

HPSF HOW-TO

by Rainer Klute

1. How To Use the HPSF APIs

This HOW-TO is organized in three sections. You should read them sequentially because the later sections build upon the earlier ones.

1. The [first section](#) explains how to read the most important standard properties of a Microsoft Office document. Standard properties are things like title, author, creation date etc. It is quite likely that you will find here what you need and don't have to read the other sections.
2. The [second section](#) goes a small step further and focusses on reading additional standard properties. It also talks about exceptions that may be thrown when dealing with HPSF and shows how you can read properties of embedded objects.
3. The [third section](#) tells how to read non-standard properties. Non-standard properties are application-specific triples consisting of an ID, a type, and a value.

1.1. Reading Standard Properties

Note:

This section explains how to read the most important standard properties of a Microsoft Office document. Standard properties are things like title, author, creation date etc. Chances are that you will find here what you need and don't have to read the other sections.

The first thing you should understand is that properties are stored in separate documents inside the POI filesystem. (If you don't know what a POI filesystem is, read the [POIFS documentation](#).) A document in a POI filesystem is also called a **stream**.

The following example shows how to read a POI filesystem's "title" property. Reading other properties is similar. Consider the API documentation of `org.apache.poi.hpsf.SummaryInformation` to learn which methods are available!

The standard properties this section focusses on can be found in a document called `\005SummaryInformation` located in the root of the POI filesystem. The notation `\005` in the document's name means the character with the decimal value of 5. In order to read the title,

an application has to perform the following steps:

1. Open the document `\005SummaryInformation` located in the root of the POI filesystem.
2. Create an instance of the class `SummaryInformation` from that document.
3. Call the `SummaryInformation` instance's `getTitle()` method.

Sounds easy, doesn't it? Here are the steps in detail.

1.1.1. Open the document `\005SummaryInformation` in the root of the POI filesystem

An application that wants to open a document in a POI filesystem (POIFS) proceeds as shown by the following code fragment. (The full source code of the sample application is available in the *examples* section of the POI source tree as *ReadTitle.java*.)

```
import java.io.*;
import org.apache.poi.hpsf.*;
import org.apache.poi.poifs.eventfilesystem.*;

// ...

public static void main(String[] args)
    throws IOException
{
    final String filename = args[0];
    POIFSReader r = new POIFSReader();
    r.registerListener(new MyPOIFSReaderListener(),
        "\005SummaryInformation");
    r.read(new FileInputStream(filename));
}
```

The first interesting statement is

```
POIFSReader r = new POIFSReader();
```

It creates a `org.apache.poi.poifs.eventfilesystem.POIFSReader` instance which we shall need to read the POI filesystem. Before the application actually opens the POI filesystem we have to tell the `POIFSReader` which documents we are interested in. In this case the application should do something with the document `\005SummaryInformation`.

```
r.registerListener(new MyPOIFSReaderListener(),
    "\005SummaryInformation");
```

This method call registers a `org.apache.poi.poifs.eventfilesystem.POIFSReaderListener` with the `POIFSReader`. The `POIFSReaderListener` interface specifies the method `processPOIFSReaderEvent` which processes a document. The class `MyPOIFSReaderListener` implements the `POIFSReaderListener` and thus the `processPOIFSReaderEvent` method. The eventing POI filesystem calls this method when it finds the `\005SummaryInformation` document. In the sample application

HPSF HOW-TO

`MyPOIFSReaderListener` is a static class in the `ReadTitle.java` source file.

Now everything is prepared and reading the POI filesystem can start:

```
r.read(new FileInputStream(filename));
```

The following source code fragment shows the `MyPOIFSReaderListener` class and how it retrieves the title.

```
static class MyPOIFSReaderListener implements POIFSReaderListener
{
    public void processPOIFSReaderEvent(POIFSReaderEvent event)
    {
        SummaryInformation si = null;
        try
        {
            si = (SummaryInformation)
                PropertySetFactory.create(event.getStream());
        }
        catch (Exception ex)
        {
            throw new RuntimeException
                ("Property set stream \"" +
                 event.getPath() + event.getName() + "\": " + ex);
        }
        final String title = si.getTitle();
        if (title != null)
            System.out.println("Title: \"" + title + "\"");
        else
            System.out.println("Document has no title.");
    }
}
```

The line

```
SummaryInformation si = null;
```

declares a `SummaryInformation` variable and initializes it with `null`. We need an instance of this class to access the title. The instance is created in a `try` block:

```
si = (SummaryInformation)
    PropertySetFactory.create(event.getStream());
```

The expression `event.getStream()` returns the input stream containing the bytes of the property set stream named `\005SummaryInformation`. This stream is passed into the `create` method of the factory class `org.apache.poi.hpsf.PropertySetFactory` which returns a `org.apache.poi.hpsf.PropertySet` instance. It is more or less safe to cast this result to `SummaryInformation`, a convenience class with methods like `getTitle()`, `getAuthor()` etc.

