Developing Penlets

Livescribe™ Platform SDK
Version 1.5

Copyright and Trademarks

DevPenlets-PlatformSDK-1.5.0-REV-C

12/20/2010 10:52 AM

Table of Contents

Livescribe Platform SDK

The Livescribe Platform SDK is a set of development tools enabling the creation of Java applications that run on the Livescribe™ Smartpen J2ME platform and dot paper products that work with the smartpen.

Smartpen applications are known as *penlets*. Penlets can operate on Open Paper, Fixed Print Paper, or a combination of both. For a conceptual discussion of Livescribe dot paper, see *Introduction to the Livescribe Platform*.

Developing Open Paper Penlets

Open Paper applications can use any existing paper product that has Open Paper regions, including Livescribe notebooks and journals. The Piano, which ships with the Livescribe smartpen, is an example of an Open Paper application. It issues instructions to the user via the display and the audio speakers, creates the dynamic regions based on what the user draws, and captures Livescribe smartpen events on Open Paper as the user plays the piano. No pre-printed, Fixed Paper controls are involved.

The figure below shows the basic development process for penlet-only applications.

Penlet Only Development

Specify the Penlet — Java Programming — Penlet — Build — JAR file — Publish

Livescribe Java API and IDE Plug-in — Developer Tools
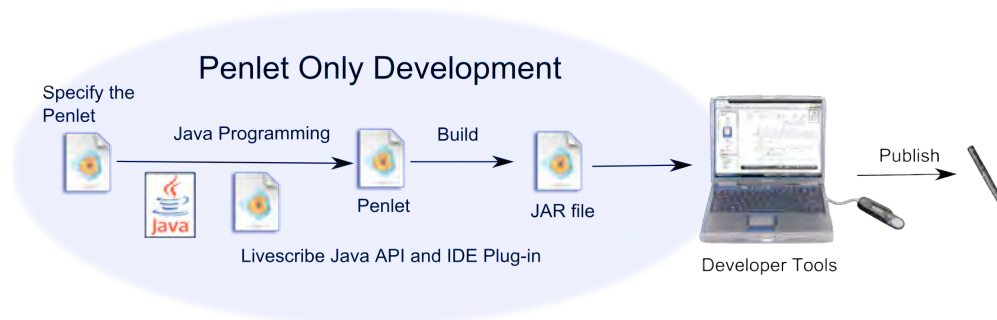
To develop Open Paper penlets, follow these general steps:

1. Install Eclipse and the Eclipse plugins from Livescribe. The plugins are packaged in a single Eclipse feature. For details, see the *Installing the Livescribe Platform SDK* document.

2. Code your penlet against the Livescribe Platform Java API. The current manual describes how to use the classes and methods exposed in the API.

3. Eclipse automatically builds your penlet. This process compiles the penlet code, pre-verifies the classes, and packages the penlet files into a JAR file.

4. When your penlet is complete, (install) the penlet JAR to the Livescribe smartpen from within Eclipse.
If you prefer, you can deploy the penlet to the Livescribe Smartpen Emulator and test your penlet and paper products on the desktop. Download and install the Livescribe Smartpen Emulator from the Livescribe Developer site, and read the *Livescribe Smartpen Emulator User Guide* for details.

5. Test the penlet code installed on the Livescribe smartpen against an Open Paper region of a Livescribe notebook.

6. Iterate through steps above until your Livescribe smartpen application is complete and tested.

Developing Fixed Print Penlets

A Fixed Print penlet uses dot paper that Livescribe has licensed to you, often including an association with your penlet. This Fixed Print paper is known as the *paper product* for your penlet. It contains the *static* regions that you define and to which you assign specific functionality in your penlet. Users of your penlet cannot access these static regions on the generic Open Paper notebooks and journals from Livescribe. Printed images normally indicate the location of static regions on your paper product.

In addition to static regions, your paper product can also support *dynamic* regions—that is, areas that are defined at run-time as the user interacts with the paper. A powerful paper product often combines static and dynamic regions. As an example, consider a generic Livescribe notebook as a paper product for Livescribe Paper Replay application. The controls along the bottom are static regions defined and shaped by the Livescribe engineers during development of Paper Replay. The blank space in the middle allows users to create dynamic regions as they take

notes while Paper Replay is running. Similarly, your paper products may consist of a combination of static and dynamic regions.

Note: Your penlet can create dynamic regions on any portion of licensed dot paper that does not have static regions defined by a particular penlet. The dynamic regions will belong to the current penlet. When a user taps on one of those regions in the future, that penlet will be activated and will receive notification of region tap via the penDown event.

Livescribe Paper Designer

Fixed Print penlet development uses the same Livescribe IDE as Open Paper penlets: Eclipse with a custom Eclipse feature developed by Livescribe. In addition to the Penlet Editor and Penlet Project type that you will have used for developing Open Paper penlets, you will need to use the Livescribe Paper Designer and the Paper Project type to create the paper product for your Fixed Print penlet. This tool allows you to define static regions, define the pages including artwork and Livescribe dots, produce test pages containing development-only dots, and request the production dots for your paper product from the Livescribe Pattern Server. The final output of the tool is a Postscript file that you can print out. The result will be your paper product, complete with your licensed dots.

To develop Fixed Print penlets, follow these general steps:

1. Use the Penlet Editor to develop Fixed Print penlets. For details, see the manual titled *Getting Started with the Livescribe Platform SDK*.

2. Define a paper product for your Fixed Print penlet, using the Livescribe Paper Designer of the Livescribe IDE. Define one or more static areas on the page(s) of your paper product. For details, see the manual titled *Developing Paper Products*.

3. Code your penlet against the Livescribe Platform Java API. This manual describes how to use the classes and methods exposed in the API.

4. Eclipse automatically builds your penlet. This process compiles the penlet code, pre-verifies the classes, and packages the penlet files into a JAR file.

5. When your penlet is complete, deploy (install) the penlet JAR to the Livescribe smartpen from within Eclipse.
    If you prefer, you can deploy the penlet to the Livescribe Smartpen Emulator and test your penlet and paper products on the desktop. Download and install the Livescribe Smartpen Emulator from the Livescribe Developer site, and read the *Livescribe Smartpen Emulator User Guide* for details.

6. Print out one or more test pages, using the Livescribe Paper Designer. Test pages contain developer dots, which can be used for the development process only.

7. Test the penlet code installed on the Livescribe smartpen against the test pages.

8. Iterate through steps above until your Fixed Print penlet is complete and tested. Do not continue until you are certain your penlet operates as you wish.

9. Request production pages from the Livescribe Pattern Server, using the Livescribe Paper Designer. The dots on these pages are licensed to you and will work with your penlet only.

