Differences between the GWG 1v4 and 2015 specifications

Authors	David van Driessche Executive Director, Ghent Workgroup Chief Technical Officer, Four Pees davidvd@publigence.com
Date	27 January 2016
Status	v1

Table of Contents

1 Introduction

The GWG 1v4 specifications from the Ghent Workgroup have long been the de facto standard for exchange of PDF documents in graphic arts. Evolving from such a standard is always difficult, yet we feel it's necessary to keep up with changes in technology. The GWG 2015 specifications were born as the result of that evolution.

1.1 What about 2012?

The 2015 specifications aren't the first round of improvements done; in 2012 we already released significantly different specifications. Those were used a lot in testing and in some countries were used as the basis for national standards. In most cases, however, the specifications themselves and the technological and business landscape were not ready for widespread adoption.

While the 2012 specifications exist and will continue to exist, we don't expect them to be widely implemented.

1.2 Transition

Because of the fundamental differences between the 1v4 and 2015 specifications, we don't expect an instant transition. We expect that the 2015 specifications will be picked up as acceptable specifications and that only after a (long) transition period 1v4 will no longer be accepted. A number of national organizations already announced plans to that effect.

1.3 Clarity

Change is easier if it is clearly explained *what* changes, and *why* it changes. The following chapters do exactly that: document the changes between 1v4 and 2015 in clear and easily understandable language. Each chapter highlights one particular area that has changed and discusses the reasons behind the change and the repercussions that change may have in today's workflows.

This is not a technical document; it was not written to help developers implement the 2015 specifications (other documentation exists for that and we advise you to contact the Ghent Workgroup if you need more information).

1.4 Test

This is a major change! As explained above, we expect you to have a transition period to go from 1v4 to 2015. We also expect that you will have to test your workflows, applications and procedures to make sure they all support the new specifications; especially so if you allow RGB images.

The Ghent Workgroup can help with that aspect as well; the Ghent Output Suite was created specifically to test workflows and show which aspects may not be up to spec. It's free and you can download it from the Ghent Workgroup site at http://www.gwg.org/download/test-suites/ghent-output-suite/

2 Variants

Different workflows require different PDF documents; examples are easy to find:
- Advertisements are typically single-page PDF documents, books always multi-page
- Newspapers often accept the use of spot colors, magazines often don't
- Documents for sheet fed workflows can be multi-page, web workflows typically single-page

To cover such differences, the Ghent Workgroup has always released different *variants* for its specifications.

2.1 What changes?

The number of variants in the 2015 specifications is very different than in the 1v4 specifications.

2.2 Why?

There are two reasons for this:
- A number of variants in the 1v4 specifications were hardly used or converged until there were very few differences between them. Such specifications were merged in the 2015 specifications.
- The variants that remained after this process were effectively doubled to create one variant that supports color-managed RGB for images and one that doesn't.

2.3 Consequences…

If you were using a 1v4 variant that no longer exists, you will need to select a variant in 2015 that comes as close as possible to what now exists. The following table provides an overview of each 1v4 specification and what we believe is the best match in the 2015 specifications.

MagazineAds_1v4	GWG_MagazineAds_2015 CMYK
NewspaperAds_1v4	GWG_NewspaperAds_2015 CMYK
SheetCmyk_1v4	GWG_SheetCmyk_2015 CMYK
SheetCmykVeryHiRes_1v4	GWG_SheetCmyk_2015 CMYK
SheetSpotLoRes_1v4	GWG_SheetSpot_2015 CMYK
SheetSpotHiRes_1v4	GWG_SheetSpot_2015 CMYK
WebCmykNews_1v4	GWG_WebCmykNews_2015 CMYK
WebCmykHiRes_1v4	GWG_WebCmyk_2015 CMYK
WebSpotNews_1v4	GWG_WebSpot_2015 CMYK
WebSpotHiRes_1v4	GWG_WebSpot_2015 CMYK

If additionally, you want to allow color-managed RGB for images, you need to select the "CMYK + RGB" variants.

3   Reliance on PDF/X

All Ghent Workgroup specifications are built on ISO PDF/X standards. In other words, if a PDF document is compliant with a Ghent Workgroup specification, it is also necessarily compliant to the PDF/X standard that particular specification is built on. The Ghent Workgroup specifications impose *additional* requirements, but never *remove* requirements from PDF/X.