The `PropertySetFactory.create` method may throw all sorts of exceptions. We'll deal with them in the next sections. For now we just catch all exceptions and throw a `RuntimeException` containing the message text of the origin exception.

If all goes well, the sample application retrieves the title and prints it to the standard output. As you can see you must be prepared for the case that the POI filesystem does not have a title.

```
final String title = si.getTitle();
if (title != null)
    System.out.println("Title: \"" + title + "\"");
else
    System.out.println("Document has no title.");
```

Please note that a Microsoft Office document does not necessarily contain the `\005SummaryInformation` stream. The documents created by the Microsoft Office suite have one, as far as I know. However, an Excel spreadsheet exported from StarOffice 5.2 won't have a `\005SummaryInformation` stream. In this case the applications won't throw an exception but simply does not call the `processPOIFSReaderEvent` method. You have been warned!

1.2. Additional Standard Properties, Exceptions And Embedded Objects

Note:

This section focusses on reading additional standard properties. It also talks about exceptions that may be thrown when dealing with HPSF and shows how you can read properties of embedded objects.

A couple of **additional standard properties** are not contained in the `\005SummaryInformation` stream explained above, for example a document's category or the number of multimedia clips in a PowerPoint presentation. Microsoft has invented an additional stream named `\005DocumentSummaryInformation` to hold these properties. With two minor exceptions you can proceed exactly as described above to read the properties stored in `\005DocumentSummaryInformation`:

- Instead of `\005SummaryInformation` use `\005DocumentSummaryInformation` as the stream's name.
- Replace all occurrences of the class `SummaryInformation` by `DocumentSummaryInformation`.

And of course you cannot call `getTitle()` because `DocumentSummaryInformation` has different query methods. See the Javadoc API documentation for the details!

In the previous section the application simply caught all **exceptions** and was in no way interested in any details. However, a real application will likely want to know what went wrong and act appropriately. Besides any IO exceptions there are three HPSF resp. POI specific exceptions you should know about:

NoPropertySetStreamException:

This exception is thrown if the application tries to create a `PropertySet` instance from a stream that is not a property set stream.

(`SummaryInformation` and `DocumentSummaryInformation` are subclasses of `PropertySet`.) A faulty property set stream counts as not being a property set stream at all. An application should be prepared to deal with this case even if it opens streams named `\005SummaryInformation` or `\005DocumentSummaryInformation` only. These are just names. A stream's name by itself does not ensure that the stream contains the expected contents and that this contents is correct.

UnexpectedPropertySetTypeException

This exception is thrown if a certain type of property set is expected somewhere (e.g. a `SummaryInformation` or `DocumentSummaryInformation`) but the provided property set is not of that type.

MarkUnsupportedException

This exception is thrown if an input stream that is to be parsed into a property set does not support the `InputStream.mark(int)` operation. The POI filesystem uses the `DocumentInputStream` class which does support this operation, so you are safe here. However, if you read a property set stream from another kind of input stream things may be different.

Many Microsoft Office documents contain **embedded objects**, for example an Excel sheet on a page in a Word document. Embedded objects may have property sets of their own. An application can open these property set streams as described above. The only difference is that they are not located in the POI filesystem's root but in a **nested directory** instead. Just register a `POIFSReaderListener` for the property set streams you are interested in. For example, the `POIBrowser` application in the contrib section tries to open each and every document in a POI filesystem as a property set stream. If this operation was successful it displays the properties.

1.3. Reading Non-Standard Properties

Note:

This section tells how to read non-standard properties. Non-standard properties are application-specific ID/type/value triples.

1.3.1. Overview

Now comes the real hardcode stuff. As mentioned above, `SummaryInformation` and `DocumentSummaryInformation` are just special cases of the general concept of a property set. This concept says that a **property set** consists of properties and that each **property** is an entity with an **ID**, a **type**, and a **value**.

