



PDF CanOpener© by WindJack Solutions, Inc.

Version 2.0

PDF CanOpener is an Adobe® Acrobat® plug-in for analyzing and manipulating the internals of PDF Documents. It displays a PDF Document's tree of PDF COS objects in an Explorer style structure that is both intuitive and easy to use. PDF CanOpener also provides a selection tool for identifying drawn objects on the PDF page, displaying details about the object (fonts, color spaces, etc), and navigating to the object in the COS Object tree. PDF CanOpener not only displays this difficult to access information, it also gives the user total control to edit, add, delete, rename, copy, and paste these objects at any location within the COS Object Tree. Please see the Adobe® PDF Reference for information on the PDF file structure and object model.



CAUTION: Use with care. PDF CanOpener can be used to permanently alter or destroy information that may render a document un-viewable. This plug-in should only be used for analysis and development purposes and never to manipulate a working document. Backup all documents before use.

Requirements:

- Adobe® Acrobat® version 5.0, 6.0, or 7.0
- Microsoft® Windows® NT 4.0, XP, or 2000

Activation:

PDF CanOpener can be activated for any currently open document by one of three methods (See **Figure 1**).

1. Selecting the “PDF CanOpener” menu item under the “Advanced” menu in Acrobat 6.0 and 7.0, or the “Tools” menu in Acrobat 5.0.
2. Clicking on the PDF CanOpener icon  from the “Advanced Editing” toolbar.
3. Using the PDF CanOpener Selector Tool  on any page or page object in a document.

Once activated, it will remain so until toggled off by either selecting the menu item or clicking the toolbar button. PDF CanOpener can also be deactivated by various activities that may cause the connection between the display window and the document to be broken. These activities include moving the PDF CanOpener display window between documents, following a “link” to another document, and other similar actions. Activation is easily re-established by using one of the three methods listed above. In most cases the state of the previously displayed document tree is maintained.

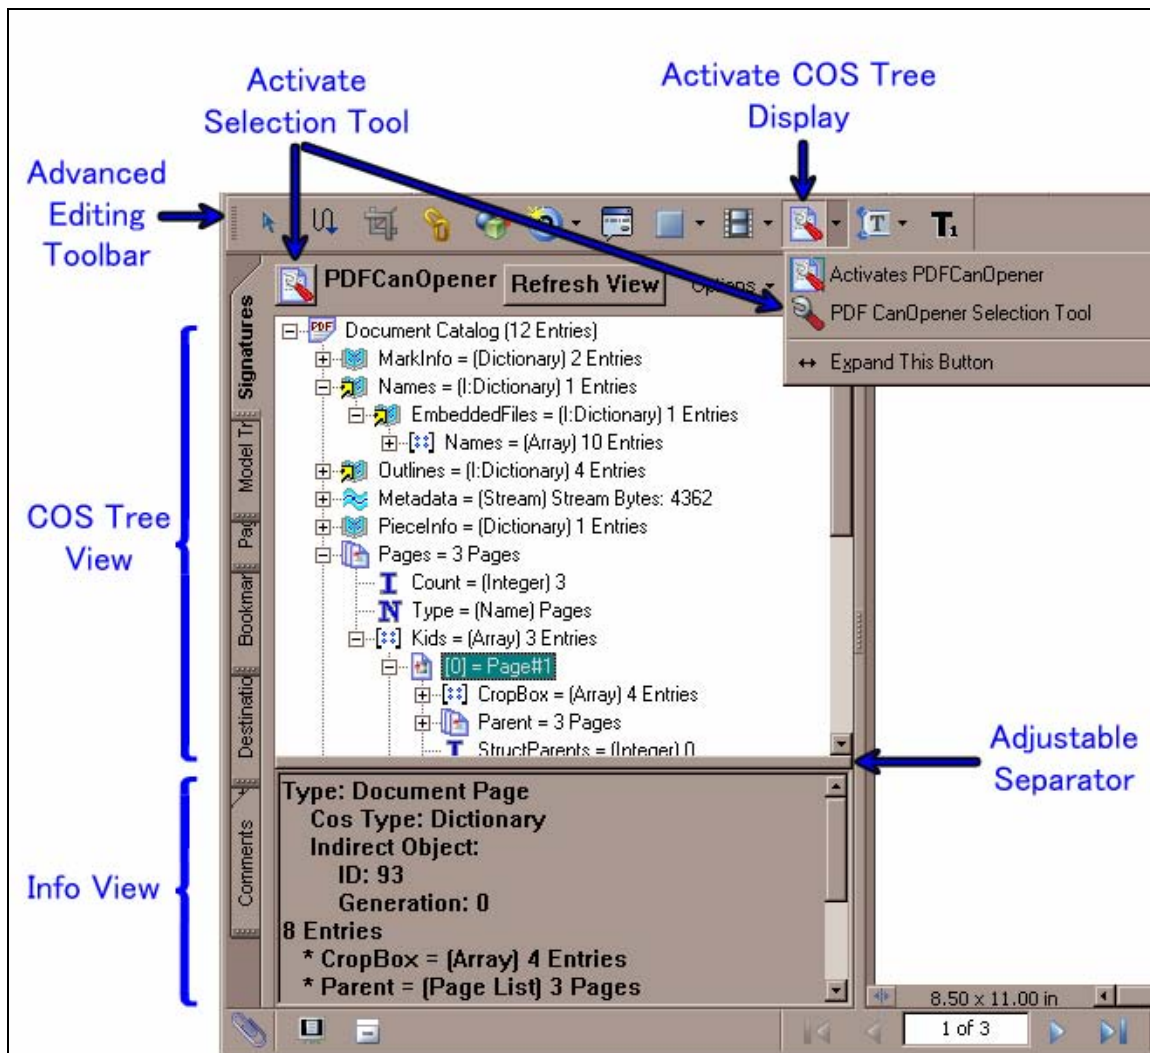


Figure 1 – Activating and viewing the PDF CanOpener Display

The Display:

Once activated, the PDF CanOpener display can appear on any one of the ordinary Acrobat navigation windows (see Preferences). Several of the most popular navigation windows are typically displayed to the left of a document as a set of tabbed windows. However, they may also appear as modeless floating windows. The Acrobat Navigation windows are a convenient location for drawing the PDF CanOpener display. The original information in these windows, as well as the functionality of these windows, are not affected in any way. Simply deactivating PDF CanOpener uncovers the original window.

The display window is subdivided into a *tree view* and an *info view*. These sub-windows can be resized relative to each other with the movable bar that separates them. The *tree view* displays the tree of COS objects for the current document. The root of this tree is always the document catalog. A right mouse click in this window will display a popup context menu of operations that can be applied to the COS objects in the tree. The *info view* displays information about the currently selected node. It can also display the rendered image of an XObject.

