
AS5 SUBTITLE FORMAT

By Rodrigo Braz Monteiro, Niels Martin Hansen and David Lamparter
This work is licensed under a Creative Commons Attribution-Share Alike 3.0 License.



July 10, 2007

Contents

1	Abstract	2
2	AS5 Files	3
2.1	File Format	3
2.2	File Structure	3
2.2.1	[AS5]	4
2.2.2	[Events]	5
2.2.3	[Styles]	6
2.2.4	[Resources]	7
3	Style Overrides	8
3.1	General Information on Override Tags	8
3.2	Sub Station Alpha Tags	9
3.3	Advanced Sub Station Alpha Tags	9
3.4	AS5 Property Tags	9
3.4.1	\left	9
3.4.2	\right	9
3.4.3	\top	9
3.4.4	\bottom	9
3.4.5	\bordstyle	9
3.4.6	\effect	9
3.4.7	\relative	9
3.4.8	\vertical	10
3.5	AS5 Distortion Tags	10
3.5.1	\distort	10
3.5.2	\baseline	10
3.5.3	\bls	10
3.5.4	\fsc	11
3.5.5	\fay	11
3.5.6	\fax	11
3.6	AS5 Rastering Tags	11
3.6.1	\#vc	11
3.6.2	\#blend	11
3.6.3	\iclip	11
3.6.4	\blur	11
3.7	AS5 Advanced Tags	11
4	Renderer Behaviour Specification	12
5	Container Multiplexing Specification	13
5.1	Matroska	13
	References	14

1 Abstract

This document specifies the *AS5 Subtitle Format*, developed jointly by the Aegisub[1] and asa[2] teams in order to replace the old *Sub Station Alpha*[3] subtitle format and its extensions:

- Advanced Sub Station Alpha (ASS) implemented by Gabest in VSFilter[5]
- Advanced Sub Station Alpha 2 (ASS2), also implemented by Gabest in VSFilter
- Advanced Sub Station Alpha 3 (ASS3) implemented by equinox in asa.

The goal is to create a flexible, easy to understand and powerful subtitle format that can be used in hardsubs or multiplexed into Matroska Video[7] files as softsubs.

AS5 has no official meaning. The “A” can stand for Aegisub, asa, ASS or Advanced, the “S” for Subtitles, and the 5 is a reference to the fact that it’s a major improvement over SSA4 format (from which ASS, ASS2 and ASS3 derive). The full name of the format is “AS5 Subtitle Format”.

2 AS5 Files

2.1 File Format

All AS5 files are *REQUIRED* to comply with the three requirements below:

- Be encoded with one of *UTF-8*[8], *UTF-16 Big Endian* [9] or *UTF-16 Little Endian Unicode Transformation Formats*. UTF-8 is preferred.
- Not to have any character below Unicode code point U+20, except for U+09, U+0A, U+0D. That is, it must be a plain-text file.
- All lines must end with Windows line endings, that is, U+0D followed by U+0A.

These requirements are important so the AS5 format can be edited in most plain-text editors across most operating systems and languages without problems. The character set of a subtitle file can be autodetermined by its Byte-Order Mark or by the value of the first two bytes. See below.

When used as a standalone file, the extension should be `.AS5`. When multiplexed into a Matroska container, the Codec ID should be `S_TEXT/AS5`.

TODO: Get clearance from the Matroska team to use that Codec ID.

2.2 File Structure

The file is divided in *sections*, which are uniquely identified by a string inside square brackets, in a line of its own. From that point on, every next line is considered to be part of the last found section until another section is found. There is no end-of-section termination mark; they always end at the start of the next one or at the end of the file. *Section names are case sensitive*.

Each section is divided in lines, each line representing one command or definition. Empty lines (that is, lines only containing a line ending) *MUST* be ignored by the parser. It is recommended that programs generating AS5 files insert a blank line at the end of each section to increase readability. There *MUST* always be a blank line at the end of the file (as every line is required to end in a line break).

Each line in a section takes the general form of *Type: data1,data2,...,dataN*. An unknown *Type* *MUST* be ignored by a parser. Subtitle editing programs *SHOULD* keep such ignored lines in the file after re-saving it. Note that the space after the colon is *mandatory*.

There are two sections which are required, *[AS5]* and *[Events]*, the former being the equivalent of *[Script Info]* in previous formats. If either of those sections is missing, the file is invalid and (*MUST*) be refused by the parser. Any other section can be omitted from the file, and need not be implemented by all parsers. However, any unknown section *MUST* be preserved in the file by a subtitle editing program when it re-saves a file with sections that it does not recognize. It can, however, be removed at the user's discretion.