Okay, that was still rather easy. However, to make things more complicated, Microsoft in its infinite wisdom decided that a property set shall be broken into one or more **sections**. Each section holds a bunch of properties. But since that's still not complicated enough, a section may have an optional **dictionary** that maps property IDs to **property names** - we'll explain later what that means.

The procedure to get to the properties is the following:

1. Use the **PropertySetFactory** class to create a `PropertySet` object from a property set stream. If you don't know whether an input stream is a property set stream, just try to call `PropertySetFactory.create(java.io.InputStream)`: You'll either get a `PropertySet` instance returned or an exception is thrown.
2. Call the `PropertySet`'s method `getSections()` to get the sections contained in the property set. Each section is an instance of the `Section` class.
3. Each section has a format ID. The format ID of the first section in a property set determines the property set's type. For example, the first (and only) section of the `SummaryInformation` property set has a format ID of `F29F85E0-4FF9-1068-AB-91-08-00-2B-27-B3-D9`. You can get the format ID with `Section.getFormatID()`.
4. The properties contained in a `Section` can be retrieved with `Section.getProperties()`. The result is an array of `Property` instances.
5. A property has a name, a type, and a value. The `Property` class has methods to retrieve them.

1.3.2. A Sample Application

Let's have a look at a sample Java application that dumps all property set streams contained in a POI file system. The full source code of this program can be found as `ReadCustomPropertySets.java` in the *examples* area of the POI source code tree. Here are the key sections:

```
import java.io.*;
import java.util.*;
import org.apache.poi.hpsf.*;
import org.apache.poi.poifs.eventfilesystem.*;
import org.apache.poi.util.HexDump;
```

The most important package the application needs is `org.apache.poi.hpsf.*`. This package contains the HPSF classes. Most classes named below are from the HPSF package. Of course we also need the POIFS event file system's classes and `java.io.*` since we are dealing with POI I/O. From the `java.util` package we use the `List` and `Iterator` class. The class `org.apache.poi.util.HexDump` provides a methods to dump byte arrays as nicely formatted strings.

```
public static void main(String[] args)
```

HPSF HOW-TO

```
throws IOException
{
    final String filename = args[0];
    POIFSReader r = new POIFSReader();

    /* Register a listener for *all* documents. */
    r.registerListener(new MyPOIFSReaderListener());
    r.read(new FileInputStream(filename));
}
```

The POIFSReader is set up in a way that the listener MyPOIFSReaderListener is called on every file in the POI file system.

1.3.3. The Property Set

The listener class tries to create a PropertySet from each stream using the PropertySetFactory.create() method:

```
static class MyPOIFSReaderListener implements POIFSReaderListener
{
    public void processPOIFSReaderEvent(POIFSReaderEvent event)
    {
        PropertySet ps = null;
        try
        {
            ps = PropertySetFactory.create(event.getStream());
        }
        catch (NoPropertySetStreamException ex)
        {
            out("No property set stream: \"" + event.getPath() +
                event.getName() + "\"");
            return;
        }
        catch (Exception ex)
        {
            throw new RuntimeException
                ("Property set stream \"" +
                 event.getPath() + event.getName() + "\": " + ex);
        }

        /* Print the name of the property set stream: */
        out("Property set stream \"" + event.getPath() +
            event.getName() + "\"");
    }
}
```

Creating the PropertySet is done in a try block, because not each stream in the POI file system contains a property set. If it is some other file, the PropertySetFactory.create() throws a NoPropertySetStreamException, which is caught and logged. Then the program continues with the next stream. However, all other types of exceptions cause the program to terminate by throwing a runtime exception. If all went well, we can print the name of the property set stream.

1.3.4. The Sections

The next step is to print the number of sections followed by the sections themselves:

```

/* Print the number of sections: */
final long sectionCount = ps.getSectionCount();
out("  No. of sections: " + sectionCount);

/* Print the list of sections: */
List sections = ps.getSections();
int nr = 0;
for (Iterator i = sections.iterator(); i.hasNext(); )
{
    /* Print a single section: */
    Section sec = (Section) i.next();

    // See below for the complete loop body.
}

```

The `PropertySet`'s method `getSectionCount()` returns the number of sections.

To retrieve the sections, use the `getSections()` method. This method returns a `java.util.List` containing instances of the `Section` class in their proper order.