Copying Text from the *info view window*

It is often useful to select and copy text from the *info view window*. Unfortunately, Acrobat reserves use of the Ctrl-c operation for its own purposes, so PDF CanOpener provides two alternate methods for copying text from this window.






1. When text is selected and the keyboard focus is on the *info view window*, pressing the “c” key is equivalent to Ctrl-c.
2. A right click with the cursor inside the *info view window* will display a popup menu with a single “copy” item.

COS Object Symbols:

Each COS Object type defined by the PDF Specification is represented in the tree as an icon for easy identification. Additionally, some special objects are also given icons or more descriptive labels.

Standard Objects


Scalar Types:

-  - Boolean
-  - Integer
-  - Fixed
-  - Name
-  - String

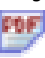



Container Types:

-  - Dictionary
-  - Array
-  - Stream

Object Modifier:

-  - Indicates an object is indirect. Can overlay any basic object type icons except Stream (Streams are indirect by definition).

Special Objects:

-  - Document Catalog (Root Object)
-  - Page List
-  - Page Object
-  - Annotation Object

Special Objects, without Icons:

- **XObjects** – Both Form and Image XObjects are labeled as such and given special treatment in some contexts.
- **Rectangles** – In some contexts, an array is known to contain values that describe a rectangle. In these cases the array elements are labeled appropriately with *left*, *top*, *right*, and *bottom*.

The objects that create this context are:

- **Annotations** – The “Rect” entry of the annotation dictionary.
- **Form XObject** – The “BBox” entry of the stream attributes dictionary.
- **Pages** – Several entries in the page dictionary.
 - MediaBox
 - CropBox
 - BleedBox
 - TrimBox
 - ArtBox
- **Matrix** – The “Matrix” entry in Form XObject attributes dictionary. The elements of this array are labeled to indicate the elements of a PDF Matrix Object.
- **Annotation Flags** – The individual flags are listed in the *info view* when this object is selected on the *tree view*. This object has a special editor with check boxes for each individual flag.
- **Page Content Streams** – Labeled as “Contents”. Given special treatment in the Stream editor as a viewable object.
- **Page Content Arrays** – Labeled as “Contents”. Given special treatment in the Stream editor as a viewable, but not editable, object. Special entry provided in the COS Tree Context Menu for display in the Stream Editor.

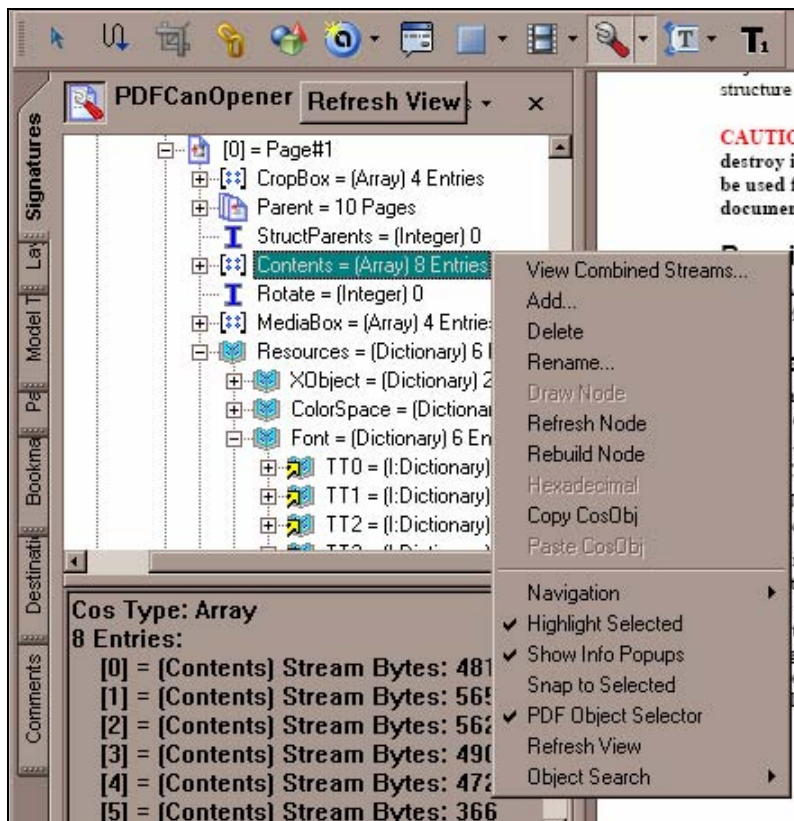


Figure 2 – COS Object Context Menu.

The COS Tree Context Menu:

Right clicking anywhere in the *tree view* will display a popup context menu. Items on the menu are enabled depending on the context of the right click. If the cursor is over a tree node, the menu items that can be applied to that COS object are enabled. If the cursor is not over any tree nodes, then only the globally applicable items are enabled. The first menu item is usually either “Edit” or “Add.” “Add” for Dictionary and Array objects, and “Edit” for all others.

Menu Items:

Object Operations:

- **View Combined Streams**
This option is only visible if the current node is an Array of Page Content Streams. It displays combined Content for the current node in the Stream Editor with all editing options turned off. Pressing the **Shift** key and double clicking on an Array of Page Content Streams has the same effect as selecting this option.
- **Edit/Add:**
An appropriate dialog is displayed for editing or adding items to the specified COS object.
- **Delete:**
Removes the specified COS Object from the document.

- **Rename:**
Renames the object by first removing it from the dictionary object that contains it and then adding it back in with the new name. Direct objects are essentially destroyed and rebuilt.
note: This operation is only valid for objects contained within a COS Dictionary.
- **Draw Node:**
Draws the graphic representation of the selected COS object in the *COS Tree*.
note: This operation is only valid for drawable objects: Form and Image XObjects, Content Streams, Patterns, Thumbnails, and Type3 Char Procs.
- **Refresh Node:**
The specified node and all expanded descendants are re-examined. Child nodes are added to, or deleted from the *tree view* to match the current object structure of the document.
- **Rebuild Node:**
The specified node is deleted and rebuilt. All expanded descendants are closed.
- **Hexadecimal:**
Toggles the display of the COS object value to a hexadecimal format.
note: This operation is only valid for Integers, Strings, and Streams.
- **Copy CosObj:**
Copies the specified COS object to the PDF CanOpener clipboard.
- **Paste CosObj:**
Pastes the COS object previously copied to the PDF CanOpener clipboard into the specified COS container object, either a *Dictionary* or an *Array*.
- **Navigation – Tree→ Doc:**
The document view is scrolled to the page containing the specified object.
note: This item is only valid for Page and Annotation Objects.