3.1   What changes?

The 1v4 specifications are built on PDF/X-1a (ISO PDF/X-1a:2001), the 2015 specifications are built on PDF/X-4 (ISO PDF/X-4:2010). This is perhaps the single biggest change between 1v4 and 2015; PDF/X-4 is a much more modern standard than PDF/X-1a. It allows a whole range of PDF features that previously were not allowed, but it also imposes additional restrictions; it tightens up some of the areas where PDF/X-1a was relatively lax.

The following is an incomplete list of changes in PDF/X-4:

• Transparency is allowed.
• Optional content (sometimes called "layers") is allowed.
• Color-managed color spaces are allowed.
• Stricter rules exist for overall PDF file structure, fonts and XMP metadata.

It is important to note that the 2015 specifications often impose additional constraints on these features allowed in PDF/X-4.

3.2   Why?

The ISO PDF/X-1a standard is a very mature and reliable standard; it's well understood and implemented by a wide range of solutions. It is however also an old standard that blocks many modern practices. To give a few examples:

• Transparency has to be flattened. Transparency flattening is a complex process that causes a number of real problems in modern workflows: editing (even of text) is often made completely impossible, color transformations often cause problems, repurposing of documents is made more difficult and – perhaps most obvious – transparency flattening sometimes leads to visual defects (white lines, color differences…).
• Optional content can contribute significantly to the power of some modern workflows (to handle regional versions, different language versions, deal with non-print content…) but it is not allowed in PDF/X-1a based PDF documents.
• Having (color-managed) RGB content for certain elements in a PDF document could significantly improve the visual quality of PDF documents in certain workflows (such as when displayed on mobile devices, or when printed with large gamut devices in digital printing or large format workflows). PDF/X-1a allows only CMYK and spot colors.

3.3   Consequences…

The change from PDF/X-1a to PDF/X-4 is an important one and can have substantial consequences for your workflows.

3.3.1   Transparency

Because transparency is now allowed, your workflows need to be able to deal with that.

• Preferably all steps in your workflow support transparency correctly, or
• Transparency flattening will have to be implemented at some point.

3.3.2 Optional content

Optional content allows PDF documents in which part of the content is hidden. Your workflows must be able to honor those settings. More explanation is given in chapter 6 ("Optional content").

3.3.3 Color

It's obvious that being able to handle color spaces other than CMYK, gray and spot color (as was allowed in 1v4) is a considerable challenge. The 2015 specifications make this task easier by implementing different variants of the specifications; some only allowing the "old" color spaces and some allowing additional color-managed color spaces for some objects. Those requirements and their implications are described in chapter 4 ("Color").

3.3.4 Stricter rules

Overall it's good to have stricter rules to make sure the PDF documents you receive are properly built. However, those stricter rules may occasionally lead to additional warning or error messages being shown. That may introduce an additional educational challenge in your organization.

4 Color

In most standards and specifications actively used today, the color spaces in a PDF document are restricted to gray, CMYK and spot color (named colors). This was the case for the 1v4 and the 2012 specifications from the Ghent Workgroup as well.

4.1 What changes?

In the 2015 specifications, we relaxed the color restrictions in *some* variants. For each market segment two variants are defined: one named "CMYK" and one named "CMYK + RGB".

Depending on the variant, the following restrictions now apply:

- **CMYK**:
  The restrictions are essentially the same as in the 1v4 specifications; colors are restricted to gray, CMYK and spot colors.
- **CMYK + RGB**: For most objects in PDF documents, the restrictions are essentially the same as in the 1v4 specifications. For real images (the type of images traditionally used to represent photos) the restrictions are relaxed: those images are also allowed to use ICC-based RGB, CalRGB, and Lab.

4.1.1 Color managed

Remark that all content in *all* variants is color managed. Compliant PDF documents have a CMYK output intent that defines the CMYK objects in the file. And included RGB images use only color managed RGB flavors (RGB with an ICC profile attached, CalRGB which is a well-defined RGB color space and Lab). So in all cases you know exactly what color you are dealing with, which gives color software and RIPs the necessary information to handle it correctly.

4.2 Why?

Introducing color-managed RGB in a workflow brings a number of challenges, especially when it's in files where there's also transparency used. The 2015 specifications try to strike a balance between allowing RGB in those cases where it's useful, and not allowing it in all other places.

