379-1 and SMPTE ST 379-2.

In general a VC-2 stream can contain both video coded as frames and video coded as fields. In each VC-2 sequence the method being used is signaled by `picture_coding_mode` in the Sequence Header. In Operating Mode A all of the Sequence Headers must be byte-for-byte identical throughout the entire wrapped VC-2 stream and so in Operating Mode A the video will either be coded as frames throughout or as fields throughout. Even if Operating Mode A is not in use and the Sequence Headers are not identical throughout the wrapped VC-2 stream then the considerations in Section 7.3 mean that the video must be coded as frames throughout or as fields throughout (so that the wrapped VC-2 stream can be accurately described by the Essence Descriptor / Sub-Descriptor).

An Edit Unit could contain additional Sequence Header data units, as well as Auxiliary data units and Padding data units. SMPTE ST 2042-1 requires that all Sequence Headers in a VC-2 sequence are byte-for-byte identical and so all of the Sequence Headers in an Edit Unit will be byte-for-byte identical.

The sizes of the Edit Units can be constant or variable.

7.3 Wrapping a VC-2 Stream Comprised of Multiple VC-2 Sequences (Informative)

If the wrapped VC-2 stream contains multiple VC-2 sequences then the Sequence Headers could be non-identical throughout the wrapped VC-2 stream. Therefore, for example, `video_parameters` could not be fixed throughout the wrapped VC-2 stream. However, Section 11 requires the use of a single CDCI Picture Essence Descriptor and a single VC-2 Sub-Descriptor and these must describe the entire wrapped VC-2 stream accurately. Therefore the amount of permitted variation in the Sequence Headers is limited.

VC-2 properties mapped to mandatory elements of the Essence Descriptor / Sub-Descriptor will have to be fixed throughout the wrapped VC-2 stream. VC-2 properties mapped to optional elements of the Essence Descriptor / Sub-Descriptor can potentially have more than one value across the wrapped VC-2 stream: if this is the case then these optional elements of the Essence Descriptor / Sub-Descriptor will have to be omitted.

In Operating Mode A all Sequence Headers will be identical throughout the wrapped VC-2 stream.

8 KLV Coding of the VC-2 Stream

8.1 Picture Element Key

The value of the Picture Element Key shall be 06.0E.2B.34.01.02.01.01.0D.01.03.01.15.xx.yy.zz where the values of bytes 14 and 16 shall be as defined by SMPTE ST 379-1 and SMPTE ST 379-2, and the value of byte 15 shall be as given in Table 1.

Note 1: Byte 14 is a count of the Picture Elements in the Picture Item.

Note 2: Byte 16 is the Essence Element Number (an identifier) of this Picture Element in the Picture Item.

Table 1 – Byte 15 of the VC-2 Picture Element Key

Byte 15 value (hex)	Meaning
10h	MXF-GC Frame-wrapped VC-2 Picture Essence
11h	MXF-GC Clip-wrapped VC-2 Picture Essence

8.2 Picture Element Length (Informative)

It is MXF convention to use 4-byte BER long-form encoded (i.e. 83h.xx.yy.zz) lengths for frame-based wrapping and to use 8-byte BER long-form encoded (i.e. 87h.aa.bb.cc.dd.ee.ff.gg) lengths for clip-based wrapping. SMPTE ST 377-1 allows the use of any length of BER long-form encoding up to and including 9 bytes.

8.3 Picture Element Value (Informative)

SMPTE ST 379-1 and SMPTE ST 379-2 define how the value field of each Picture Element KLV triplet is composed for both frame-wrapping and clip-wrapping:

- In frame-wrapping the value field of each Picture Element KLV triplet is a single Edit Unit.
- In clip-wrapping the value field of the Picture Element KLV triplet is a concatenation of all Edit Units: that is, it contains the entire wrapped VC-2 stream.

9 VC-2 Essence Container Label

The VC-2 Essence Container Label shall be 06.0E.2B.34.04.01.01.0D.0D.01.03.01.02.15.xx.00 where byte 15 shall be as shown in Table 2.

Table 2 – Byte 15 of VC-2 Essence Container Label

Byte 15 value (hex)	Meaning
01h	MXF-GC Frame-wrapped VC-2 Picture Essence
02h	MXF-GC Clip-wrapped VC-2 Picture Essence

10 VC-2 Picture Essence Compression Label

The VC-2 Picture Essence Compression Label shall be: 06.0E.2B.34.04.01.01.0D.04.01.02.02.03.03.01.00

11 Essence Descriptors for VC-2

The VC-2 Picture Essence shall be described by exactly one CDCI Picture Essence Descriptor.