Global Operations:

- **Navigation – Doc→ Tree:**
Cos Tree display is scrolled to the current document Page Object.
- **Highlight Selected:**
Enables/Disables highlighting of PDF Objects in the document view for the selected tree node. Very useful for determining which tree node is associated with which PDF Object.
note: This item is only valid for Annotation Objects.
- **Show Info Popups:**
Enables/Disables the info popups for objects identified by the Object Selector.
- **Snap to Selected:**
Enables/Disables automatic navigation from the tree to the document for the selected tree node.
note: This item is only valid for Page and Annotation Objects.
- **PDF Object Selector:**
Activates/Deactivates the PDF Object Selector Tool.
- **Refresh View:**
Forces the current page to refresh its' contents. May not affect Annotations.
- **Object Search→ Find:**
Displays the Object Search Dialog.
- **Object Search → Show Results:**
Displays the Results Dialog from the previous Object Search.

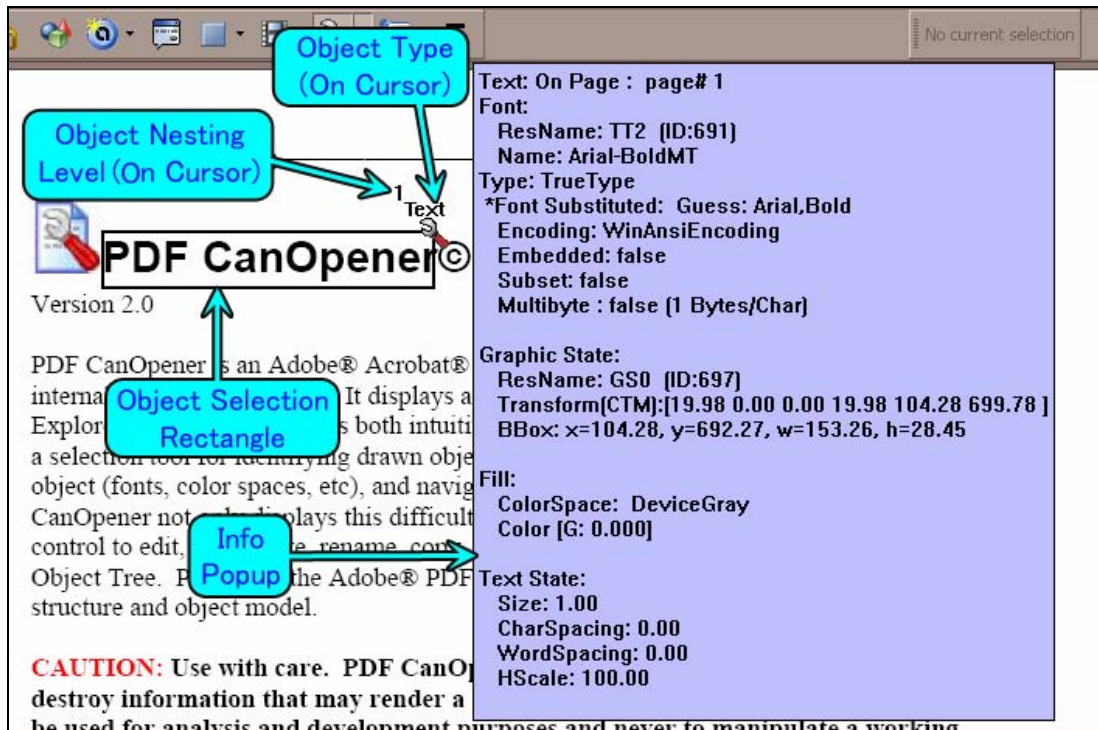





Figure 3 – Parts of the PDF CanOpener Object Selector Tool





PDF Object Selector Tool:




The PDF Object Selector is the primary tool for locating, navigating to, and displaying information on document objects in the COS Object Tree. It can be accessed from one of three places:

1. From the “Advanced Editing” toolbar on the PDF CanOpener flyout buttons .
2. From the PDF CanOpener Logo Button  in the title area of the PDF CanOpener Display.
3. From the COS Tree context menu.

Once activated, the document cursor changes to . As the cursor passes over objects on the document they are highlighted and the cursor changes to indicate the type of object and nesting level. Object highlighting can seem slow at times because of frequent redraws on large and complex pages. Highlighting can be turned off in the PDF CanOpener Preferences. If the cursor hovers over a single object for more than one second the info popup for that object is displayed.

Seven types of objects are recognized by the Selector Tool.

1. **Page Object** – No Highlight, No Popup. Displays the  cursor.
2. **Annotation Objects** – Blue Highlight. Displays the  cursor.
3. **Form XObjects** – Green Highlight. Displays the  cursor.
4. **Image XObjects** – Red Highlight. Displays the  cursor.

5. **Text Runs** – Black Highlight. Displays the  cursor.
6. **Drawing Paths** – Yellow Highlight. Displays the  cursor.
7. **Shading Objects** – Light Blue Highlight. Displays the  cursor

All objects are recognized whether they are on the page or inside an annotation, and even if they are hidden.

Finding nested and overlapped objects

All of these objects may overlay one another, effectively hiding them from the Selector Tool. The Selector Tool normally chooses the first object it finds (the top object). There are two methods of moving the selection through the stacked objects.

1. Press the **Arrow Down Key** to select the next object and the **Up Arrow Key** to select the previous object. When the bottom of the stack is reached the cursor selects the current page, and then starts over at the top of the stack.
2. Press and hold the **Ctrl key**, then Left click the mouse to select the next overlapped object. When the bottom of the stack is reached the cursor selects the current page, and then starts over at the top of the stack.

As the selection is moved through the overlapped objects the stack level is displayed on the lower left corner of the cursor. In **Figure 4** an overlap level of 3 is shown. This means that the selected path is the third object down in the stacking order.

Objects may also be nested inside other objects. Currently, the selector tool only recognizes XObject and Annotation nesting. The nesting level is displayed on the upper left corner of the cursor. In **Figure 4** a nesting level of 1 is shown. This means that the selected path object is inside of an XObject that is on the page. The associated XObject would have a nesting level of 0, since it is referenced in the main page content stream. However, the overlap level of the containing XObject would be 4, i.e. the 4th object in the stacking order. This is because the selector tool walks all content of a container (the XObject) first, before selecting the container. Another way to think of this is that as the Selector Tool walks the object hierarchy looking for a match with the cursor location, it matches contained objects on the way down, and the content containers on the way back up the tree. The nesting level indicator can be turned off from the Preferences Dialog. The Overlap level indicator cannot be turned off.