10. Test the production pages with your penlet on your Livescribe smartpen.

Developing a Penlet

Although many kinds of penlets can be written for the Livescribe Platform, the basic anatomy of all penlets is the same. This section examines the `Penlet` class, details the penlet lifecycle management, and lists the packages included in the Livescribe Java API.

Main Penlet Class

The first step to create a penlet is to extend `com.livescribe.penlet.Penlet`. The smartpen runtime will only instantiate classes that derive from `Penlet`. The life cycle of these objects is managed by the smartpen runtime.

From this point onward, we refer to the class you write that directly extends `com.livescribe.penlet.Penlet` as your `Penlet` subclass. The term merely alludes to the importance of this class. You should *not* infer that the class has a `main` method. . There is *no* `public static void main` method in a penlet. Penlets are like MIDlets: they are launched by the runtime system, which manages their life cycle by calling specific methods.

One Penlet Active at a Time

Only one penlet can be active on the Livescribe smartpen at a time. Keep this in mind while you read about the penlet's life cycle. When the user selects another penlet, the runtime *deactivates and destroys the current penlet* and switches to the requested penlet. When the runtime switches back to the first penlet, the second one becomes inactive and is destroyed, and the first one is reinitiated and reactivated.

You should be aware that certain static regions are able to call Livescribe system functionality *without* deactivating and destroying your penlet. Examples of these regions include printed volume change and mute controls. For further details, see "Using Standard Livescribe Controls" in *Developing Paper Products*.

Penlet Life Cycle

The Livescribe system runtime manages the life cycle of a penlet in the following manner:

1. Instantiates the constructor of your `Penlet` subclass.

2. Calls `initApp`.

3. Calls `activateApp`.

4. Calls appropriate event handlers in registered listeners as various smartpen events arrive.

5. Calls `deactivateApp` when an event causes the runtime to move your penlet from the active state, due to activation of another penlet. When a penlet is inactive, the runtime system might also

    destroy it by calling `destroyApp`. See Developer Tasks in each Life Cycle Method for more information.

6. Calls `activateApp`, if the penlet becomes active again.

7. Calls appropriate event handlers in registered listeners as Livescribe smartpen events arrive.

8. Calls `deactivateApp` and `destroyApp` when the Livescribe smartpen is shut down, provided the penlet is running at that time.

Important Lifecycle Considerations

The above lifecycle flow is a general case, and in the future might change to support new features or improve performance.

Livescribe guarantees the following lifecycle events:

- `initApp` is the first method invoked after the class `init` method when a penlet is being created. `initApp` leaves the penlet in the "inactive" state.

- `activateApp` is invoked whenever a penlet is in the "inactive" state and the smartpen processes an activation event (that is, an active region tapped, menu invoked, or a snapback event occurred). `activateApp` leaves the penlet in the "active" state

- `deactivateApp` is invoked whenever the penlet is in the "active" state and the runtime wants to deactivate it. `deactivateApp` leaves the penlet in the "inactive" state.

- `destroyApp` is invoked whenever the penlet is in the "inactive" state and the runtime wants to destroy it. `destroyApp` is the last method that will be invoked on the penlet after which the class instance will be destroyed. The `destroyApp` method is invoked before the smartpen is powered off.

Other Considerations:

When developing your penlets, consider these things as well.

- `activateApp`/`deactivateApp` pairs may be invoked many times on a single penlet instance.

- `initApp`/`destroyApp` pairs may be invoked many times during a given power cycle of the smartpen.

- Smartpen firmware 1.0 through 1.7 invoked `destroyApp` once per smartpen power cycle (although this is not guaranteed).

- Smartpen firmware 2.0 and above invokes `destroyApp` after every `deactivateApp` (although this is not guaranteed).

The following diagram highlights the state transitions during the life cycle of the penlet.

initApp is the first method invoked after the class init method when a penlet is being created

initApp()

Penlet Initialized

activateApp is invoked whenever a penlet is in the "inactive" state and the smartpen processes an activation event:
- an active region tapped
- menu invoked
- snapback event occured

activateApp()

Penlet Active

deactivateApp is invoked whenever the penlet is in the "active" state and the runtime wants to deactivate it.

deactivateApp()

Penlet Inactive

destroyApp is invoked whenever the penlet is in the "inactive" state and the runtime wants to destroy it. destroyApp is the last method that will be invoked on the penlet after which the class instance will be destroyed. The destroyApp method is invoked before the smartpen is powered off.
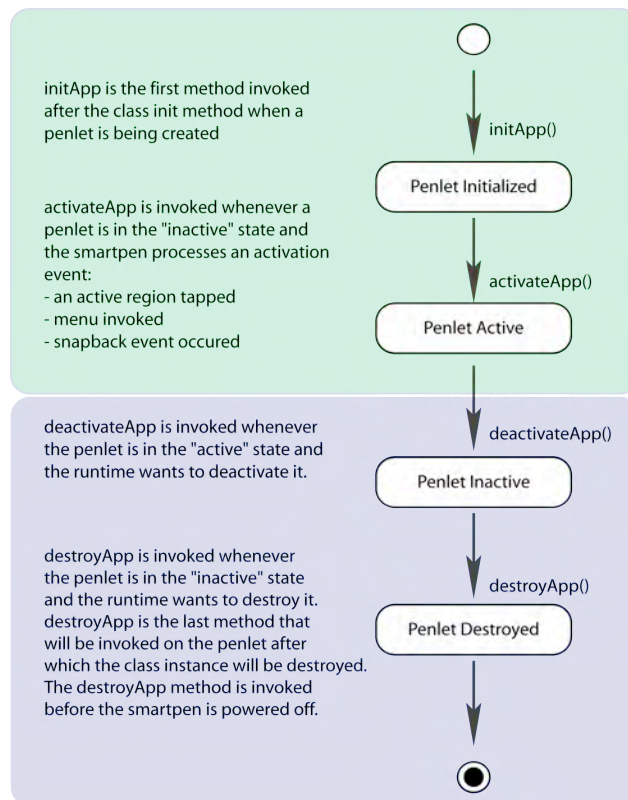
destroyApp()

Penlet Destroyed

Developer Tasks in each Life Cycle Method

Following is a list of the life cycle methods and the tasks that developers should perform in each.

Constructor of the Main Penlet Class

Developers can create any application-specific resources that they may need in their penlet. These resources should relate to the specific domain of your penlet. If any of your objects are part of the penlet profile and require a `PenletContext`, you must instantiate them in the `initApp` method instead.

The Four Life Cycle Callbacks

There are four callback methods for which your `Penlet` subclass must provide the method bodies. At the appropriate moment in the penlet life cycle, the runtime system calls these methods, sometimes passing back a value to your penlet. The four life cycle callback methods are: `initApp`, `activateApp`, `deactivateApp`, and `destroyApp`.