Finally, there is a special type of undefined group, *[Private:PROGNAME]*, which *MUST* be *ENTIRELY* preserved by other programs when re-saving it. This is used to store program-specific data. For example, Aegisub would create a group called *[Private:Aegisub]* to store its data inside. This type of group should be identified by the fact that it starts with "*[Private:.*".

Note that *Format:* lines from the previous formats are not admitted in AS5. If the parser finds any of them, it MUST halt parsing.

The sections MAY be written in any order, with the exception of the *[AS5]* section which MUST always be the first section.

2.2.1 [AS5]

This MUST be the first section in every AS5 file. If the very first line of the file is not *[AS5]*, the file MUST be rejected by the parser as invalid. Note, however, that the first line is allowed to contain a Byte-Order Mark (BOM), which is the character U+FEFF encoded in the encoding used for the rest of the script[10]. The first four bytes will therefore be:

- 0xEF 0xBB 0xBF 0x5B - UTF-8 (with BOM)
- 0x5B 0x41 0x53 0x53 - UTF-8 (without BOM)
- 0xFF 0xFE 0x5B 0x00 - UTF-16 LE (with BOM)
- 0x5B 0x00 0x41 0x00 - UTF-16 LE (without BOM)
- 0xFE 0xFF 0x00 0x5B - UTF-16 BE (with BOM)
- 0x00 0x5B 0x00 0x41 - UTF-16 BE (without BOM)

It is possible, therefore, to determine the encoding of the file by checking its first two bytes.

This section is used to declare several script properties that affect its parsing and rendering. All properties are stored in the format *Name: data*, with one property per line. This section MUST always declare the following properties:

- *ScriptType*: Should always be set to *AS5*, for this particular version of the specification. If this contains a value that the parser does not understand, it MUST abort parsing.
- *Resolution*: Should contain the script resolution in *WxH* format. For example, for a 640x480 script, this should say “*Resolution: 640x480*”. Note that this does not need to correspond to the video resolution, however, subtitles MUST be rendered on such a coordinate space. That is, in a 640x480 script, `\pos(320,240)` always represents the center of the script, no matter the resolution of the video it’s being drawn on. Also, in a 100x100 script, a radius 50 circle centered on the center will always take half of the height and half of the width of the video, even if that means being distorted if drawn on a video with a non-1:1 aspect ratio (for example, a 640x480 video).

The following items MAY also be used; they are not required, but are recommended. They all have default values:

- *Generator*: The name of the program that generated this script, e.g. “*Generator: Aegisub*”. Default value is empty. This should be ignored by the renderer, but might be useful for inter-editing-program interaction.

- **Wrapping:** The line wrapping style. This can be “Manual”, in which case only `\n` can break lines or “Automatic”, in which the renderer chooses how to break them. The default is “Automatic”. Note that if this is set to manual, the line can NEVER be broken at anywhere other than forced line breaks, even if it means that the line will become unreadable because it goes outside the display area.
- **Extensions:** A comma-separated list of all extensions being used in this file. At the moment, there are no extensions available. Renderers should read this to enable any extensions that they might support. Editing programs MUST keep this field intact, unless the user chooses otherwise. Scripts WILL break if the list of extensions is suddenly lost.
- **Credits:** Credits for the people who worked on this subtitle file. Purely for informational purposes and SHOULD be ignored by the renderer. Subtitling programs SHOULD be able to display these credits to the user.
- **Title:** The title of this script. Purely for informational purposes and SHOULD be ignored by the renderer. Subtitling programs SHOULD be able to display this title to the user.

Unlike in the previous incarnations of the format, storing private data here is not allowed, which means that this section MUSTnot contain any properties not listed here. Any application-specific or otherwise private data MUST be stored in `[Private:PROGNAME]` groups instead, as mentioned above.