The sample code shows a loop that retrieves the `Section` objects one by one and prints some information about each one. Here is the complete body of the loop:

```

/* Print a single section: */
Section sec = (Section) i.next();
out("  Section " + nr++ + ":");
String s = hex(sec.getFormatID().getBytes());
s = s.substring(0, s.length() - 1);
out("    Format ID: " + s);

/* Print the number of properties in this section. */
int propertyCount = sec.getPropertyCount();
out("    No. of properties: " + propertyCount);

/* Print the properties: */
Property[] properties = sec.getProperties();
for (int i2 = 0; i2 < properties.length; i2++)
{
    /* Print a single property: */
    Property p = properties[i2];
    int id = p.getID();
    long type = p.getType();
    Object value = p.getValue();
    out("      Property ID: " + id + ", type: " + type +
        ", value: " + value);
}

```

1.3.5. The Section's Format ID

HPSF HOW-TO

The first method called on the `Section` instance is `getFormatID()`. As explained above, the format ID of the first section in a property set determines the type of the property set. Its type is `ClassID` which is essentially a sequence of 16 bytes. A real application using its own type of a custom property set should have defined a unique format ID and, when reading a property set stream, should check the format ID is equal to that unique format ID. The sample program just prints the format ID it finds in a section:

```
String s = hex(sec.getFormatID().getBytes());
s = s.substring(0, s.length() - 1);
out("    Format ID: " + s);
```

As you can see, the `getFormatID()` method returns a `ClassID` object. An array containing the bytes can be retrieved with `ClassID.getBytes()`. In order to get a nicely formatted printout, the sample program uses the `hex()` helper method which in turn uses the POI utility class `HexDump` in the `org.apache.poi.util` package. Another helper method is `out()` which just saves typing `System.out.println()`.

1.3.6. The Properties

Before getting the properties, it is possible to find out how many properties are available in the section via the `Section.getPropertyCount()`. The sample application uses this method to print the number of properties to the standard output:

```
int propertyCount = sec.getPropertyCount();
out("    No. of properties: " + propertyCount);
```

Now its time to get to the properties themselves. You can retrieve a section's properties with the method `Section.getProperties()`:

```
Property[] properties = sec.getProperties();
```

As you can see the result is an array of `Property` objects. This class has three methods to retrieve a property's ID, its type, and its value. The following code snippet shows how to call them:

```
for (int i2 = 0; i2 < properties.length; i2++)
{
    /* Print a single property: */
    Property p = properties[i2];
    int id = p.getID();
    long type = p.getType();
    Object value = p.getValue();
    out("    Property ID: " + id + ", type: " + type +
        ", value: " + value);
}
```

1.3.7. Sample Output

The output of the sample program might look like the following. It shows the summary information and the document summary information property sets of a Microsoft Word

document. However, unlike the first and second section of this HOW-TO the application does not have any code which is specific to the SummaryInformation and DocumentSummaryInformation classes.

```
Property set stream "/SummaryInformation":
```

```
No. of sections: 1
```

```
Section 0:
```

```
Format ID: 00000000 F2 9F 85 E0 4F F9 10 68 AB 91 08 00 2B 27 B3 D9 .....O..h.....+
```

```
No. of properties: 17
```

```
Property ID: 1, type: 2, value: 1252
```

```
Property ID: 2, type: 30, value: Titel
```

```
Property ID: 3, type: 30, value: Thema
```

```
Property ID: 4, type: 30, value: Rainer Klute (Autor)
```

```
Property ID: 5, type: 30, value: Test (Stichwörter)
```

```
Property ID: 6, type: 30, value: This is a document for testing HPSF
```

```
Property ID: 7, type: 30, value: Normal.dot
```

```
Property ID: 8, type: 30, value: Unknown User
```

```
Property ID: 9, type: 30, value: 3
```

```
Property ID: 18, type: 30, value: Microsoft Word 9.0
```

```
Property ID: 12, type: 64, value: Mon Jan 01 00:59:25 CET 1601
```

```
Property ID: 13, type: 64, value: Thu Jul 18 16:22:00 CEST 2002
```

```
Property ID: 14, type: 3, value: 1
```

```
Property ID: 15, type: 3, value: 20
```

```
Property ID: 16, type: 3, value: 93
```

```
Property ID: 19, type: 3, value: 0
```

```
Property ID: 17, type: 71, value: [B@13582d
```

```
Property set stream "/DocumentSummaryInformation":
```

```
No. of sections: 2
```

```
Section 0:
```

```
Format ID: 00000000 D5 CD D5 02 2E 9C 10 1B 93 97 08 00 2B 2C F9 AE .....+
```

```
No. of properties: 14
```

```
Property ID: 1, type: 2, value: 1252
```

```
Property ID: 2, type: 30, value: Test
```

```
Property ID: 14, type: 30, value: Rainer Klute (Manager)
```

```
Property ID: 15, type: 30, value: Rainer Klute IT-Consulting GmbH
```

```
Property ID: 5, type: 3, value: 3
```

```
Property ID: 6, type: 3, value: 2
```

```
Property ID: 17, type: 3, value: 111
```

```
Property ID: 23, type: 3, value: 592636
```

```
Property ID: 11, type: 11, value: false
```

```
Property ID: 16, type: 11, value: false
```

```
Property ID: 19, type: 11, value: false
```

```
Property ID: 22, type: 11, value: false
```

```
Property ID: 13, type: 4126, value: [B@56a499
```

```
Property ID: 12, type: 4108, value: [B@506411
```

```
Section 1:
```

```
Format ID: 00000000 D5 CD D5 05 2E 9C 10 1B 93 97 08 00 2B 2C F9 AE .....+
```

```
No. of properties: 7
```

```
Property ID: 0, type: 0, value: {6=Test-JaNein, 5=Test-Zahl, 4=Test-Datum, 3=Test
```

```
Property ID: 1, type: 2, value: 1252
```

```
Property ID: 2, type: 65, value: [B@c9ba38
```

```
Property ID: 3, type: 30, value: This is some text.
```

```
Property ID: 4, type: 64, value: Wed Jul 17 00:00:00 CEST 2002
```