The initApp method

Before calling `initApp`, the runtime creates a `PenletContext` object and makes it accessible from the `context` field of your `Penlet` subclass. A penlet has only one `penletContext` object, which encapsulates various runtime characteristics of the currently executing penlet.

Factory methods that you call on this `PenletContext` object instantiate and return references to key objects such as event listeners, the Handwriting Recognition engine context, and the collection of regions for the current page. We will return to these objects as we discuss the other life cycle callback methods.

The code in this method will be executed just once, so you should include tasks that need to be done only once—at the beginning of your penlet's life cycle.

The canProcessOpenPaperEvents Method

If your penlet works on Livescribe Open Paper, you must override the default behavior of `canProcessOpenPaperEvents` method. By default, the method returns false. To enable your penlet to work on Open Paper, override this method and return true. If you neglect this step, every time your user taps on Open Paper, the runtime system will deactivate your penlet and switch to Notes Mode. If you intend on supporting Open Paper

functionality and do not override this method, your penlet will be unusable.

The activateApp Method

The runtime system calls `activateApp` immediately after `initApp` *and* whenever the penlet becomes active after having been inactive. You should use this method to restore application state that the newly activated penlet needs to run.

The runtime passes the reason for the activation to the `activateApp` method. Constants identifying the various activation reasons are defined in the `com.livescribe.penlet.Penlet` class. The ones of immediate concern for the new developer are:

`Penlet.ACTIVATED_BY_MENU`	This event is sent to the penlet when the user launches the penlet via the Main Menu.
`Penlet.ACTIVATED_BY_EVENT`	This event is sent to the penlet when the user taps on an area.

Your penlet can call `eventId` on the `Event` passed to this method by the system and compare it to these constants.

Symmetrical Method Calls

One approach to coding this method is to make it symmetrical with the `deactivateApp` method. Do tasks here that you will undo in `deactivateApp`. For instance, this is the recommended time to do the following:

- Add your event listeners, by calling the appropriate `add*Listener` methods on the `PenletContext` object. In your first penlets, these will probably include:

- `addMenuEventListener` is required for your penlet to display an application menu and receive events when the user selects a menu item.

- `addPenTipListener` is required for your penlet to receive such events as `PenDown`.

- `addStrokeListener` to receive strokes from the system as the user draws and writes on dot paper.

When adding any of these event listeners, pass in the object implementing the corresponding listener interface. For example, you will need the object that implements the `MenuEventListener` to pass to `addMenuEventListener`. In simple penlets that have one monolithic class, you will pass a reference to your `Penlet` subclass, using the *this* reference.

Keep in mind that you will "undo" these actions in the `deactivateApp` by calling `remove*Listener` methods.

Asymmetrical Method Calls

On the other hand, some method calls in `activateApp` do not have a symmetrical call in `deactivateApp`. For example, you may decide to get a `Display` object here so that you can display the application menu. To do this, call `getDisplay` on the `PenletContext` object. However, there is no corresponding call to "release" this object.

The deactivateApp Method

When a user switches the smartpen from one penlet to another, the smartpen runtime system calls `deactivateApp` and possibly `destroyApp` as well.

Because your penlet may be destroyed when it is inactive, you should store any relevant application data in the body of deactivateApp and/or destroyApp to preserve it for future use. As a general rule:

- Any states and data structures that can be easily and quickly reconstructed should be freed when deactivateApp is called.

- Other states and data structures, such as serializations, should be stored by destroyApp.

You can use the `PropertyCollection` class or the `PenletStorage` class for these purposes. A `PropertyCollection` object allows the developer to set and get developer-defined properties that the system saves on the Livescribe smartpen. The `PenletStorage` class provides developers access to storage on the Livescribe smartpen file system.

The runtime passes a constant to this method that describes the reason for the deactivation. Your code can test for the reasons that matter to

your penlet. In your first penlets, it is usually okay not to worry about these constants and provide a single response to your penlet being deactivated. As you develop more expertise in penlet development, you may be interested in some of the following deactivation reasons:

- `Penlet.DEACTIVATED_BY_SHUTDOWN`

- `Penlet.DEACTIVATED_BY_APPSWITCH`

- `Penlet.DEACTIVATED_BY_NOTESMODE`

You should release resources that you have retrieved during the execution of your penlet so that the memory they occupy can be used by the next activated penlet. There are some symmetrical calls in this method which undo some of the calls you made in `activateApp`. Some of the event listeners fall into this category.

Asymmetrical Calls in deactivateApp

There are some calls you make in `deactivateApp` that do not have a symmetrical call in `activateApp`.

For instance, consider the Handwriting Recognition engine. It uses many resources to do the work of analyzing user strokes and producing a word. You should release these resources in the `deactivateApp` method. There are special methods to do this, including calling `clearStrokes` and `dispose` on the engine context object. However, the instantiation of the engine probably did not take place in the `activateApp` method. Because HWR resources are relatively large, you may have decided not to instantiate the HWR engine until some user event occasioned the need for it.

The destroyApp Method

The smartpen may call `destroyApp` *whenever the penlet is in an inactive state.* For example, deactivation occurs when switching from one penlet to another. The runtime system may also destroy the penlet at this time. When destroyed, your penlet's initialized state will be lost and its Java classes unloaded. You should store any relevant application data in the body of `deactivateApp` and/or `destroyApp` to preserve it for future activation. See The deactivateApp Method.

Livescribe Platform Java API

Penlets are Java applications that use the Livescribe Platform Java API (Application Programming Interface). It is based on the Java Platform, Micro Edition (Java ME) and CLDC (Connected Limited Device Configuration). The platform is a Java ME stack for with the following components:

MMAPI 1.2 (JSR 135)
Livescribe Smartpen Profile
CLDC 1.1 (JSR 139)

The Smartpen Profile, created by Livescribe, consists of the following packages:

com.livescribe.afp	com.sun.cldc.io
com.livescribe.configuration	com.sun.cldc.io.j2me.socket
com.livescribe.display	com.sun.cldc.isolate
com.livescribe.event	com.sun.cldc.util
com.livescribe.geom	com.sun.cldc.util.j2me
com.livescribe.icr	com.sun.cldchi.io
com.livescribe.penlet	com.sun.cldchi.jvm
com.livescribe.storage	com.sun.cldchi.test
com.livescribe.ui	java.io
com.livescribe.util	java.lang
com.livescribe.buttons	java.lang.ref
com.livescribe.i18n	java.util
com.livescribe.io	javax.microedition.io
com.livescribe.quickCommand	javax.microedition.media
com.sun.cldc.i18n	javax.microedition.media.control
com.sun.cldc.i18n.j2me	javax.microedition.media.protocol
com.sun.cldc.i18n.uclc	

In addition to the classes always available in the runtime profile, there are a number of extra classes that can be used by any penlet. These classes are added to the jar for a penlet automatically by the Livescribe

build system if they are used by the penlet. The extra classes are contained in the following packages:

`com.livescribe.ext.plugins`
`com.livescribe.ext.ui`
`com.livescribe.ext.util`

For technical details on packages, their classes and methods, please consult, the Livescribe Javadoc in the SDK.