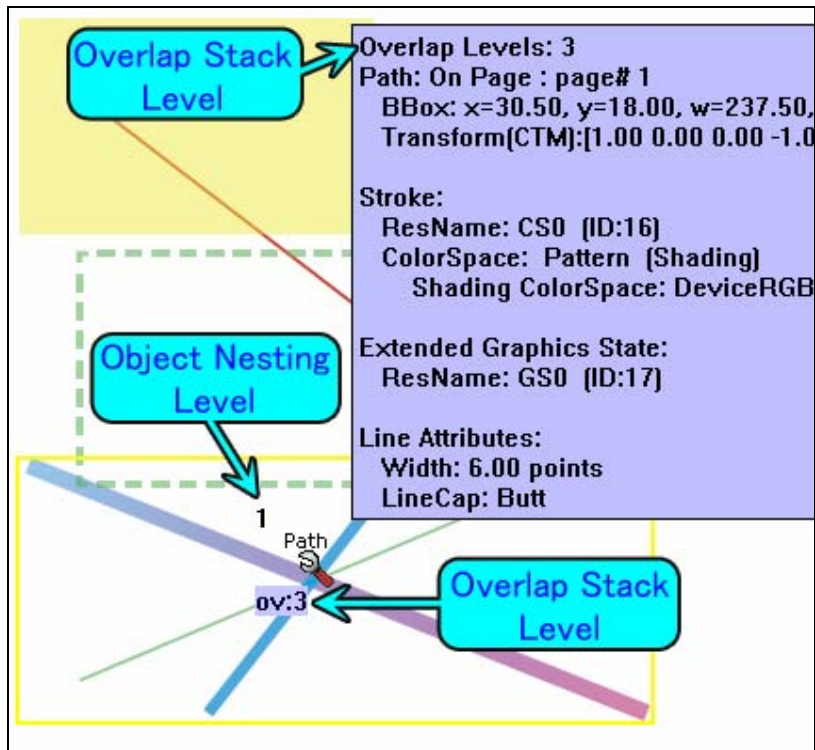


Figure 4 – Selecting Overlapped Objects. The PDF page is composed of several overlapped shapes.

When the Selector Tool is clicked on one of these objects the corresponding tree node is selected and displayed. A second click on the same object will cause its' node in the COS tree to expand. If the tool is clicked on a point in a document page that does not contain any Objects, the COS tree node corresponding to that page is selected and expanded. If a document is clicked on when this tool is active, the PDF CanOpener Display will be activated for that document if it has not already been activated. Annotations, XObjects, and Shading Objects are represented by actual COS objects in the PDF file. Text runs and Drawing paths on the other hand, are drawing operators in either a page content stream or a Form XObject. They do not have a corresponding object in the COS Tree. When a Text Run is clicked on the COS Tree is navigated to the font used in the text run. When a Path is clicked, the COS Tree is navigated to the content stream or XObject it is contained in.

Object Selector Context Menu

There are several features associated with objects selected by the Object Selector Tool. These features are listed on a menu that is accessed by placing the selector tool over an object on the PDF page and pressing the right mouse button. There is a standard set of menu items (**Figure 5**) and special menu items for different object types (**Figure 6**).

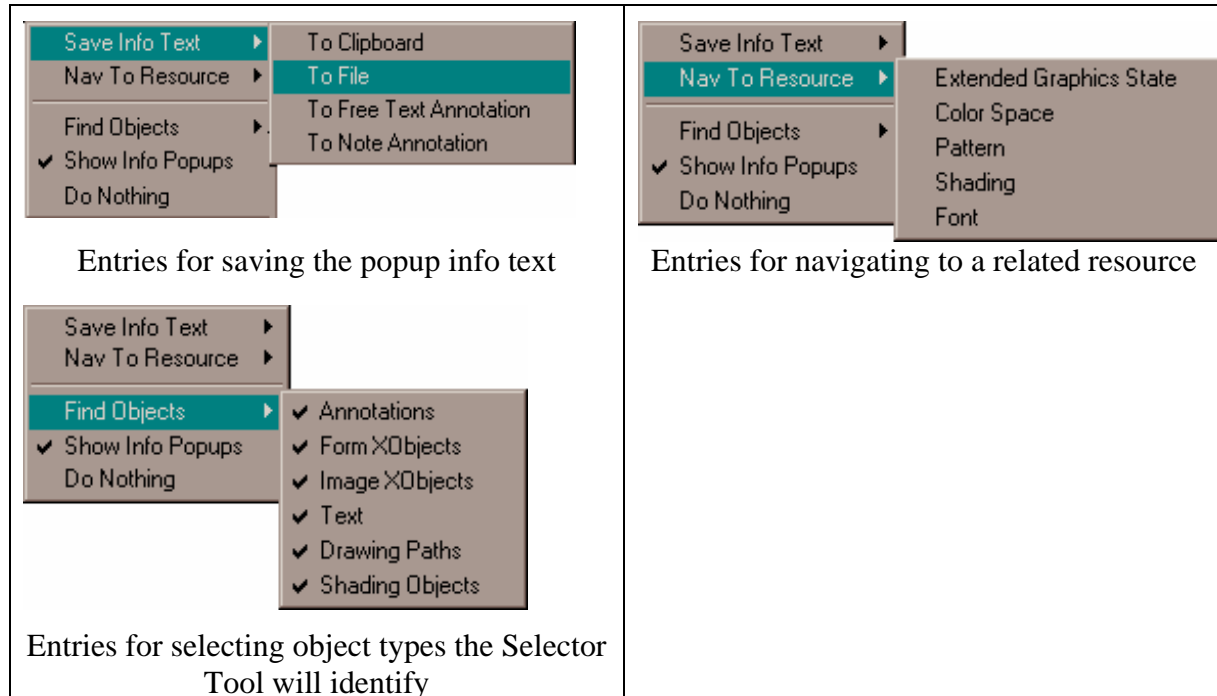


Figure 5 – General entries for the Selector Content Menu

General Menu Items:

These menu items are always available on the right click menu for the Selector Tool.

- **Save Info Text:**
Displays a sub-menu of 4 options for saving the text that appears in the Info Popup, for the currently selected object.
 - **To Clipboard:** Copies the text to the system clipboard.
 - **To File:** Displays the *File Save* dialog and writes to a “.txt” file.
 - **To Free Text Annotation:** Places a Free Text (Text Box) annotation onto the page in the exact position and size of the Info Popup. A callout arrow points to the original cursor location.
 - **To Note Annotation:** Places a Note annotation on the page at the cursor position.