The elements in the CDCI Picture Essence Descriptor shall be set as defined in Annex B.

Note 1: The use of the CDCI Picture Essence Descriptor applies even when wrapping VC-2 encoded RGB pictures. This is because the CDCI Picture Essence Descriptor is being used to describe the YC_1C_2 interface to VC-2. A VC-2 encoder will treat RGB pictures and CDCI pictures in the same way.

Note 2: Descriptor elements that describe the physical layout of bytes of data are omitted because they describe byte packing of uncompressed data. This is irrelevant in the case of VC-2.

Note 3: VC-2 cannot encode an alpha channel.

11.1 VC-2 Sub-Descriptor

The VC-2 Sub-Descriptor shall be a subclass of Sub-Descriptor.

The Key for the VC-2 Sub-Descriptor shall be 06.0E.2B.34.02.xx.01.01.0D.01.01.01.01.01.74.00 where byte 6 shall be as defined in SMPTE ST 377-1.

Note: The value of byte 6 depends on the lengths of the values of all of the elements in the set. The lengths of the elements defined means that the value of byte 6 will be 53h.

The additional elements defined by the VC-2 Sub-Descriptor shall be as specified in Table 3.

Table 3 – Additional Elements defined by the VC-2 Sub-Descriptor

Item Name	Type	Len	Item UL	Req ?	Meaning
VC-2 Major Version	UInt8	1	06.0E.2B.34 01.01.01.0E 04.01.06.07 01.00.00.00	Req	VC-2 Major Version Number as defined in SMPTE ST 2042-1
VC-2 Minor Version	UInt8	1	06.0E.2B.34 01.01.01.0E 04.01.06.07 02.00.00.00	Req	VC-2 Minor Version Number as defined in SMPTE ST 2042-1
VC-2 Profile	UInt8	1	06.0E.2B.34 01.01.01.0E 04.01.06.07 03.00.00.00	Req	VC-2 Profile as defined in SMPTE ST 2042-1
VC-2 Level	UInt8	1	06.0E.2B.34 01.01.01.0E 04.01.06.07 04.00.00.00	Req	VC-2 Level as defined in SMPTE ST 2042-1
VC-2 Wavelet Filters	Array of UInt8	8+n	06.0E.2B.34 01.01.01.0E 04.01.06.07 05.00.00.00	Opt	The distinct values of **`state`**`[wavelet_index]` across the entire wrapped VC-2 stream (each value used in the wrapped VC-2 stream appears only once in this array). This identifies the wavelet filters used.
VC-2 Sequence Headers Identical	Boolean	1	06.0E.2B.34 01.01.01.0E 04.01.06.07 06.00.00.00	Opt	A flag to indicate whether all VC-2 Sequence Headers in the entire wrapped VC-2 stream are byte-for-byte identical
VC-2 Edit Units Are Complete Sequences	Boolean	1	06.0E.2B.34 01.01.01.0E 04.01.06.07 07.00.00.00	Opt	A flag to indicate whether every Edit Unit comprises a single valid VC-2 sequence in its entirety

11.1.1 Understanding the VC-2 Sub-Descriptor (Informative)

The variable length coding used in VC-2 means that the values of profile, level, major and minor version and the wavelet indexes can be of arbitrary size. There are currently no values defined for these properties that would exceed the size of a UInt8. However, if a future version of VC-2 uses values for these properties that cannot be stored in the types defined here, this standard will need to be updated.

At the time of writing some VC-2 decoders do not support the entire restricted set of wavelet filters specified by the Default Levels defined in SMPTE ST 2042-2. The wavelet filters used in the wrapped VC-2 stream are surfaced in the Sub-Descriptor to enable an MXF decoder to identify an appropriate VC-2 decoder.

The value of **state**[wavelet_index] can vary from one Picture data unit to the next. Therefore, to populate the VC-2 Wavelet Filters array the entire VC-2 stream must be read.

If VC-2 Sequence Headers Identical is True then the wrapped VC-2 stream could be converted to a single VC-2 sequence. In Operating Mode A this element will be present and have the value True.

If VC-2 Edit Units Are Complete Sequences is True then the wrapped VC-2 stream is a valid and complete VC-2 stream. Additionally, Edit Units can be reordered and the resultant VC-2 stream will still be valid and complete. In Operating Mode A this element will be present and have the value True.