Handling Smartpen Events

Penlets are event-driven applications. Events can be actions that occur in response to a user manipulating the Livescribe smartpen on dot paper. Or events can be various system activities that are of interest to the penlet.

The smartpen system accepts input from the hardware or from the user and notifies the penlet by calling appropriate event handlers. In some cases, such as menu actions and system-generated activities, these events are encapsulated as descendants of the `Event` class. Thus, there are menu event objects and system event objects, which the system passes to the event handler. In other cases, there are no `Event` subclasses to encapsulate the events. The system calls the appropriate event handler and passes in other useful parameters instead, such as `Region` objects and time stamps of user strokes.

The major kinds of user actions and system activities include:

Menu Events	Sent by the system when the user taps on a Nav Plus symbol.
Strokes	Sent by the system when the user draws or writes on Livescribe dot paper. A stroke is the unbroken line (or "curve") traced between the time the user puts the smartpen tip down on dot paper and the time they pick it up.
Pen Down Events	Sent by the system when the user puts the Livescribe smartpen tip down on Livescribe paper.
HWR Results	Sent by the system when the Handwriting Recognition (HWR) engine has an intermediate or final result (i.e., a string containing the word or pattern that the engine produced as its best analysis of the word the user just wrote).
Area Entering and Exiting Notification	Sent by the system when a user is in the middle of creating a stroke with the Livescribe smartpen on dot paper and the stroke enters or leaves a region.
System Events and Hardware Events	Sent by the system to communicate system occurrences (such as the user just muted the speaker) or hardware state (battery

	level, available flash memory for storage, etc.)

Creating Event Listeners

The developer implements the appropriate listener interface, providing functionality for the event-handling method(s). Then, the developer adds the listener to the `PenletContext` object. In simple penlets, the `Penlet` subclass generally implements the various listeners. In more complex penlets, a particular listener may be implemented by a helper or manager class.

Although there are a variety of listener interfaces and the event handler methods they specify, the most important event handlers for beginning penlet developers are:

handleMenuEvent in MenuEventListener

The `handleMenuEvent` method in `com.livescribe.events.MenuEventListener` is called by the system when the user taps up, down, right, left or center on a Nav Plus. The system passes a `MenuEvent` object to the method. You can check what kind of menu event occurred by calling the `getID` method on the `MenuEvent` object and comparing the return value with the constants defined in the `MenuEvent` class: MENU_UP, MENU_DOWN, MENU_RIGHT, MENU_LEFT, and MENU_SELECT.

To select the currently visible menu item in the OLED display, the user clicks on the right arrow of the Nav Plus. Thus, developers should pay close attention that they handle the MENU_RIGHT events appropriately. Note that, despite the name, MENU_SELECT events have nothing to do with selecting menu items. Instead, they are sent when the user taps the center of a Nav Plus.

penDown in PenTipListener

The `penDown` method in `PenTipListener` is called by the system when the user places the smartpen tip down on dot paper. You will be handling this event even in your first penlets. There is no `Event` class that encapsulates this user action. Instead, the system passes in the following parameters:

`time`	This is a long integer value that indicates when the pen down occurred.
`region`	This is a `Region` object that identifies the unique rectangle on a page of dot paper on which the pen down occurred. If there are no regions where the pen down occurred, the system sets the `regionId` to 0.
`pageInstance`	A `PageInstance` object that identifies the particular page of dot paper on which the pen down occurred. The system instantiates all `PageInstance` objects for you and passes the appropriate one to certain event handlers in which you might need the page instance. For example, when coding the `penDown` or `strokeCreated` event handlers, you would need the current `PageInstance` object in order to create a `StrokeStorage` object and retrieve individual strokes made by the user.

There are other event handlers in `PenTipListener` that the basic penlets leave as no-ops, including:

- `penUp`
- `doubleTap`
- `singleTap`

These can be very useful in more sophisticated penlets.

strokeCreated in StrokeListener

The `strokeCreated` method in `StrokeListener` is called by the system when the user completes a stroke on dot paper. There is a `Stroke` class to encapsulate strokes, although the `strokeCreated` does not pass a `Stroke` object to this method. Instead, it passes the same parameters as does the `penDown` event handler.

Note that when a user creates a stroke, the system calls `penDown`, `penUp`, and `strokeCreated`. Since `strokeCreated` encompasses both a pen up and a pen down, you should think carefully about what happens when a stroke occurs on an existing region.

For instance, your region may be designed for tapping. But users will sometimes jerk the smartpen slightly when attempting a tap. This movement will probably cause a stroke event to be sent to your penlet. In that case, you should implement the `strokeCreated` method the same as the `penDown` method.

Regions and Areas

Regions

An active region is an active expanse of contiguous dots on Livescribe dot paper. A user can tap on a region and get a response from the penlet that owns the region. If the penlet is not running, the runtime system will launch it. If the penlet is just deactivated, the system will activate it.

An active region is often simply referred to as a region. *Active region* and *region* are synonymous.

Static Regions

A static region is a region specified by the developer in the paper product definition (called an AFD). The dot paper usually has a printed graphic to indicate the location, shape, and usage of a static region. For example, the Paper Replay control bar at the bottom of each page in a Livescribe notebook is a group of static regions. Static regions are sometimes called *Fixed Print regions*.

Dynamic Regions

A dynamic region is a region created during run time when a user taps on unclaimed dot space. (Unclaimed dot space is known as *Open Paper*.) The penlet creates dynamic regions that encompass the written input and can be tapped on to trigger behavior in the penlet. For example, in Piano, the user creates dynamic regions when drawing piano keys and rhythm and instrument buttons. In Paper Replay, the user creates dynamic regions as the user takes notes during a recording. Later, the user taps on a note and the associated point in the audio starts to play.

Overlapping Regions

Regions may overlap. In that case, the z-order of regions determines in which order the events are delivered. The region with highest z-order receives the events first. Then the region with next highest z-order. And so on. If a region has the occlusive bit set, then regions with lower z-orders do not receive events.

Developers assign z-order to static regions during penlet development. Dynamic regions often receive a z-order at run time such that the most recently drawn region is "on top"—that is gets a higher z-order than older regions. However, this behavior is up to the developer to implement and is not required.

Areas

While a region is a physical entity on dot paper, an *area* is a logical concept. An area defines functionality that should occur when a user taps on (otherwise interacts with) a region. Best practice dictates that each area trigger only one such action.