- **Nav To Resource:**
Displays a sub-menu of supporting resource types associated with the currently selected object. Which items are placed on the menu depends on the type and the available resources for the selected object. For example, when the Selector Cursor shows “page”, this menu will be empty. This item is used to navigate the COS Tree to the selected resource.
 - **Extended Graphics State:** All types except Annotation.
 - **Color Space:** All types except Annotations and Form XObjects.
 - **Pattern:** Drawing Paths and Text objects.
 - **Shading:** Drawing Paths and Text objects.
 - **Font:** Only Text Objects.
- **Find Objects:**
Displays a sub-menu of object types that are identified by the Selector Tool. Toggles detection of these types on and off.
 - **Annotations**
 - **Form XObjects**
 - **Image XObjects**
 - **Text**
 - **Drawing Paths**
 - **Shading Objects**
- **Show Info Popups:**
Indicates and controls whether Info Popups are displayed when the Selector Tool hovers over a drawn object on the page.
- **Do Nothing:**
Performs no actions. This item is a convenience for exiting the menu without making a selection. Clicking the cursor outside of the menu or pressing the Esc key will also exit the menu. However these actions may also cause undesired results in the PDF and COS Tree. This item is also a convenient way to place focus on the page view window without causing anything else to happen.

Special Menu Items:

Special menu items are provided for two objects, the Text Object and XObjects.

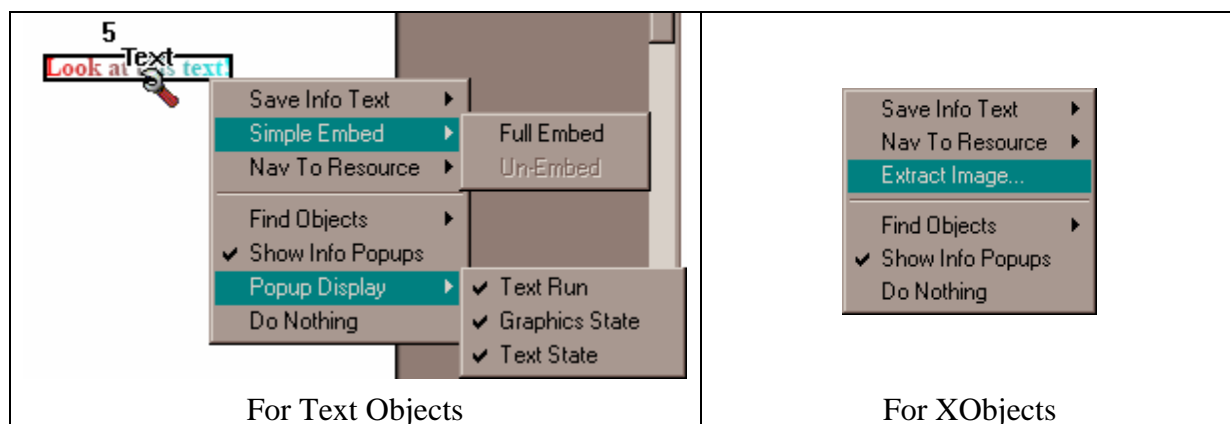


Figure 6 – Special entries in the context menu

Text Object:

- **Simple Embed:**

Provides for simple font embedding and un-embedding. PDF CanOpener does not do font sub-setting. These options are enabled only if the operation can be performed on the specified font. It does the operation only for the font associated with the selected text run. A font must be embeddable (ex: CID and Type3 fonts are not embeddable; CID fonts require more complex operations and Type3 fonts are a different kind of animal) and a similar font must be found on the system for the Embed options to be enabled. The enabling of the Un-Embed option is much more liberal, since this is much easier to do.

 - **Full Embed:** Copies a known font function into the FontDescriptor dictionary for the font.
 - **Un-Embed:** Deletes the font function from the FontDescriptor dictionary for the font.
- **Popup Display:**

The Info Popup for Text Runs is usually very large since it has the most parameters associated with it. To reduce the size of this popup some options are provided for turning unwanted information off and on.

 - **Text Run:** Information about the characters in the text run, number of bytes, number of characters, and the character codes used.
 - **Graphics State:** Information about the transformation matrix, location and bounding box of the text on the page.
 - **Text State:** Font geometry parameters, including the point size, scaling, and separation, as well as other parameters. Only parameters specified in the content stream are displayed.

XObject:

- **Extract Image:** XObjects (both image and form) can be rendered separately by the Acrobat SDK. This option takes advantage of this SDK capability to create bitmap files. Objects are rendered at their natural point size, one point/pixel, on a white background.

Note: Acrobat 7.0, 7.05, and 7.07 contain a bug which prevents transparent objects from being rendered properly under certain circumstances. Acrobat 5 and 6 don't have this problem. Since the Stream Editor uses the exact same functionality to render images, the rule of thumb is that if it cannot be rendered properly in the Stream Editor then it cannot be rendered as a bitmap file.

Object Search:

The search dialog is activated from a menu option on the [COS Tree context menu](#).

Even the simplest PDF has a large and complex object structure. In order to make finding a particular object easier PDF CanOpener provides a search dialog. Two search methods are provided, search by indirect ID number and search by text matching. These options are limited by the types of COS objects to which they can be applied.