12 Index Tables (Informative)

VC-2 is intra-frame coded. Therefore the features of the MXF Index Entry Array that are provided to support indexing of inter-frame coded essence (for example: prediction flags, Temporal Offset, Key-Frame Offset) will use the default values given in SMPTE ST 377-1.

The semantics for Index Table Stream Offsets for clip-wrapped and frame-wrapped essence are different; see SMPTE ST 379-1 and SMPTE ST 379-2.

13 Operating Mode A

An implementation may choose to implement the restrictions which are set out below for “Operating Mode A”. If all of these constraints are implemented then the implementation shall be said to conform to “Operating Mode A” of this standard.

In Operating Mode A (in addition to the other normative provisions of this standard):

- Each Edit Unit shall comprise a single valid VC-2 sequence in its entirety.
- All Sequence Headers shall be byte-for-byte identical throughout the entire wrapped VC-2 stream.
- The above characteristics shall be signaled using the relevant VC-2 Sub-Descriptor elements.

13.1 Understanding Operating Mode A (Informative)

In Operating Mode A, each Edit Unit necessarily begins with a Parse Info header (with a previous_parse_offset value of zero) followed by a Sequence Header and ends with an end of sequence Parse Info header. This means that the wrapped VC-2 stream is always a valid and complete VC-2 stream.

The following elements will therefore be present in the VC-2 Sub-Descriptor and will have the following values:

- VC-2 Edit Units Are Complete Sequences: True
- VC-2 Sequence Headers Identical: True

If these elements are present in the VC-2 Sub-Descriptor and they have these values then it is necessarily the case that Operating Mode A is being used.

Operating Mode A places no restrictions on the wavelet transform (coding) parameters used in the wrapped VC-2 stream. This means that any of the wavelet transform parameters could have more than one value across the wrapped VC-2 stream: this is a scenario which could not be (efficiently) supported by all VC-2 decoders.

Annex A Using the Wrapped VC-2 Stream (Informative)

This standard defines the mapping of a valid VC-2 stream (or part thereof) into a generic container in a MXF file. As such, it is assumed that the input to the device performing the MXF wrapping is a valid VC-2 stream (or part thereof). The definition of an Edit Unit of a wrapped VC-2 stream allows video decoding to begin at any Edit Unit.

If Operating Mode A is used then any frame-level editing operations such as cutting, splicing, or re-ordering of Edit Units can be performed and the resulting VC-2 stream will still be entirely valid. As such, the use of Operating Mode A is recommended.

If Operating Mode A is not used then when the MXF file is unwrapped the VC-2 data will not form a valid and complete VC-2 stream in all circumstances. This could be acceptable. However, this could cause some difficulties for the device taking this VC-2 data as input (such as a VC-2 decoder). This annex sets out some of the scenarios which could occur and the minor alterations needed to the VC-2 data in order to form a valid and complete VC-2 stream.

Implementers of this standard are advised to familiarize themselves with the VC-2 sequence syntax defined in SMPTE ST 2042-1.

A.1 Scenarios

A.1.1 Unwrapping the entire wrapped VC-2 stream

If the entire wrapped VC-2 stream is unwrapped (that is, if the Essence for all the Edit Units is concatenated in order) it will be a valid VC-2 stream (or part thereof).

A.1.2 Unwrapping part of the wrapped VC-2 stream

If only a portion of the wrapped VC-2 stream is unwrapped then the result will not always be a valid VC-2 stream. At the least, either the Parse Info header at the start of the stream could have an incorrect value for the `previous_parse_offset` property (and so is not a valid Parse Info header for the beginning of a VC-2 sequence) or the end of sequence Parse Info header could be missing.

A.1.3 Editing

VC-2 is intra-frame coded. This makes it possible to reorder the video frames of the wrapped VC-2 stream by simply reordering the Edit Units. However, this will not necessarily result in a valid VC-2 stream. Notably the Parse Info header parse offsets and the sequence of `picture_number` values could become invalid. Additionally, the rule that all Sequence Headers in a sequence must be byte-for-byte identical could be broken if editing is performed on a wrapped VC-2 stream that consists of more than one VC-2 sequence where each VC-2 sequence consists of more than one video frame.

A.2 Relevant VC-2 Sequence Details

A.2.1 End of Sequence Parse Info header

The wrapped VC-2 stream could not conclude with an end of sequence Parse Info header. An essential part of generating / restoring an end of sequence Parse Info header (where this is necessary) is to calculate and write the correct value for the `previous_parse_offset` (which is dependent on the particular VC-2 sequence).