The developer assigns an area to each region. The same area can be assigned to multiple regions. For example, all the Record buttons in Paper Replay control bars have the same area, because they all perform the same action.

Region Ids and Area Ids

A Region Id is an internal 64-bit number that uniquely identifies a region to a smartpen. The Region Id encodes: Area Id, Instance Id, Occlusiveness, and Z-order, among other things.

An Area Id is a 16-bit positive integer that is a subset of the Region Id. An Area Id must be unique within a penlet, but one Area Id can be assigned to multiple regions. All regions that are owned by the same penlet and have the same Area Id will have the same functionality.

Developers are responsible for assigning Area Ids to their regions. Static regions are assigned Area Ids when the paper product is defined. Dynamic regions are usually assigned Area Ids in event handlers such as `penDown`.

The system reserves Area Id of 0 to denote Open Paper—that is, dot space that has not been claimed by a region. Thus, a developer starts assigning Area Ids at 1.

Another way of thinking of an *area* is as a collection of regions that have the same area Id.

Associating a Penlet and a Region

A region must be associated with the penlet that should be activated when a user taps on the region. The Instance Id identifies the penlet thus associated. Since some penlets may be instantiated multiple times, each running instance of a penlet has its own Instance Id. For instance, the Piano application is instantiated separately every time a user draws a new piano. Thus, several Piano instances may exist simultaneously.

An instance Id is a 16-bit positive integer and is encoded as part of the Region Id.

Dynamic Regions and Instance Ids

Dynamic regions are assigned an Instance Id by the system when the region is created. The system encodes the Instance Id in the Region Id.

Static Regions and Instance Ids

Static regions are assigned an Instance Id in a slightly more complex way.

Application Class Name and Application Id

Each penlet is uniquely identified by its Java class name, such as `com.livescribe.paperreplay`. Such class names, however, can be rather unwieldy, so a more efficient identifier has been devised.

The Application Id is a 16-bit positive integer that the developer must create and assign to the penlet. The AFD for a paper product maps these Application Ids to Application Class Names.

When defining static regions, the developer associates an Application Id with one or more static regions. The Application Id is coded into the Region Id at development time. Only static regions have Application Ids.

Application Ids and Instance Ids.

At runtime, a user taps a region. Here's how the smartpen system responds:

1. Reads the Application Id.

2. Looks up the Application Class Name in the AFD for the paper product.

3. Looks up the Instance Id for that Application Class Name.

4. Modifies the Region Id, replacing the Application Id with the Instance Id.

5. Activates (or launches and activates) the appropriate instance of the appropriate penlet.

The developer deals with these identifiers at different points in his development process. When defining a region in your paper product, you associate an Application Id and Application Class Name with the region. When your code accesses the Region Id in an event handler, however, only the Instance Id is retrievable.

An Example

To clarify things a little, let's consider an example. Assume a smartpen has two applications, Paper Replay and Timer, with the following Instance Ids at run time. (The Instance Ids likely vary from one smartpen to another.)

Application Class Name	Instance Id
com.livescribe.paperreplay	10
com.livescribe.timer	11

During application development, our developers mapped the following Application Ids and Application Class Names in the AFDs for the Livescribe notebooks. They could have chosen any numbers for the Application Ids, as long as each was unique within an AFD.

Application Class Name	Application Id
com.livescribe.paperreplay	2
com.livescribe.timer	1

A static region, like the Stop button for Paper Replay, has Area Id = 4, which is a global value defined by Livescribe for standard controls.

Looking at the second table, you can see that region must have an Application Id = 2. At run time, the Static Region is tapped by the user; the event thrown will have Area Id= 4 and Instance Id = 10 (In the first table, you can see that the Instance Id for Paper Replay is 10.)

For more details on associating Application Ids with penlets and assigning them to static regions, please read *Developing Paper Products.*

Accessing Standard Livescribe Controls

Some functionality provided by the smartpen system and by the bundled applications (such as Paper Replay) is accessible from within your penlet. You access this functionality by using the Standard Livescribe Controls in your paper product.

Livescribe publishes a list of the standard *Area Id* for each such control, as well as the *Application Class Name* of the associated system module or bundled application. Use of Standard Livescribe Controls does not require Application Ids.

For details, please read "Standard Livescribe Controls" in *Developing Paper Products.*

Uniqueness of a Region ID

You may be wondering how Regions Ids can be kept unique in the following situation:

- Two regions are defined in the same page of a paper product.

- The regions are associated with the same penlet.

- The regions are not distinguished by a z-order value. (They do not occupy the same dots on the same page.)

In this case, the smartpen system ensures that each `Region` object has a unique region ID. It uses the z-order value for this purpose, since the z-order is otherwise unused. Normally, this is all a matter of system "bookkeeping" and need not concern you.

Working with Static Regions

If your application has static regions, it also has a paper product. The best tool for developing paper products is the Livescribe Paper Designer. For a detailed discussion of paper creation, including static regions, please see *Developing Paper Products.*

For information on handling events, see Handling Smartpen Events in this manual.

Working with Dynamic Regions

This section discusses how to create a dynamic region and how to respond to user taps on a region

Creating a Dynamic Region

There are three main steps to creating a dynamic region.

1. Get a bounding box (which is a Rectangle object). See Get a Bounding Box.

2. Assign an Area Id and create a Region object. See New Dynamic Region: Assigning Area Id and Adding Region to Collection.

3. Attach the Region to the RegionCollection for that page. See New Dynamic Region: Assigning Area Id and Adding Region to Collection.

Get a Bounding Box

The location of a region on dot paper is defined by a bounding box. Whatever actual shape the user writes or draws the resulting bounding box is a `Rectangle` object. There are three ways to get a bounding box for user input.

- Get a bounding box for a stroke from the ICR engine.

- Get a bounding box for a stroke from a `StrokeStorage` object.

- Get a bounding box for a group of strokes, using the `StrokeStorage` and `Stroke` classes.

The ICR Engine: When you are using the ICR engine to analyze user writing, the engine determines a bounding box that encompasses the user's written word. You can get that bounding box by calling the `getTextBoundingBox` method on the `ICRContext` object. The method returns a `Rectangle` object, which is the bounding box containing the written word. You normally make that call in the Handwriting Recognition event handler `hwrUserPause`.

The StrokeStorage Class: When you are not using the ICR engine to analyze user writing, you must instantiate the `StrokeStorage` class. A `StrokeStorage` object contains all strokes that meet two conditions: (1) the strokes were made on the current page of dot paper and (2) the strokes belong to the current penlet.

The `StrokeStorage` object has a `getStrokeBoundingBox` method that returns a `Rectangle` object representing the stroke's bounding box. You normally call that method in the `strokeCreated` event handler.