So why images? Images (photos) is where you see the biggest problems today. If you prepare a photo for the wrong CMYK flavor and it has to be adjusted afterward, you don't have its full quality anymore. Even worse, if you prepare a photo for CMYK and it has to be used in an RGB workflow afterwards (published on the Internet or on a mobile device for example), it looks awfully dull compared to photos kept as RGB.

4.3 Consequences...

If you decide to only accept the CMYK variants, your challenges around color will be very similar to what you have today. The main difference is the possibility that files will still have live transparency which may make them more challenging to process, but that isn't really related to color.

If you decide to accept the CMYK + RGB variants, your potential gains are the biggest but there are definitely additional challenges. In that case, you have to either convert files to the CMYK flavor you need during production, or you need to make sure that your RIPs will be able to handle files with RGB images correctly.

5 File size

When we started developing the Ghent Workgroup specifications, file size (the size of a file in bytes) was much more important than it is today. Both in terms of storage and transmission time, technology has progressed significantly.

5.1 What changes?

In 1v4 we included warnings for documents that were not properly compressed (PDF documents support all kinds of advanced compression, but you're not obligated to use them while generating the PDF) and for images that had a resolution that was too high.

In 2015, we have removed all such warnings.

5.2 Why?

There are two reasons to remove those warnings (and to advise vendors to remove associated fixes):

• Having larger PDF documents is not experienced as so much of an issue anymore.
• The result of these warnings was that many vendors applied fixes to the PDF documents they processed; most notoriously image down sampling. These fixes often lead to a reduction in image quality and they interfered with workflows where security features were used.

5.3 Consequences...

This sounds like a relatively minor issue because the Ghent Workgroup specifications contained only warnings, not errors. But in many cases vendors use these specifications to decide which fixes they will apply on PDF documents by default as well.

As a result, PDF documents (especially those containing lots of high-resolution images) may be bigger.

6 Optional content

PDF documents can contain "optional content" – often referred to as "layers". Optional content is content that can easily be made visible or invisible either by the user or automatically in a workflow (visible when the document is viewed, for example, but invisible when it's printed).

6.1 What changes?

In 1v4 optional content is forbidden as it didn't yet exist in the PDF version 1v4 is built on. In 2015, we decided to allow optional content, with some restrictions for some variants:

• The variants destined for advertisements allow optional content but restrict its use so a single unique view on the document exists. This was done to make sending around advertisements less risky.
• All other variants allow any kind of optional content.

6.2 Why?

Though optional content is not useful in *all* workflows, there are a growing number of workflows where their use was desirable. Examples are publications with regional differences, publications containing different language versions, workflows where optional content is used to communicate non-printing information…

6.3 Consequences...

All processing software and workflows must be intelligent enough to handle optional content correctly. In most cases, this means that the software and hardware must know that *hidden* content should not be shown / printed.

7 OpenType fonts

Fonts come in many formats; one of the most modern font formats is called OpenType. It supports many advanced typographic features and has provisions to support the many writing systems used throughout the world.

7.1 What changes?

In the 2015 specifications OpenType is fully supported; in 1v4 OpenType was not supported as it didn't yet exist in the PDF version 1v4 is built on.

7.2 Why?

Not allowing OpenType fonts implies that, if you use OpenType fonts in your layout application, those have to be converted to Type 1 or TrueType fonts during export. Because OpenType fonts embedded in PDF documents have not been allowed for many years, this is exactly what all layout applications do today (and so likely what you do as well even though you might not realize it).

Keeping the real OpenType font would have advantages though (even if only because there is one less conversion step taking place), and because OpenType should be well supported by now, we decided to no longer restrict the use of OpenType fonts in the 2015 specifications.

7.3 Consequences...

As no layout applications we know of currently support embedding OpenType fonts in PDF documents, the impact of this change should – at least initially – be very small. The Ghent Output Suite (see the note on testing) contains a patch to test support for all font formats allowed in the 2015 specifications if you want to make sure you'll have no problems going forward.

8 16-bit images

Originally PDF only supported images with up to 8 bits per color channel; that was later extended to also support images using 16 bits per color channel. The higher number of bits means more accurate color can be expressed in such images.

8.1 What changes?

In the 2015 specifications 16-bit images are fully supported; in 1v4 16-bit images were not supported as they didn't yet exist in the PDF version 1v4 is built on.