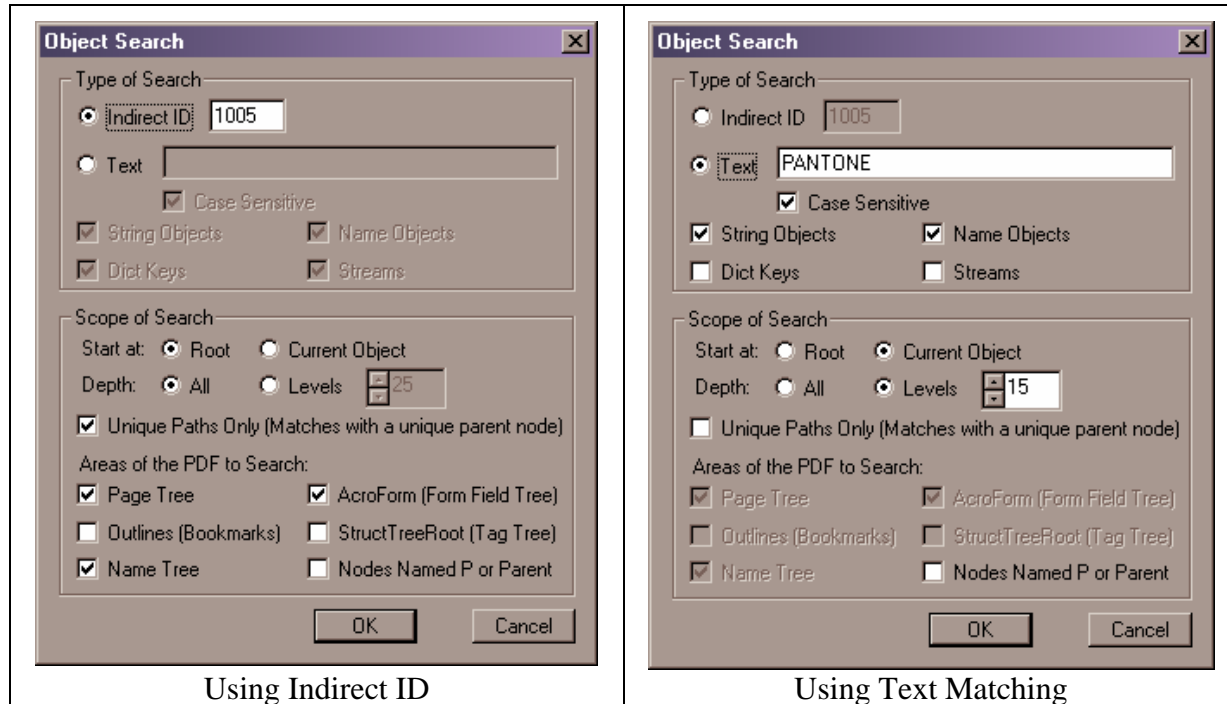


Figure 7 – Search dialog options

Indirect ID Search

There is only one input required for this option, the indirect ID number. If the currently selected node in the COS tree is an indirect node then its ID is used as the default.

Text Search

The text match can be case sensitive or case insensitive. Case sensitive searches are slightly faster. There are 4 locations where text may be found in a PDF: Name Objects, String Objects, Streams, and Dictionary keys. In order to improve performance, options are provided for selecting the COS Object types searched.

Scoping Options

Three scoping options are provided: starting object, tree depth, and unique paths. These options narrow the scope of the search for improving performance and excluding unwanted information.

Starting Object:

The search starts with either the root node (the Catalog) or the currently selected object in the COS Tree, and moves downward. However, it's possible for indirect objects to move the search into other parts of the COS Tree or even to nodes above the starting node.

Tree Depth Limit:

The search may continue to the bottom of the object tree or be limited. Each time the search pushes into an Array, Dictionary, or Stream, the search level is incremented. If a limit of 0 levels is set, then only the current node is searched. This only makes sense for a text search on a stream.

Unique Paths Only

The COS Object Tree may contain references to the same indirect object in many locations. This creates a situation where there are both infinite search loops and multiple redundant paths to the same object. PDF CanOpener tracks all search paths to remove loops and gives the user the option to block or allow redundant path selection. Redundancy is checked by keeping track of the parent nodes of all found objects. If unique paths are requested the search rejects any paths containing any parent nodes already encountered in the search. This action will yield a list of all unique paths to the matching node. Otherwise, all paths to matching nodes are collected.

Areas of PDF to Search

The search can be narrowed to only look in certain areas of the PDF when the search is started at the PDF Root node (the document Catalog). These areas include the Page Tree, the AcroForm dictionary (the field tree), the Bookmark tree, the Name Tree, and the StructTreeRoot. Of these areas the StructTreeRoot is the most troublesome because it has indirect references to all page objects in the PDF as well as references to other objects like fields and XObjects. This level of cross linking can deliver an annoying number of search matches and seriously limit performance. For this reason, both the StructTreeRoot and the Bookmarks are taken out of the search by default.

Many nodes in the PDF contain references to their parent. These references will cause the search to loop back on itself. The search algorithm has logic to prevent continuous looping, but the search can be made faster and return more relevant results if these parent nodes are ignored. These nodes are most often named “P” or “Parent” so an option is provided for explicitly including nodes with these names. By defaulted this setting is turned off and it is recommended that this option never be turned on.

Search Results Dialog:

This dialog is displayed when selected by the user or on completion of a search. It shows a list of paths through the COS Object Tree, starting at the root node. The path syntax is similar to an ordinary URL path with some special path separators for the different types of container objects in a PDF.

- // – Root node
- /*name* – Node is the *name* entry in a Dictionary
- [*n*] – Node is the *n*th entry in an Array
- ~*name* – Node is the *name* entry in a Stream Attributes Dictionary

Clicking on the “Find in Tree” button or Double clicking on a path in the list will open the *tree view* to that path. This dialog is a floating window, so it will not block other operations in Acrobat.

Stream Editor:

While streams are only one of the 8 PDF COS object types, they are the most complex and possibly most important of these objects. As such, a special Stream Editor is provided in PDF CanOpener. A Stream Object is composed of a block of binary data and a set of dictionary entries. Streams can also be encoded with up to 6 daisy chained filters. To deal with this complexity 3 dialog boxes associated with the Stream Editor. The first dialog displays information about the stream and acts as a jump station to the other two dialogs. The second dialog is for editing the block of binary data in either a text or hex editor. PDF CanOpener provides a basic editor for this purpose, but both external hex and text editors can be used instead of the default, and this is recommended to take advantage of advanced options these editors can provide. The external editors are set up in the PDF CanOpener Preferences. The third dialog is for handling the stream's data source, compression filters or encoders, and other features.

Stream Viewer:

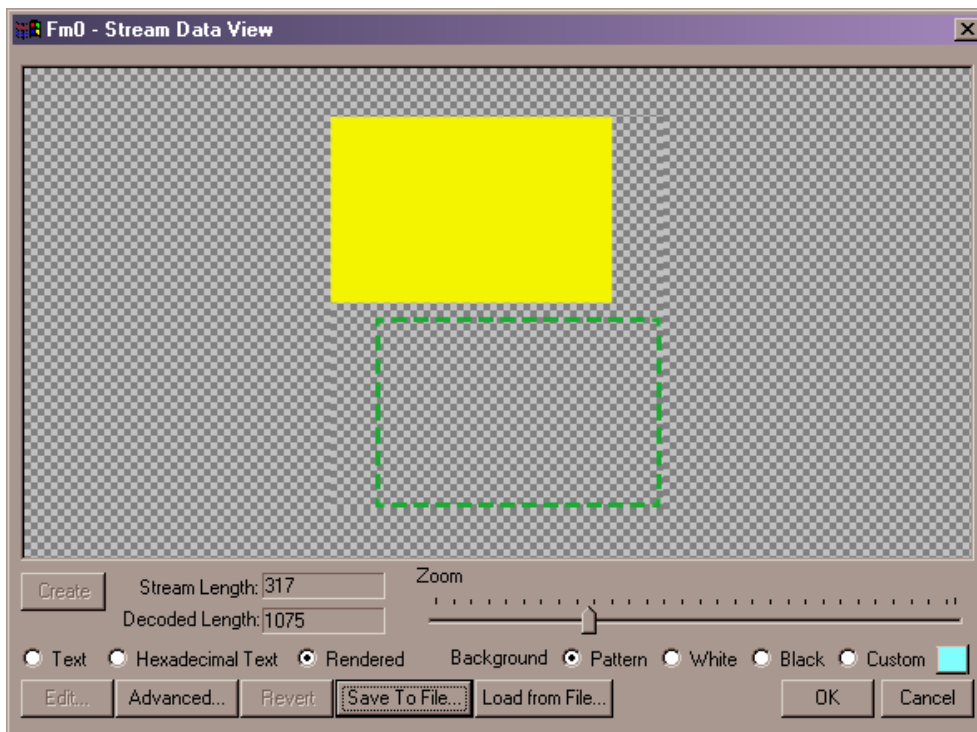


Figure 8 – Viewing a Form XObject in Render mode with a pattern background.