A.2.2 Parse Info header parse offsets

Each Parse Info header includes byte offsets to the previous and next Parse Info headers in the VC-2 sequence. For some types of VC-2 stream, editing operations at the Edit Unit level will cause the parse offsets to be incorrect; for other types of VC-2 stream this issue will not arise.

Special attention needs to be given to the parse offsets in both the Parse Info header at the start and at the end of a VC-2 sequence.

A.2.3 `picture_number` sequence

The first 4 bytes of the Picture header of each VC-2 Picture data unit is a `picture_number`. This value will increment by one for each successive Picture data unit in a VC-2 sequence. If the Edit Units of a wrapped VC-2 stream are reordered, then the `picture_number` of some or all Picture data units could need to be rewritten.

Annex B CDCI Picture Essence Descriptor Mapping (Normative)

The elements in the CDCI Picture Essence Descriptor shall be set as defined in Table B.1. This table shall be interpreted as follows:

- If an element is absent from the table then this annex imposes no constraints on either the presence or value of the element.
- If an element is present in the table then it shall be constrained as per the “Constraints” column.
    - If no constraint on the presence of the element is listed then this annex imposes no constraint on the presence of the element.
    - If no constraint on the value of the element is listed then this annex imposes no constraint on the value of the element.

Values derived from VC-2 properties shall be converted to the type required for use in the CDCI Picture Essence Descriptor.

Note 1: Section 7.3 explains potential complexities involved in determining Essence Descriptor values in certain situations.

Note 2: The variable length coding used in VC-2 means that some integer VC-2 properties can be of arbitrary size.

Note 3: It is possible for the VC-2 properties to exceed the size of the type to which they are mapped in the CDCI Picture Essence Descriptor. This is considered extremely unlikely since it would involve, for example, frame sizes with dimensions in excess of 4 billion pixels or sample depth greater than 32 bits. Implementers are advised to guard against the potential overflow of the types used in the CDCI Picture Essence Descriptor.

Note 4: The Transfer Characteristic value for `video_parameters[transfer_function] = 0` refers to Recommendation ITU-R BT.2020. However, this value is also applicable for a transfer characteristic in accordance with Recommendation ITU-R BT.709 because the parameters are equivalent.

Table B.1 – Constraints on Elements of the CDCI Picture Essence Descriptor

Element Name	Constraints
SubDescriptor UIDs	Element shall be present This array shall include exactly one Strong Reference to an instance of the VC-2 Sub-Descriptor
Sample Rate	{ `video_parameters[frame_rate_numer]` , `video_parameters[frame_rate_denom]` }
Essence Container	See Section 9
Codec	Element shall not be present

<table>
<tr><th>Element Name</th><th>Constraints</th></tr>
<tr><td>Signal Standard</td><td>[No additional constraint]

Note: Determining a value for this is complicated. It will not necessarily be directly related to <code>base_video_format</code>. If the Signal Standard element is present, implementations must take care to ensure that the value matches the related properties of the VC-2 stream.</td></tr>
<tr><td>Frame Layout</td><td>
<table>
<tr><th><code>video_parameters[source_sampling]</code></th><th><code>picture_coding_mode</code></th><th>Frame Layout value</th></tr>
<tr><td>0</td><td>0</td><td>FULL_FRAME</td></tr>
<tr><td>1</td><td>0</td><td>MIXED_FIELDS</td></tr>
<tr><td>0</td><td>1</td><td>SEGMENTED_FRAME</td></tr>
<tr><td>1</td><td>1</td><td>SEPARATE_FIELDS</td></tr>
</table>
</td></tr>
<tr><td>Stored Width</td><td><code>video_parameters[frame_width]</code></td></tr>
<tr><td>Stored Height</td><td>
<table>
<tr><th><code>picture_coding_mode</code></th><th>Stored Height value</th></tr>
<tr><td>0</td><td><code>video_parameters[frame_height]</code></td></tr>
<tr><td>1</td><td><code>video_parameters[frame_height] / 2</code></td></tr>
</table>

Note: This contradicts the table in SMPTE ST 377-1 G.2.7, which states that when Frame Layout has a value of SEGMENTED_FRAME then Stored Height is set to the frame height. In this case the value in the table in G.2.7 has been ignored because it contradicts the definition of Stored Height given in SMPTE ST 377-1 itself.
</td></tr>
<tr><td>Stored F2Offset</td><td>Element shall not be present</td></tr>
<tr><td>Sampled Width</td><td>As for Stored Width</td></tr>
<tr><td>Sampled Height</td><td>As for Stored Height</td></tr>
</table>