Getting Bounding Box for Several Strokes: The bounding box returned by `getStrokeBoundingBox` is the smallest rectangle in which the stroke will fit. If you wish to create a bounding box that encompasses several strokes, proceed in the following manner: Get a bounding box for the current stroke. Then get the next stroke and get its bounding box. Create the union of those two bounding boxes. Continue until you reach the last stroke. You normally determine the union of several bounding boxes in the `strokeCreated` event handler.

The following code snippet captures written strokes and creates a `Shape` union which represents the smallest possible rectangle containing all strokes:

```
public void strokeCreated(long startTime, Region areaID
                          PageInstance pageInstance) {

   this.currentStroke = strokeStorage.getStroke(startTime);

   // Initialize the container Shape
   if (null == this.currentRect) {
      this.currentRect = this.currentStroke.getBoundingBox();
   }
```

```
        // Add the stroke to the container Shape
    else {
        this.currentRect = Shape.getUnion(this.currentRect, this.currentStroke);
        this.currentRect = this.currentRect.getBoundingBox();
    }
}
```

Both the `Stroke` and `Rectangle` classes extend the `Shape` class. All shapes have `getBoundingBox` and `getUnion` methods. You can learn about these classes in the Javadoc for the `com.livescribe.geom` package.

New Dynamic Region: Assigning Area Id and Adding Region to Collection

The Area ID of a region determines how a penlet responds to a user tap. When the user taps on Open Paper (unclaimed dot space), the area ID that the system passes to `strokeCreated` is 0. The developer must create an area ID and assign it to the new dynamic region. That dot space is now claimed by your penlet and will have the Area Id you specified. Regions that have the same area ID will have the same behavior.

A `RegionCollection` object contains all the regions belonging to the current page of dot paper. When a user creates a stroke on Open Paper, the event handler must create a new `Region` for the current stroke and add the `Region` to the `RegionCollection`. The region is now "active." When the user taps on that region in the future, the penlet will respond as designed by the developer.

The following code snippet from the Translator Demo sample project demonstrates creating a unique Area ID for a new dynamic area and adding a new `Region` to the `RegionCollection`. Note that you must pass in the bounding box of the region when calling the `addRegion` method.

```
Rectangle rect = this.hwrEngine.getTextBoundingBox();
RegionCollection rc = this.context.getCurrentRegionCollection();
int wordAID = getAreaIdForWord(...);
addDynamicArea(rect, workAID,ac);


private static void addDynamicArea(Rectangle rect, int aid,
                                   RegionCollection rc) {
    int centerX = (rect.getX() + (rect.getWidth() >> 1));
    int centerY = (rect.getY() + (rect.getHeight() >> 1));
    int areaId = (aid & AREA_ID_MASK)
                | ((centerY & AREA_ID_CENTER_Y_MASK)
                    << AREA_ID_BITS);
```

```
    Region regionID=new Region(areaId, centerX, false, false);
    rc.addRegion(rect, regionID, false);
}
```

Responding to User Taps on Regions

A penlet responds to a user tap on a region by implementing the `penDown` event handler. The system passes the following parameters: `time`, `region`, and `pageInstance`. These are a `long` value representing the time stamp of the tap, the `Region` object in which it occurred, and the `PageInstance` object representing the current page of dot paper.

Developers generally implement the `penDown` event handler in the following manner:

1. Check if the `penDown` occurred on Open Paper. If so, simply return. The `strokeCreated` event handler should create the region for a stroke on Open Paper.

2. A developer can determine if the event is on Open Paper by calling the `getInstance` method on the `Region` object passed to `penDown`. The instance ID is a unique integer created by the system to manage penlets that are installed on a smartpen. If the value is 0, no penlet owns that dot space; it is Open Paper.

3. If `getInstance` returns a non-zero value, the current penlet owns the region on which the `penDown` occurred. The developer retrieves the area ID and calls appropriate functionality. Often, the response is a sound on the smartpen speaker or a display on the smartpen OLED.

    Note: If functionality for a region should be activated when the user either draws into it *or* taps on it, rather than *just* when the user taps on it, you should use the `regionEnter` event rather than the `penDown` event.

Displaying on the Smartpen OLED

The penlet can display on the Livescribe smartpen OLED in the following ways:

Application Menu and RIGHT_MENU Items

Many penlets have an application menu that displays available items, one item at a time. The user scrolls through the menu by tapping the up and down arrows on the Nav Plus. When the desired menu item appears, the user selects it by tapping on the right arrow. The penlet then responds with a submenu, a sound and display, or other functionality.

The responsibilities for implementing the application menu are shared by the system and the penlet. The system handles display of the current menu item and display transitions (the "upward/downward scrolling effect") from item to item. The developer codes the movement of the focus through the application menu and, of course, the response to a `MENU_RIGHT` event.

Creating an Application Menu: Developers enable an application menu for their penlets as follows:

1. Implement the `BrowseList.Item` interface as a static member class of the `Penlet` subclass.

2. Instantiate that static member class, once for each item in the application menu.

3. Instantiate the `BrowseList` class, passing in a vector of `BrowseList.Item` objects.

4. Call the `setCurrent` method on the current `Display` object, passing the `BrowseList` object as a parameter.

Moving Focus Through the Application Menu: The system displays the menu items, but the developer must handle moving the current focus to items in the `BrowseList` object in response to `MENU_UP` and `MENU_DOWN` events.

1. Call the `focusToNext` or `focusToPrevious` methods, as appropriate.

2. Call the `setCurrent` method on the current `Display` object, passing the `BrowseList` object as a parameter. (Not required, if current `Display` is already set to the `BrowseList` object.)

Handling MENU_RIGHT Event: When the application menu is displaying in the smartpen OLED, the user can tap the right arrow of a Nav Plus. The penlet must handle the MENU_RIGHT event. One response is to play a sound and display text to the smartpen OLED.

Displaying in Response to a User Tap on a Region

When a user taps on a region, the `regionEnter`, `penDown`, `penUp`, `singleTap`, `doubleTap`, and `regionExit` event handlers are called by the system. Generally, all penlets implement `penDown`. The other handlers are implemented fully or as no-ops, according to the design of the penlet developer. Many useful penlets handle user taps on regions by implementing `penDown` or `regionEnter` only.

One response to a user tap is to display text and/or images to the smartpen OLED. Such "tap and display" functionality is very common in a penlet. The code looks like this:

The following code snippet sets a `ScrollLabel` as the current `Displayable` and draws the specified text to the `Display`.

```
if (this.display.getCurrent() != this.labelUI) {
    this.display.setCurrent(this.labelUI);
}
this.labelUI.draw(textToDraw, true);
```

The current display will remain on the smartpen OLED until the penlet calls `setCurrent` again (or the system switches to another penlet in response to user actions.)