Stream data is displayed in this dialog box when the stream is selected for editing from the COS tree. The window can be resized to accommodate better viewing of the data. Three choices of data viewing formats are provided.

Data Viewing Options

The stream data may be viewed in one of three formats, text, hexadecimal text, or a rendered image. Not all of these options are available for all streams.

1. **Text** – Data is displayed in the character format specified by the data itself, if applicable and the character set is available on the system. The display of binary data is unpredictable and may not be viewable at all.

2. **Hexadecimal Text** – Data is converted to a hexadecimal string and each line is prefixed with the address of the first byte in the line. The address is shown in angle brackets i.e., <0x0010> for byte 16. Each line is organized into groups of 8 bytes. The number of groups in each line is determined by the width of the dialog box.
3. **Rendered** – This option is only valid for Page Content Streams, Page Content Arrays, XObjects, Tiling Patterns, Thumbnails, and Type3 font Character Procs. When selected, the object is drawn in the window. Page Content Arrays cannot be edited since they are not streams. This is just a convenience for viewing the combined streams. The image zoom can be controlled with either a slider control or the “-” (zoom down) and “=” (zoom up) keys when the input focus is on the drawing window. The image can be panned with either the scroll bars or by pressing the left mouse button and dragging the cursor across the image.

Rendering Background

When **Rendered** is selected as the Data Viewing option the background controls become visible. These options select the how the background is drawn before the stream is rendered so the best, alternate views of the image can be seen. Four choices are provided:

1. **Pattern** – This is the default option. It is a pattern of small light and dark gray squares.
2. **White** – A solid white background.
3. **Black** – A solid black background.
4. **Custom** – A solid color chosen by the user.

Action Buttons

- **“Edit” Button** – Activates the Stream Edit Dialog. Enabled only when **Text** or **Hexadecimal** is selected for the Data Viewing option. If external editors have been selected then the external text editor will be activated when the **Text Data View** option is selected and the external hex editor will be activated when the **Hexadecimal Data View** option is selected. If an external editor is used it must be closed (exited) before the data will be imported back into the Stream Viewer.
- **Advanced Button** – Activates the Advanced Edit Dialog for controlling the stream’s data source and encoding.
- **“Revert” Button** – Enabled after stream has been changed in the editor. Reverts the displayed data and all other stream parameters to the original values.
- **“Save To File” Button** – Write stream data to a file in one of five formats.
 1. *Decoded* – Decoded stream bytes written to the file.
 2. *Encoded* – Stream bytes are encoded with the stream filters and written to the file.
 3. *Encoded and Encrypted* – Stream bytes are encoded with the stream filters, then encrypted with any security handler that has been selected for this stream before being written to the file.
 4. *Hexadecimal String* – Data is converted to hexadecimal text. Data can be written out as a continuous string of hexadecimal characters or formatted into columns of 8 bytes each. A Byte address can be prefixed to each line.
 5. *JavaScript Icon for Acrobat* – For all XObjects. Data from the rendered XObject is written out to a JavaScript file with all code necessary to produce an icon that can be used in an Acrobat JavaScript. A dialog box is provided to select colors directly from the rendered image for creating an alpha channel

mask. Color data in the JavaScript is limited to the default color space of the system on which the image is rendered.

- **“Load From File” Button** – Read data from a file directly into the viewer. Data is formatted according to the current settings.
- **“Create” Button** – This button is enabled when the viewer contains data for which a new Stream Object has not yet been created. Pressing this button will cause a new Stream Object to be created within the viewer. This new stream does not affect the original. It is necessary to regenerate the stream to see changes in the rendered view of an XObject. This option can also be used to validate stream filter parameters before saving the new stream to the document.

Other Features

- **Stream Length Display** – Both the decoded byte length of the stream and the stored byte length in the PDF File are displayed.
- **Zoom Slider** – Displayed when the Data View is set to **Rendered**.

Default Stream Edit Dialog:

This dialog is where the stream content is edited. However, only stream data that is displayed as text or hexadecimal can be edited here. This dialog, unlike the Stream Viewer, is not resizable. When activated, it is sized to match the Stream Viewer.

The editing capabilities of this dialog are sufficient but minimal. It is recommended that an external editor is used. Two external editors can be specified, one for text editing and one for hex editing. These editors should have low startup overhead. PDF CanOpener starts up the editor when the “Edit” button is pushed and reloads the data when the editor is closed. Just closing the file in the editor will not work. The external text editor should be for plain text. The Windows Notepad tool works well. Several of the popular code editors also work well, for example, MultiEdit Lite, and CodeView.

Hexadecimal data is displayed in the same way it is displayed in the viewer, with addresses prefixed to the lines. These addresses are only provided as a convenience. This is not a true hex editor and they do not change as text is added or deleted. They can be edited just like any other text in the window. Data returned from the edit is stripped of any text that appears between the angle brackets <>. **Deleting an angle bracket during the edit can corrupt the data. So don’t do it.**

Other Features:

- **“Unicode” Checkbox** – Controls interpretation of the edited data when it is saved back to the Stream Viewer. Can be used to convert between ANSI and Unicode. Be careful with this conversion. In many cases Unicode stream data makes no sense to Acrobat. Conversions are handled differently depending on the displayed format of the data.
 - *Character Display* – Data is converted one character to one character. For the Unicode to ANSI conversion, a loss of 1 byte/char, Acrobat provides the Best Roman Encoding available on the system.
 - *Hexadecimal Display* – Bytes are converted to either ANSI (1 byte/char) or Unicode (2 bytes/char).
- **“Recreate Stream” Checkbox** – If checked, a new stream object is created in the stream view after the OK button on this dialog is pressed.
- **“Advanced” Button** – Provides access to the “Advanced” Stream Options Dialog.

Advanced Stream Options Dialog:

This dialog contains features that control stream generation.