HPSF HOW-TO

```
Property ID: 5, type: 3, value: 27
Property ID: 6, type: 11, value: true
No property set stream: "/WordDocument"
No property set stream: "/CompObj"
No property set stream: "/1Table"
```

There are some interesting items to note:

- The first property set (summary information) consists of a single section, the second property set (document summary information) consists of two sections.
- Each section type (identified by its format ID) has its own domain of property ID. For example, in the second property set the properties with ID 2 have different meanings in the two section. By the way, the format IDs of these sections are **not** equal, but you have to look hard to find the difference.
- The properties are not in any particular order in the section, although they slightly tend to be sorted by their IDs.

1.3.8. Property IDs

Properties in the same section are distinguished by their IDs. This is similar to variables in a programming language like Java, which are distinguished by their names. But unlike variable names, property IDs are simple integral numbers. There is another similarity, however. Just like a Java variable has a certain scope (e.g. a member variables in a class), a property ID also has its scope of validity: the section.

Two property IDs in sections with different section format IDs don't have the same meaning even though their IDs might be equal. For example, ID 4 in the first (and only) section of a summary information property set denotes the document's author, while ID 4 in the first section of the document summary information property set means the document's byte count. The sample output above does not show a property with an ID of 4 in the first section of the document summary information property set. That means that the document does not have a byte count. However, there is a property with an ID of 4 in the *second* section: This is a user-defined property ID - we'll get to that topic in a minute.

So, how can you find out what the meaning of a certain property ID in the summary information and the document summary information property set is? The standard property sets as such don't have any hints about the **meanings of their property IDs**. For example, the summary information property set does not tell you that the property ID 4 stands for the document's author. This is external knowledge. Microsoft defined standard meanings for some of the property IDs in the summary information and the document summary information property sets. As a help to the Java and POI programmer, the class `PropertyIDMap` in the `org.apache.poi.hpsf.wellknown` package defines constants for the "well-known" property IDs. For example, there is the definition

```
public final static int PID_AUTHOR = 4;
```

These definitions allow you to use symbolic names instead of numbers.

In order to provide support for the other way, too, - i.e. to map property IDs to property names - the class `PropertyIDMap` defines two static methods: `getSummaryInformationProperties()` and `getDocumentSummaryInformationProperties()`. Both return `java.util.Map` objects which map property IDs to strings. Such a string gives a hint about the property's meaning. For example, `PropertyIDMap.getSummaryInformationProperties().get(4)` returns the string "PID_AUTHOR". An application could use this string as a key to a localized string which is displayed to the user, e.g. "Author" in English or "Verfasser" in German. HPSF might provide such language-dependent ("localized") mappings in a later release.

Usually you won't have to deal with those two maps. Instead you should call the `Section.getPIDString(int)` method. It returns the string associated with the specified property ID in the context of the `Section` object.