2.2.2 [Events]

The most important section, [Events], lists all the actual subtitle lines in the file. The syntax has been radically simplified from previous incarnations of the format, and now consist of only five fields. Each line is represented as:

Line: `start,end,style,user,content`

Where:

- **Start:** The start time of the line. See below for the timestamp format. A line is only displayed if the timestamp of the current frame is *greater than or equal* to the start time. That is, start time is *inclusive*.
- **End:** The end time of the line. It follows the same format as the start time. The line is only displayed if the timestamp of the current frame is *lesser than* the end time. That is, end time is *exclusive*. In particular, it means that a line whose start time is equal to its end time will never be displayed. If the end time is earlier than the start time, the renderer MAY issue a warning, but it SHOULD render the remaining lines regardless of the issue.
- **Style:** The name of the default style used for this line. See the [Style] section below. If left blank, the script’s global default style MUST be used. If an unknown style name is specified, the renderer MUST fallback to default, and MAY issue a warning.
- **User:** This field is used by the program to store program-specific data in each line. Renderers SHOULD ignore this (but MAY use it for application-specific extension features). This field SHOULD be left blank if it’s not used. Note that whatever data is stored here MUSTnot contain any commas!
- **Content:** The actual text of the line. This contains actual text and override tags. See the section on override tags for more information.

The timestamp format is h...h:mm:ss[.s...], that is, it begins with an integer of arbitrary length (up to a maximum of 4 digits) representing the number of hours, followed by a one-digit or two-digit integer representing minutes, and a floating point number representing seconds. Leading zeroes in the hours field MAY be omitted. Localization is irrelevant: a period (“.”) is always used to separate the decimal point. This way, 0:21:42.5 and 0000:21:42.5000 are equivalent, and both represent 0 hours, 21 minutes, 42 seconds and 500 milliseconds.

Spaces between each field MUST be ignored by all parsers. Any spaces at the beginning of the content line SHOULD be stripped by any editing program. A hard space (see the overrides section) or empty override block should be used if space at the start of a line is truly desirable. That is, the two following lines are syntactically identical:

```
Line: 0:2:31.57 , 0:02:34.22 , , , Hello world of {\b1}AS5{\b0}!  
Line: 0:02:31.570,00:02:34.22,,,Hello world of {\b1}AS5{\b0}!
```

2.2.3 [Styles]

This is equivalent to the [V4 Styles] (and subsequent variations) from the Sub Station Alpha format. Like [Events], it has been greatly simplified when compared to the previous formats, and now each entry contains only three fields. They are declared as:

```
Style: name,parent,overrides
```

Where:

- Name: The name of this style. Style names are not case-sensitive, but MUST be unique. A script with conflicting style names MUST be rejected by the parser. If the style name is “Default”, it will be used for all lines that omit the style name. If there is no “Default” line, the renderer default is used.
- Parent: The style from which the current style derives from. See below for more information. Leaving this field blank means that the style derives from the renderer’s default style.
- Overrides: A list of override tags to define this style. See below.

Styles work in a very different way from the way they did on previous formats (with the notable exception of ASS3, which actually implements this very same style based on this format, as “StyleEx”). Instead of setting multiple parameters across many commas, you simply specify override tags. When a line uses a style, it’s as if the overrides of the style were inserted right before the start of the line contents.

Also, a style can inherit from another style, and define new overrides which are then appended to those of the parent style. The parent style MUST have been declared BEFORE the style trying to use it as a parent. If the parent doesn’t exist or wasn’t declared yet, the parser must refuse to parse the script. This is important because otherwise you could get a “inheritance loop”, where styles derive from each other in a cycle.

For example, see the following [Styles] group:

```
[Styles]
```

```
Style: Default,,\fn(Arial)\fs20
Style: Speech,,\fn(Respublica)\fs24\bord2\shad2\4a#80\2c#000000
Style: Actor1,Speech,\1c#B9C5E3
Style: Actor2,Speech,\1c#FFB3CF
Style: UglinessItself,Default,\fn(Comic Sans MS)
```

In the above fragment, the first style defines the Default style that will be used on all lines that don't set any style and the second style defines a base speech style that will be used for all actors (note that it doesn't inherit from Default, even though Default overrode the renderer's default, that one is still used for style definitions).

The third and fourth styles are based on the second, and simply assign different colours to it. They will both have all properties of Speech, and only differ in primary colour. Finally, the last example shows how to derive from the overridden default. In this case, font size would be 20 points, regardless of renderer's default.

The two Actor styles could have been defined without a parent style as follows:

```
[Styles]
Style: Actor1,,\fn(Respublica)\fs24\bord2\shad2\4a#80\2c#000000\1c#B9C5E3
Style: Actor2,,\fn(Respublica)\fs24\bord2\shad2\4a#80\2c#000000\1c#FFB3CF
```

Since all that deriving a style from another does is append the new tags to the end of the previous, this way of declaring styles is identical to the one above, but is more verbose.

If no Default style is defined, the renderer **MUST** choose its own defaults to render the text with. These are entirely arbitrary and can be set to anything, but the renderer **SHOULD** let the user set his own defaults. A simple Sans-Serif font with white text and black borders is recommended.