8.2 Why?

We believe the technology is in place to support such images without any problem. It is therefore not useful to restrict them from being used in PDF documents.

8.3 Consequences…

If you work with relatively modern software or hardware, you shouldn't have any problems. The Ghent Output Suite (see the note on testing) contains a patch to test support for 16-bit images as allowed in the 2015 specifications if you want to make sure.

9 Total Ink Coverage

In many output processes using ink, there is an amount of ink at which things start to break down: because the ink no longer dries in time, because the paper gets too wet and breaks apart, or simply because the quality of the print starts to suffer. Because of this, workflows typically define a maximum *Total Ink Coverage (TAC)* that has to be respected.

9.1 What changes?

In some variants, small changes to the acceptable TAC limit have been introduced: specifically, both the GWG_SheetCmyk_2015 and GWG_SheetSpot_2015 now have a limit of 320% wherein the equivalent 1v4 specifications it was 340%.

The biggest change, however, is that the 1v4 specifications allowed a very casual and incomplete calculation of the TAC in PDF documents. In the 2015 specifications, it is demanded that the effective TAC is calculated; This is usually done by rendering all relevant objects into an image (respecting object properties such as transparency and overprint) and then examining that image.

9.2 Why?

The small changes in TAC limit came through years of experience with the 1v4 specifications and the belief that a slightly lower TAC in certain variants would more accurately reflect real production limits.

The fact that the 2015 specifications now demand true TAC calculations sounds like a no-brainer. It's important to realize that we did not demand this in the 1v4 specifications because it was deemed impractical with the technology available at the time, but that no longer is the case today.

9.3 Consequences...

If you test the same file against a 1v4 specification and a 2015 specification, the result may differ. Whether that represents a real production problem is something you will have to evaluate in your workflow. The results from software using the 2015 specifications will normally be more accurate.

10 Output Intent and ICC profiles

PDF documents can contain an *Output Intent* that specifies what the intended output is for that particular document; the output intent also defines the color space of objects defined in CMYK in the PDF document. In PDF/X-4, an additional requirement states that an ICC profile representing that intended output must be embedded in the output intent in the PDF document.

10.1 What changes?

The 2015 specifications are compliant with PDF/X-4 so PDF documents must contain an output intent and that output intent must contain a valid ICC profile. There are two important changes:

• First of all, the ICC profile v4 format is now fully supported.
• And secondly, there no longer is a list of *acceptable* ICC profiles for each variant.

10.2 Why?

As far as the ICC profile v4 format is concerned, we felt the technology is in place to support it and didn't see any reason to disallow it.

The list of acceptable ICC profiles is a more difficult issue. That list was created in order to give users guidelines on what was acceptable and what was not for a specific variant. In practice however, we found that the list didn't add a lot of value and that the ICC profiles changed too quickly for it to make sense in a standard that has to be supported over a long period of time.

10.3 Consequences…

Based on the tests performed by the Process Control Subcommittee of the Ghent Workgroup, the ICC profile v4 format is well-supported by relatively recent software and hardware; as a result we don't expect major problems with v4 profiles.

As far as the list of acceptable ICC profiles is concerned, you will have to (as you always had to) check that the ICC profile in the output intent of the documents that you receive is, in fact, a good (enough) ICC profile for your workflows.

11 Single image pages

One of the big advantages of PDF is that it doesn't require all page content to be rasterized (converted into an image); text and vectors can be maintained and can be output with maximum quality.

But PDF also allows converting all objects on a page into a single image; such a page – that contains nothing but a single image that everything has been rendered into – is called a single image page in the 2015 specifications.

11.1 What changes?

Single image pages where not mentioned in the 1v4 specifications; in the 2015 specifications they get their own limits for minimum image resolution.

11.2 Why?

If a page contains only one image and no other text or line art elements, that likely means text and vectors have been rendered into that image. This happens, for example, if the PDF document has been opened and saved in Adobe Photoshop.

In such cases, it is desirable that the resolution of the image is higher than one would typically need for photos in that variant.

11.3 Consequences…

In all variants the minimum resolution where an error is produced is the same as the minimum resolution for real images (photos) in that variant, so you're unlikely to get additional errors.

But because the minimum resolution where you get a warning is much higher than the equivalent minimum resolution for real images (photos) in that variant, you're much more likely to get warnings for PDF documents using single image pages.