Element Name	Constraints
SampledXOffset	0 Note: Stored Width and Sampled Width are the same for VC-2 pictures.
SampledYOffset	0 Note: Stored Height and Sampled Height are the same for VC-2 pictures.
DisplayHeight	[No additional constraint] Note: It is recommended that this element be present and its value be the same as for Stored Height. However, there could be application specific reasons to set this element to an alternative value (for example it could be used to reflect the "clean area" communicated in the VC-2 Sequence Header properties).
DisplayWidth	[No additional constraint] Note: It is recommended that this element be present and its value be the same as for Stored Width. However, there could be application specific reasons to set this element to an alternative value (for example it could be used to reflect the "clean area" communicated in the VC-2 Sequence Header properties).

Aspect Ratio	[No additional constraint]
	
	Note 1: It is recommended that the value of this element be the simplified (reduced) fraction equal in value to `(A / B)` where:
	
	`A = video_parameters[pixel_aspect_ratio_numer] * video_parameters[clean_width]`
	
	`B = video_parameters[pixel_aspect_ratio_denom] * video_parameters[clean_height]`
	
	Note 2: Following this recommendation for Aspect Ratio as well as the recommendations for DisplayHeight and DisplayWidth will result in pixel aspect ratio that is different to that signaled in the VC-2 properties for some video formats (depending on the exact interpretation of these elements). There could be application specific reasons to set Aspect Ratio to an alternative value (for example to address this discrepancy).

<table>
<tr><th>Element Name</th><th>Constraints</th></tr>
<tr><td>Video Line Map</td><td>[No additional constraint]

Note: It is recommended that the value of this element be set correctly for the video interface (such as SDI) typically used to carry the relevant video signal standard, if it is possible to determine the video signal standard from the VC-2 properties (one such method would involve comparing the VC-2 properties to the values listed in SMPTE ST 2042-1 Table C.1). If this is not practical then it is recommended that:

• if Frame Layout is FULL_FRAME then set Video Line Map to {1, 0}
• otherwise set Video Line Map to {1, A} where:

<code>A = (video_parameters[frame_height] / 2) + 1</code></td></tr>
<tr><td>Alpha Transparency</td><td>Element shall not be present

Note: VC-2 cannot encode an alpha channel.</td></tr>
<tr><td>Transfer Characteristic</td><td>
<table>
<tr><th><code>video_parameters[transfer_function]</code></th><th>Transfer Characteristic value</th></tr>
<tr><td>0</td><td>06.0E.2B.34.04.01.01.0E.04.01.01.01.01.09.00.00</td></tr>
<tr><td>1</td><td>06.0E.2B.34.04.01.01.06.04.01.01.01.01.05.00.00</td></tr>
<tr><td>2</td><td>06.0E.2B.34.04.01.01.06.04.01.01.01.01.06.00.00</td></tr>
<tr><td>3</td><td>06.0E.2B.34.04.01.01.08.04.01.01.01.01.07.00.00</td></tr>
<tr><td>4</td><td>06.0E.2B.34.04.01.01.0D.04.01.01.01.01.0A.00.00</td></tr>
<tr><td>5</td><td>06.0E.2B.34.04.01.01.0D.04.01.01.01.01.0B.00.00</td></tr>
</table>
</td></tr>
<tr><td>Image Alignment Offset</td><td>Element shall not be present</td></tr>
<tr><td>Image Start Offset</td><td>Element shall not be present</td></tr>
<tr><td>Image End Offset</td><td>Element shall not be present</td></tr>
</table>