2.2.4 [Resources]

The new *[Resources]* section can be used to store information on external file resources, such as images and fonts. The general syntax is:

```
Resource: type,name,path
```

Where:

- Type: Must be either "font" or "image". Any other types **MUST** be ignored by the parser.
- Name: An unique name identifying this resource. For fonts, it must correspond to the font name, e.g., "Verdana". For images, it's the name that the file will be referred as in the rest of the script. If there is already a resource with this same name, the parser **MUST** abort the parsing.
- Path: The location of the file relative to the subtitles. This **MUST** be a relative path for external .as5 files, or a container-specific string for AS5 multiplexed into a container. The relative path **MUST** use forward slashes and be case-sensitive, in order to avoid UNIX compatibility issues.

3 Style Overrides

3.1 General Information on Override Tags

As with previous formats, AS5 uses override tags to set the style for lines. Also, it uses those same tags to set style definitions themselves (see above). Although many tags were imported from *Advanced Sub Station Alpha*, do not assume that they behave exactly the same. Some had their behavior changed or properly defined. Also, AS5 defines many new tags in addition to the old ones.

All tags must be inserted between a pair of curly brackets (`{}`), except on style definitions. A pair can contain any number of override tags inside it. They should be listed one after the other, with no spaces or any other kind of separator between them. Tags then affect all text that follows it, unless re-overridden or reset by the `\r` tag. For example:

```
{\fn(Verdana)\fs26\c#FFA040>Welcome to {\b1}AS5{\b0}!
```

In the following example, the first override block affects the entire text, but only “AS5” is bolded.

Some tags might begin with a `#` in their names. This means that there are actually five variations of this specific tag, the tag with `#` replaced with a number from 1 to 4 (inclusive) or without it altogether - in that case, the tag is assumed to mean the 1 variation. Those numbers represent the four different colours available on any given line:

- 1 - Primary colour, used for the main face of the text.
- 2 - Secondary colour, used on karaoke. See the karaoke tags for more information.
- 3 - Border colour. This is the colour of the border that outlines the text. See the `\bord` tag for more information.
- 4 - Shadow colour. This is the colour of the shadow dropped by the text. See the `\shad` tag for more information.

So, for example, you would use `\1c` or `\c` to set the primary colour, or `\3c` to set the colour of the border. `\#c`, however, does not exist in itself.

When a tag requires a floating point parameter, the decimal part must be specified using a period (`.`); never a comma. When a tag requires a colour parameter, it is given in HTML hexadecimal code, which is `#` followed by a 6-digit hexadecimal string, where the first two digits represent the red component, the next two the green component, and the last two the blue component (`#RRGGBB`). Sub Station Alpha style (Visual Basic hexadecimal) is not supported - if a parser finds any colour in `&HBBGRR&` format, it **MUST** issue an error.

It is forbidden to write comments inside standard curly brackets. Any unknown tags **MUST** be ignored, and anything that doesn't begin with a backslash **MUST** be considered an error. For inline comments, you need to use a special variation, in which the first character inside the overrides block is an asterisk (`*`). Renderers **MUST** completely ignore any text inside such blocks. For example:

```
{\fn(Verdana)\fs26\c#FFA040>Welcome to {\b1}AS5{\b0}!{*It's a nifty format, isn't it?}
```

3.2 Sub Station Alpha Tags

TODO: Write me

3.3 Advanced Sub Station Alpha Tags

TODO: Write me

3.4 AS5 Property Tags

These tags replace the old style and dialogue settings that were rarely used and generally only made the file more verbose and harder to read.

3.4.1 \left

TODO: Write me

3.4.2 \right

TODO: Write me

3.4.3 \top

TODO: Write me

3.4.4 \bottom

TODO: Write me

3.4.5 \bordstyle

TODO: Write me

3.4.6 \effect

TODO: Write me. Is this really desirable?

3.4.7 \relative

TODO: Write me

3.4.8 `\vertical`

TODO: Write me

3.5 AS5 Distortion Tags

These are tags characterized by the fact that they distort the shape of the text itself. They were designed to enhance the flexibility of the format while dealing with unusually-shaped imagery.

3.5.1 `\distort`

Usage:

`\distort(x1,y1,x2,y2,x3,y3)`

Description: The distort tag allows you to apply an arbitrary distortion to the block that follows it. It takes three coordinate pairs that, along with the origin (at the current baseline position) specify a quadrilateral.