Displaying a Message to the User

Penlets also use a `ScrollLabel` object whenever they need to display a message to the user, whether in response to a user tap or not. The calls are identical to the preceding section.

Displaying Text or Image or Both

The `draw` method of the `ScrollLabel` supports the display of text, image, or both. Penlets call the appropriate version of the overloaded `draw` method.

```
void draw(java.lang.String text, Image img, boolean scroll)
void draw(Image img, java.lang.String text, boolean scroll)
```

You can read more about the `ScrollLabel` class in the Javadoc for `com.livescribe.ui` package.

Playing Sounds

Sounds that a penlet plays are resources packed in the penlet's JAR file. If you place the sound files in your penlet project in the `res\audio\` folder, the Ant build system will automatically put them in the JAR at that same path. The supported file formats are WAV and WavPak.

The following code snippet initializes a `MediaPlayer` object and plays an audio resource that is specified via usage of the I18N Resource Bundle.

```
MediaPlayer mediaPlayerUI;
. . .
this.mediaPlayerUI = MediaPlayer.newInstance(this);
. . .
String audioFile=bundle.getSoundResource(I18NResources.ID_SOUND_WRITE_WORD);
this.mediaPlayerUI.play(audioFile);
```

Using Bitmap Images

Like sounds, small bitmap images are resources that are stored in the penlet's JAR file. To access these resources at runtime, there is no special method in the Livescribe Smartpen Java API. Instead, you should use standard Java APIs to get a resource as a stream, as demonstrated in this code snippet:

```
Class myPenletClass = this.getClass()
myPenletClass.getResourceAsStream("images/myImage.arw")
```

Note: The above example is not localizable. A localizable equivalent is:

```
ImageResource helloworldImage =
context.getResourceBundle().getImageResource("helloworld");
```

Bitmaps for display on the smartpen OLED are small. Developers should verify that their bitmaps are discernable and communicate effectively with the user. The dimensions of the OLED are as follows:

Full Dimensions of Smartpen OLED Display

Height	18 pixels
Width	96 pixels

The penlet does not always have the full dimensions of the smartpen OLED display at its disposable. At certain system thresholds, the system uses a small portion at the right of the display to show the System Tray.

System Tray Dimensions

Height	18 pixels
Width	6 pixels

Converting to ARW Format

The smartpen uses image files with an ARW extension, which indicates a simple 1-bit file format designed for the smartpen. This format has nothing to do with the Sony image format that uses the same extension. The Livescribe Platform SDK's build system provides a way for developers to convert images to ARW.

1. Create bitmap images in the BMP, GIF, JPEG, JPG or PNG format.

2. Convert the images to ARW by simply placing them in the `src/images` folder in the penlet project.

3. The images will be automatically converted to ARW and put in the JAR in the `res/images/` folder.

Note: You can convert an image to ARW format *manually* by using the Livescribe Image Converter in the Eclipse IDE with Livescribe plugins:

1. In the Package Explorer, select the top node of your penlet project.

2. Select Livescribe > Penlet Configuration Editor.

3. In the Image Resources tab, click Add.

4. Click **Browse** to select your image file.

5. Click OK to **Add** the converted image.

    The image will be converted to ARW format and stored in the
    `res\images` folder of your penlet project.

If you will be localizing your penlet, see <u>Converting Localized Images to
ARW</u>.

Using and Converting Audio Formats

The Livescribe Platform supports three audio playback formats natively:

*   Microsoft WAV

*   WavPack

You should choose one of these formats based on audio quality, playback
features supported, and storage requirements. For all formats, only mono
and stereo are supported. For WAV and WavPack formats, the bit depth
must be 16-bit.

Sampling Rate

The sampling rate at which the smartpen plays audio files is 16 kHz.
Consequently, 16 kHz is the ideal sampling rate to use when creating
audio. Higher sampling rates are generally usable; however, they should
be avoided whenever possible because:

*   They require additional CPU cycles to play.

*   The smartpen's resampling algorithm does not provide high quality
    for "down-sampling", since it is designed for "up-sampling."

*   They are a waste of storage space on the smartpen.

Sampling rates lower than 16 kHz are allowed, but they result in a
tradeoff: audio files occupy less storage space, but have lower sound
quality.

Bitrate

A very important measurement of any audio file is *bitrate*. This refers to the amount of data consumed by the file each second. It is generally measured in bits per second (bps) or kilobits per second (kbps).

Gaplessness

Audio formats can be either gapless or not gapless. A format is gapless if audio playback can blend seamlessly from one clip to another (or have a clip loop back to its own beginning in a seamless manner). WAV and WavPack are inherently gapless because there is a one-to-one correspondence between input and output samples.

Summary of Supported Audio Formats

The following table describes characteristics for each of the audio formats supported by the Livescribe Platform.

Format	Ext	Min Mono Bitrate	HQ Mono Bitrate	Min Stereo Bitrate	HQ Stereo Bitrate	Gapless?	CPU Usage	License
MS WAV	.WAV	256 kbps	256 kbps	512 kbps	512 kbps	yes	low	free
WavPack	.WV	36 kbps	56 kbps	72 kbps	96 kbps	yes	medium	free

WAV Format

As the Microsoft audio standard, WAV is probably the most common audio format in the world. Although the WAV container supports compressed formats (commonly, ADPCM) WAV files intended for the smartpen must be 16-bit uncompressed PCM (either mono or stereo). This format provides perfect quality (within the limits of our 16 kHz sampling rate), sample-accurate seeking, gapless playback, and minimal use of the CPU during encoding. However, being uncompressed, it is very wasteful of the flash storage space on the smartpen and should only

be used when absolutely needed or when the clips are of very short duration.

Generating Files in WAV Format

Nearly all audio editing programs can generate WAV files compatible with the Livescribe Platform. Simply bear in mind the requirements: 16 kHz, 16-bit, stereo or mono.

WavPack Format

The WavPack format is an open-source audio codec that provides both lossless and lossy compression of WAV files. Like WAV, WavPack is sample-accurate and gapless. The Livescribe Platform does not currently support WavPack seeking.

The lossless mode allows WavPack to store the exact audio data provided by the WAV files, but in about half the space. The lossy mode is similar to the industry-standard AAC, but much simpler. To provide the same quality as AAC, WavPack requires about a 1/3 higher bitrate, but also uses fewer CPU cycle(s) for both decoding and encoding, because all processing is done in the time domain.

Generating Files in WavPack Format

You have two options for creating WavPack files: generate them using a WavPack-aware audio editor or convert WAV files to WavPack, using the WavPack tool in Livescribe Platform SDK.

Some audio editing programs support WavPack natively (such as Reaper and Traverso). There are WavPack plugins for the popular Adobe audio editor Audition (which also works with CoolEdit) and Steinberg's WaveLab.