Above you learned that property IDs have a meaning in the scope of a section only. However, there are two exceptions to the rule: The property IDs 0 and 1 have a fixed meaning in **all** sections:

Property ID	Meaning
0	The property's value is a dictionary , i.e. a mapping from property IDs to strings.
1	The property's value is the number of a codepage , i.e. a mapping from character codes to characters. All strings in the section containing this property must be interpreted using this codepage. Typical property values are 1252 (8-bit "western" characters) or 1200 (16-bit Unicode characters).

1.3.9. Property types

A property is nothing without its value. It is stored in a property set stream as a sequence of bytes. You must know the property's **type** in order to properly interpret those bytes and reasonably handle the value. A property's type is one of the so-called Microsoft-defined "**variant types**". When you call `Property.getType()` you'll get a long value which denoting the property's variant type. The class `Variant` in the `org.apache.poi.hpsf` package holds most of those long values as named constants. For example, the constant `VT_I4 = 3` means a signed integer value of four bytes. Examples of other types are `VT_LPSTR = 30` meaning a null-terminated string of 8-bit characters, `VT_LPWSTR =`

HPSF HOW-TO

31 which means a null-terminated Unicode string, or `VT_BOOL = 11` denoting a boolean value.

In most cases you won't need a property's type because HPSF does all the work for you.

1.3.10. Property values

When an application wants to retrieve a property's value and calls `Property.getValue()`, HPSF has to interpret the bytes making out the value according to the property's type. The type determines how many bytes the value consists of and what to do with them. For example, if the type is `VT_I4`, HPSF knows that the value is four bytes long and that these bytes comprise a signed integer value in the little-endian format. This is quite different from e.g. a type of `VT_LPWSTR`. In this case HPSF has to scan the value bytes for a Unicode null character and collect everything from the beginning to that null character as a Unicode string.

The good news is that HPSF does another job for you, too: It maps the variant type to an adequate Java type.

Variant type:	Java type:
<code>VT_I2</code>	<code>java.lang.Integer</code>
<code>VT_I4</code>	<code>java.lang.Long</code>
<code>VT_FILETIME</code>	<code>java.util.Date</code>
<code>VT_LPSTR</code>	<code>java.lang.String</code>
<code>VT_LPWSTR</code>	<code>java.lang.String</code>
<code>VT_CF</code>	<code>byte[]</code>
<code>VT_BOOL</code>	<code>java.lang.Boolean</code>

The bad news is that there are still a couple of variant types HPSF does not yet support. If it encounters one of these types it returns the property's value as a byte array and leaves it to be interpreted by the application.

An application retrieves a property's value by calling the `Property.getValue()` method. This method's return type is the abstract `Object` class. The `getValue()` method looks up the property's variant type, reads the property's value bytes, creates an instance of an adequate Java type, assigns it the property's value and returns it. Primitive types like `int` or `long` will be returned as the corresponding class, e.g. `Integer` or `Long`.

1.3.11. Dictionaries

The property with ID 0 has a very special meaning: It is a **dictionary** mapping property IDs to property names. We have seen already that the meanings of standard properties in the summary information and the document summary information property sets have been defined by Microsoft. The advantage is that the labels of properties like "Author" or "Title" don't have to be stored in the property set. However, a user can define custom fields in, say, Microsoft Word. For each field the user has to specify a name, a type, and a value.

The names of the custom-defined fields (i.e. the property names) are stored in the document summary information second section's **dictionary**. The dictionary is a map which associates property IDs with property names.

The method `Section.getPIDString(int)` not only returns with the well-known property names of the summary information and document summary information property sets, but with self-defined properties, too. It should also work with self-defined properties in self-defined sections.

1.3.12. Codepage support

FIXME (Rainer Klute):

Improve codepage support!

The property with ID 1 holds the number of the codepage which was used to encode the strings in this section. The present HPSF codepage support is still very limited: When reading property value strings, HPSF distinguishes between 16-bit characters and 8-bit characters. 16-bit characters should be Unicode characters and thus be okay. 8-bit characters are interpreted according to the platform's default character set. This is fine as long as the document being read has been written on a platform with the same default character set. However, if you receive a document from another region of the world and want to process it with HPSF you are in trouble - unless the creator used Unicode, of course.

1.3.13. Further Reading

There are still some aspects of HSPF left which are not covered by this HOW-TO. You should dig into the Javadoc API documentation to learn further details. Since you've struggled through this document up to this point, you are well prepared.