P_0 is the origin, $P_1 = (x1, y1)$ is the corner at the end of the baseline for the affected text, $P_2 = (x2, y2)$ is the point above that, and $P_3 = (x3, y3)$ is the point above P_0 . That is, they are listed clockwise from origin (P_0).

This tag can be animated with `\t`.

Implementation: This tag cannot be reduced to an affine transformation, so it cannot be expressed in Matrix form. In order to transform a given (x,y) coordinate pair to it:

1. Normalize the (x,y) coordinates to a (u,v) system, so that $P_0 = (0,0)$ and $P_2 = (1,1)$. This can be done by dividing x by the block's baseline length (bl) and y by the block height (h). The matrix for this operation is:

$$\begin{bmatrix} \frac{1}{bl} & 0 \\ 0 & \frac{1}{h} \end{bmatrix}$$

2. Apply the following formula: $P = P_0 + (P_1 - P_0)u + (P_3 - P_0)v + (P_0 + P_2 - P_1 - P_3)uv$
This can be interpreted as simple vector operations, that is, apply that once using the x coordinates and another using the y coordinates. Since the four points are constant, the coefficients can be precalculated, resulting in a very fast transformation.

3.5.2 `\baseline`

TODO: Write me

3.5.3 `\bls`

TODO: Write me

3.5.4 \fsc

TODO: Write me

3.5.5 \fay

TODO: Write me

3.5.6 \fax

TODO: Write me

3.6 AS5 Rastering Tags

These tags affect how the subtitles are rasterized, that is, they affect things such as colour, blurring, etc.

3.6.1 \#vc

TODO: Write me

3.6.2 \#blend

TODO: Write me

3.6.3 \iclip

TODO: Write me

3.6.4 \blur

TODO: Write me

3.7 AS5 Advanced Tags

These are more advanced tags, which might prove to be fairly complex to implement. They include things such as ruby text support (also known as furigana, when used with Japanese Kanji).

TODO: Write me

4 Renderer Behaviour Specification

TODO: Write this section

5 Container Multiplexing Specification

5.1 Matroska

Storage of AS5 files in Matroska files is similar to how similar formats are stored.[11] The Codec ID used is S_TEXT/AS5

First, the entire file is converted to UTF-8 (if it isn't already UTF-8). Then, all sections other than *[Events]* and *[Resources]* are stored on the *CodecPrivate* element. For the *[Resources]* section, each line is parsed and files are converted to Matroska file attachments. **TODO:** Specify this more clearly.

Finally, each line in the *[Events]* section is read and stored each in a block. The *start* and *end* fields are parsed (see the specifications on the section describing *[Events]*) and set as the *TimeStamp* and *BlockDuration* elements. The line itself is then stored in the following format:

```
Line: readOrder,style,userData,contents
```

Where *readOrder* is the number that the line had on the file. This is necessary so the file can be demultiplexed back in its original order, since lines will be stored in chronological order while inside the Matroska file. The remaining fields should just be copied from the original line.

References

- [1] Rodrigo Braz Monteiro, Niels Martin Hansen, David Lamparter et al., Aegisub. Application, 2005-2007.
<http://www.aegisub.net/>
- [2] David Lamparter, asa. Application, 2004-2007.
<http://asa.diac24.net/>
- [3] Kotus, Sub Station Alpha. Website, 1997-2003.
http://web.archive.org/web/*/http://www.eswat.demon.co.uk/substation.html
- [4] #Anime-Fansubs, Advanced Sub Station Alpha.
<http://www.anime-fansubs.org>
<http://moodub.free.fr/video/ass-specs.doc>
- [5] Gabest, VSFilter. Application, 2003-2007.
<http://sourceforge.net/projects/guliverkli/>
- [6] David Lamparter, Advanced Sub Station Alpha 3. Website, 2007.
<http://asa.diac24.net/ass3.pdf>
- [7] The Matroska project. Website.
<http://www.matroska.org/>
- [8] The Internet Society, RFC 3629, “UTF-8, a transformation format of ISO 10646”. Website, 2003.
<http://tools.ietf.org/html/rfc3629>
- [9] The Internet Society, RFC 2781, “UTF-16, an encoding of ISO 10646”. Website, 2000.
<http://tools.ietf.org/html/rfc2781>
- [10] Unicode, Inc, The Unicode Standard, Chapter 13. PDF, 1991-2000.
<http://www.unicode.org/unicode/uni2book/ch13.pdf>
- [11] The Matroska project, specification for SSA/ASS subtitle formats. Website.
<http://www.matroska.org/technical/specs/subtitles/ssa.html>