- **“Data Source” Selection** – Source of stream data.
 - *Local String* – Data taken from the stream edit dialog.
 - *Load from file* – Data is read from a file.
 - *External* – Data is external to the current file. Rather than load the data, a PDF File Specification is added to the stream dictionary.
- **“Source Data is Decoded” Checkbox** – Indicates the state of the stream data. Encoded data will be written directly into the stream. Decoded data is encoded before being written to the stream. Decoded data can require special encoding parameters. So know what you are doing before messing with stream filters.
- **“Type” and “Subtype” Textboxes** – These are Name entries in the stream attributes dictionary. They are provided here purely as a convenience.

Filter Parameter Editor:

Up to 6 daisy chained filters can be defined for encoding and decoding stream data. Parameter templates for the 6 standard PDF filters are provided. A custom filter can also be used by entering the name of a custom filter in the “Filter Type” edit box. On selecting the filter, the most common parameters used by that filter are entered into the “Filter Parameter Editor,” but no default values are provided.

The “Add,” “Edit,” and “Delete” buttons are used to manipulate the parameters of the current filter. Filter Parameters are restricted to either Integer or Boolean values. Parameter names can be changed by clicking twice (not a double click) on the parameter name in the list of filter parameters. Double clicking on a parameter is the same as clicking on the “Edit...” button.

Some Filters may require Encoding parameters that are different from the Decode parameters. This is true of the DCTDecode filter. When this filter is selected, the associated encode parameters are filled in. If the stream is an image and contains values for the Rows, Columns, and Color Space, then the encode parameters are filled in with the correct values. Otherwise, stream creation would always fail when a DCTDecoded stream was edited. This filter is the only one for which parameters are auto-generated.

Notes on Deleting:

- Deleting an object simply removes it from the container object it resides in.
- Indirect Objects may also be deleted completely from the document. An option is provided for this action on the “Delete Warning Message” dialog.
note: Completely removing an Indirect Object from the document will convert all references to this object to a NULL object. Acrobat 6.0 automatically deletes Null objects from Dictionaries when it detects them.

Notes on Editing:

- **Double Clicking on a COS tree node** has the same effect as selecting “Edit” or “Add” from the context menu.
- **Direct Scalar COS Objects** – COS objects in a PDF file cannot be directly modified. Instead, the original object is deleted from the object that contains it, either a Dictionary or an Array, and a new object is created to replace it. The effect on direct objects is the same as if the original object were modifiable.
- **Indirect Scalar COS Objects** – When an indirect scalar object is edited, a new indirect object is created. This new object replaces the reference to the original indirect object in it’s’ container. All other references to the original object remain unchanged.
- **String Objects** – Strings displayed as hex are stored in the PDF file as hex strings. When a string is edited in hex format the bytes are reinterpreted on completion of the edit into either ANSI (1 byte/char) or Unicode (2 bytes/char), depending on the setting of this switch in the string edit dialog. This conversion is done in a logical fashion. Conversions between Unicode and ANSI text are much different. See the [“Stream Edit Dialog”](#) section.
- **Dictionary Objects** – Dictionaries are added to, rather than edited, with the “Add Dictionary Entry” dialog. This dialog provides the information necessary to create a new object of any basic type, or an indirect reference to an existing object. If the Object Name matches an existing entry, that entry is replaced.
Information entered into the “Add Dictionary Entry” dialog:
 - Object type.
 - “Direct” or “Indirect” creation.
 - Object name (can be Null).
 - Content information specific to the object type.
- **Array Objects** – The “Add Array Element” dialog is identical to the “Add Dictionary Entry” dialog except that instead of providing a name for the new object it provides an index number. New objects may replace an existing array element, be inserted anywhere into the array, or be appended to the end of the array.

Preferences:

General Tab:

These options control access to potentially dangerous operations that can be performed on a COS Object.

- **Allow All Operations** – Allow user to modify or delete any PDF Object.
Default Off
- **Modify** – Allow user to modify any PDF Object not specifically marked as read only.
Default On.
- **Delete** – Allow user to delete any PDF Object not specifically marked as no-delete.
Default On.

PDF CanOpener automatically blocks some operations on some special PDF Objects. For example, a stream dictionary is set to no-delete and the stream length entry in the stream dictionary is set to read only. The “Allow All Operations” option overrides these automatic settings. If no options are selected, no edit or delete operations are allowed on any PDF Object.

These options control the Selector Tool Display

- **Show XObject Nesting Level on Cursor** – Turns the nesting level number on and off on the Object Selector cursor.
- **Show XObject Selection Highlight** – Turn Object highlighting on and off during object selection.

These options control the initial view of an object in the Stream Editor.

- **Display Page Contents as Text in Stream Editor** – When this option is unchecked the initial view of a Page Content Stream within the **Stream Editor** is **Rendered**, meaning the Page Content is drawn in the editor window. Drawing is a potentially slow process. Checking this option forces the initial view to **Text**, which loads much faster.

Note: the meaning of this option is inverted if the Control Key is pressed while the editor is being activated.

- **Display XObjects as Text in Stream Editor** – When this option is unchecked the initial view of an XObject within the **Stream Editor** is **Rendered**, meaning the XObject is drawn in the editor window. Drawing is potentially a slow process. Checking this option forces the initial view to **Text** for Form XObjects and **Hexadecimal Text** for Image XObjects.

Note: The meaning of this option is inverted if the Control Key is pressed while the editor is being activated.

Display Tab:

These options select one of the Acrobat navigation windows for the PDF CanOpener display to be drawn in.

- Signatures
- Bookmarks
- Pages
- Layers (Not Available in Acrobat 5.0)
- Content (Not Available in Acrobat 5.0)
- Fields
- Tags

Warnings Tab:

Several common warning dialogs are set up to be optional. The user can choose not to display the dialog again with a check box on the warning dialog. This Tab lists all warnings and provides check boxes to turn them on or off.

External Editors Tab:

Two external editors can be provided for use with the Stream Editor, one for editing text data and one for editing hexadecimal data. The editors are specified with a path to the executable, but they are used only when the associated checkbox is selected. Command line options can be included with the file path as long as they don't interfere with PDF CanOpener's use of the editor. For editing a file, the Stream Editor first writes the decoded stream data to temporary file. It then appends the file path to the end of the executable path specified in this preference and uses it to start a new process. The stream editor waits until the process is killed, supposedly by the user exiting the program. It waits a brief interval, then reloads and deletes the temporary file.

Because of the simplicity of this process the external editor should have a low startup and shutdown overhead. Lightweight editors are recommended.

Support:

WindJack Solutions provides email support for this product. Mail all inquiries to:

support@windjack.com

Or visit our web site at:

www.windjack.com