<table>
<tr><th>Element Name</th><th>Constraints</th></tr>
<tr><td>FieldDominance</td><td>
<table>
<tr><th>video_parameters[top_field_first]</th><th>FieldDominance value</th></tr>
<tr><td>True</td><td>1</td></tr>
<tr><td>False</td><td>2</td></tr>
</table>
Note: FieldDominance has no meaning for certain values of Frame Layout and so the inclusion of this element is discouraged in certain circumstances: see SMPTE ST 377-1 for full details.
</td></tr>
<tr><td>Picture Essence Coding</td><td>See Section 10</td></tr>
<tr><td>Coding Equations</td><td>
<table>
<tr><th>video_parameters[color_matrix]</th><th>Coding Equations value</th></tr>
<tr><td>0</td><td>06.0E.2B.34.04.01.01.01.04.01.01.01.02.02.00.00</td></tr>
<tr><td>1</td><td>06.0E.2B.34.04.01.01.01.04.01.01.01.02.01.00.00</td></tr>
<tr><td>2</td><td>06.0E.2B.34.04.01.01.0D.04.01.01.01.02.04.00.00</td></tr>
<tr><td>3</td><td>06.0E.2B.34.04.01.01.0D.04.01.01.01.02.05.00.00</td></tr>
<tr><td>4</td><td>06.0E.2B.34.04.01.01.0D.04.01.01.01.02.06.00.00</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Element Name</th><th>Constraints</th></tr>
<tr><td>Color Primaries</td><td>
<table>
<tr><th>video_parameters[color_primaries]</th><th>Color Primaries value</th></tr>
<tr><td>0</td><td>06.0E.2B.34.04.01.01.06.04.01.01.01.03.03.00.00</td></tr>
<tr><td>1</td><td>06.0E.2B.34.04.01.01.06.04.01.01.01.03.01.00.00</td></tr>
<tr><td>2</td><td>06.0E.2B.34.04.01.01.06.04.01.01.01.03.02.00.00</td></tr>
<tr><td>3</td><td>06.0E.2B.34.04.01.01.0D.04.01.01.01.03.05.00.00</td></tr>
<tr><td>4</td><td>06.0E.2B.34.04.01.01.0D.04.01.01.01.03.04.00.00</td></tr>
</table>
</td></tr>
<tr><td>Component Depth</td><td>state[luma_depth]

Note: For some VC-2 streams state[luma_depth] and state[color_difference_depth] could differ. However, Component Depth always indicates the luma component depth in this mapping.</td></tr>
<tr><td>Horizontal Subsampling</td><td>
<table>
<tr><th>video_parameters[color_diff_format_index]</th><th>Horizontal Subsampling value</th></tr>
<tr><td>0</td><td>1</td></tr>
<tr><td>1</td><td>2</td></tr>
<tr><td>2</td><td>2</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Element Name</th><th>Constraints</th></tr>
<tr><td>Vertical Subsampling</td><td>
<table>
<tr><th>video_parameters[color_diff_format_index]</th><th>Vertical Subsampling presence</th><th>Vertical Subsampling value</th></tr>
<tr><td>0</td><td>Element need not be present</td><td>1</td></tr>
<tr><td>1</td><td>Element need not be present</td><td>1</td></tr>
<tr><td>2</td><td>Element shall be present</td><td>2</td></tr>
</table>

Note: Vertical Subsampling has a default value of 1. If this element is omitted from the Essence Descriptor then a decoder is supposed to assume this default value for this element.
</td></tr>
<tr><td>Color Siting</td><td>[No additional constraint]

Note: This is not constrained by VC-2. The color difference planes are compressed independently of the luma plane.</td></tr>
<tr><td>ReversedByteOrder</td><td>Element shall not be present</td></tr>
<tr><td>PaddingBits</td><td>Element shall not be present</td></tr>
<tr><td>Alpha Sample Depth</td><td>Element shall not be present

Note: VC-2 cannot encode an alpha channel.</td></tr>
<tr><td>Black Ref Level</td><td>video_parameters[luma_offset]

Note: If the Black Ref Level element is omitted a decoder can determine the value to be used as specified in SMPTE ST 377-1 G.2.32 which relies on the Signal Standard element. However, this standard does not define a method for setting the Signal Standard element.</td></tr>
<tr><td>White Ref Level</td><td>video_parameters[luma_offset] + video_parameters[luma_excursion]

Note: If the White Ref Level element is omitted a decoder can determine the value to be used as specified in SMPTE ST 377-1 G.2.33 which relies on the Signal Standard element. However, this standard does not define a method for setting the Signal Standard element.</td></tr>
</table>

Element Name	Constraints
Color Range	`video_parameters[color_diff_excursion] + 1` Note: If the Color Range element is omitted a decoder can determine the value to be used as specified in SMPTE ST 377-1 G.2.34 which relies on the Signal Standard element. However, this standard does not define a method for setting the Signal Standard element.

Bibliography (Informative)

SMPTE ST 2042-2:2017, VC-2 Level Definitions

SMPTE RP 2047-1:2009, VC-2 Mezzanine Level Compression of 1080P High Definition Video Sources

SMPTE RP 2047-3:2016, VC-2 Level 65 Compression of High Definition Video Sources for Use with a Standard Definition Infrastructure