Converting WAV to WavPack

If your audio creation program does not export to Wav Pack, you can simply export to WAV. Then you can convert to WavPack using:

- The Livescribe Penlet Configuration Editor, accessed by selecting Livescribe > Configuration Editor in your Eclipse with Livescribe plugins IDE.

or

- The command-line WavPack encoding tool called `wavpack.exe`. It is found in the *SDKInstallDir*`\Resources\penletsdk\bin\win32` folder, where *SDKInstallDir* is the directory to which you unzipped the Platform SDK. Source files must be 16-bit, 16 kHz files.

Lossless WavPack Files

To create lossless WavPack files, use the following syntax at the Windows command-line:

```
wavpack filename.wav –x6
```

The destination file automatically receives the same name as the source file, but with the .WV extension.

Lossy WavPack Files

To create lossy WavPack files, use the following syntax at the Windows command-line:

```
wavpack filename.wav –x6 –bxx
```

where xx is the desired bitrate in kbps. For example, to generate a high-quality stereo file, use –b96.

Configuring Penlets

You can configure properties and resource files for your penlets using the Penlet Configuration Editor. You can also configure your penlets to support different locales using the Penlet Configuration Editor. For more information, see Localizing Penlet Properties.

To open the Penlet Configuration Editor, select your penlet project in the Project Explorer and right-click to open a list of available menu items. Select Penlet Configuration Editor to open it. The editor has tabbed

views for Properties, Image Resources, Audio Resources, Text Resources, and Advanced Settings.
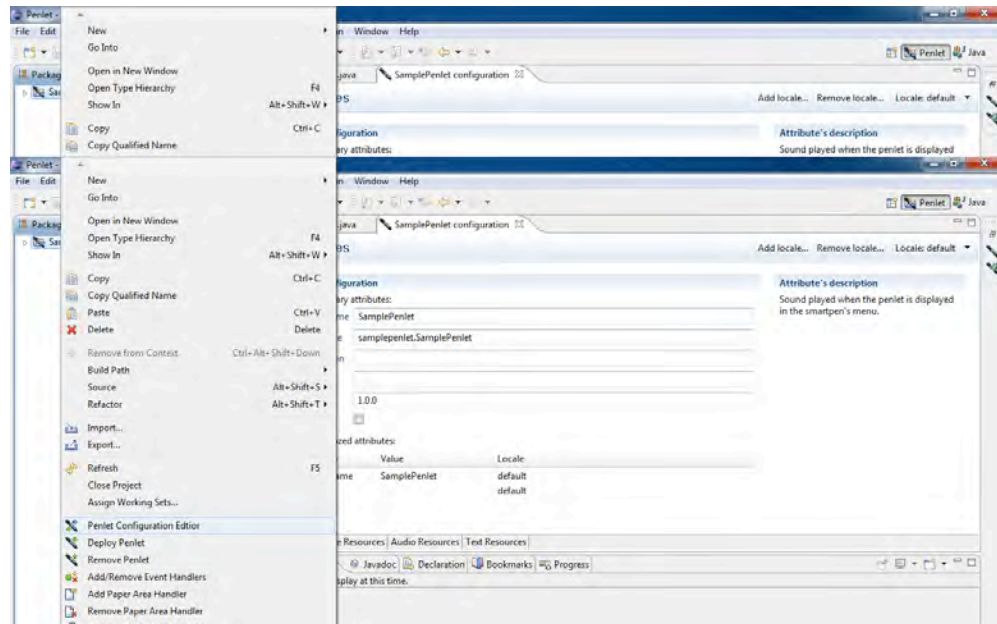
Penlet Properties

To define basic properties of your penlet, click the Properties tab in the Penlet Configuration Editor. Here you can change the menu name, launch sound, penlet name, penlet classname, free-form description of the penlet, Livescribe Store application category, penlet version, and toggle the penlet to be locked (removable) or not. To change a value, select its text and edit it in place.
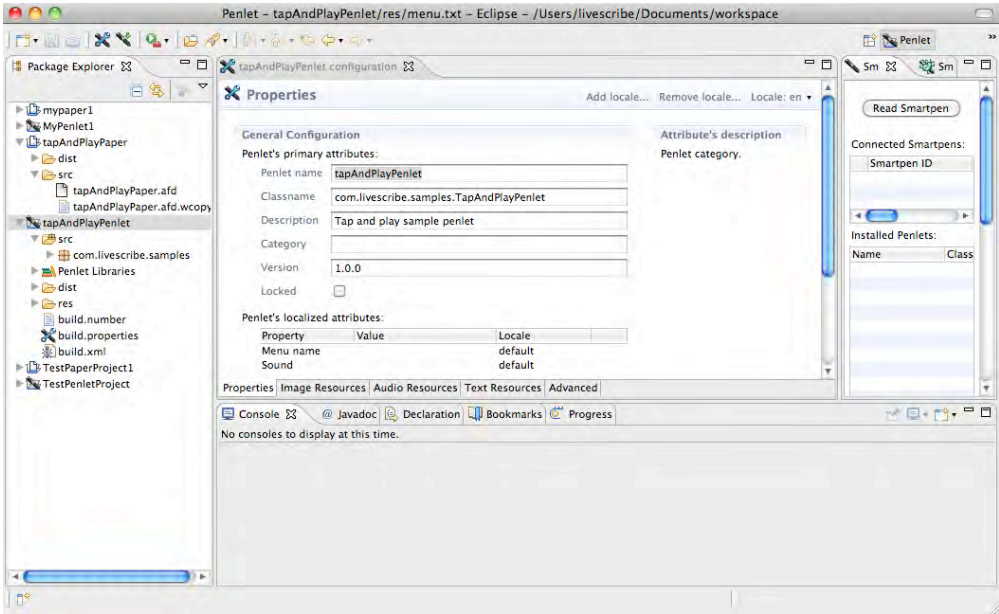
When you edit these properties using the editor, they are written to either the build.properties or the menu.txt configuration file, as appropriate. The penlet build process uses these files to create the penlet JAR.

In the penlet's source code project, the menu.txt file lives in the res folder. It is a plain text file in which you enter properties using name=value syntax. The properties include:

appclassname	Required. Name of the application class name	appclassname=Livescribe.foo.Foo
Type	Required. Set to APP	type=APP

Name	Required. Name of the penlet as it should appear in the Main Menu.	name=Foo
sound	Optional. Name of sound file that is played when your penlet's name rolls into view on Main Menu. This file should also be listed in the resources property.	sound=NP_Foo.wav

Image Resources

To define image resources for your penlet, click the Image Resources tab in the Penlet Configuration Editor. Use the Add, Rename, or Remove buttons to locate and assign your images.
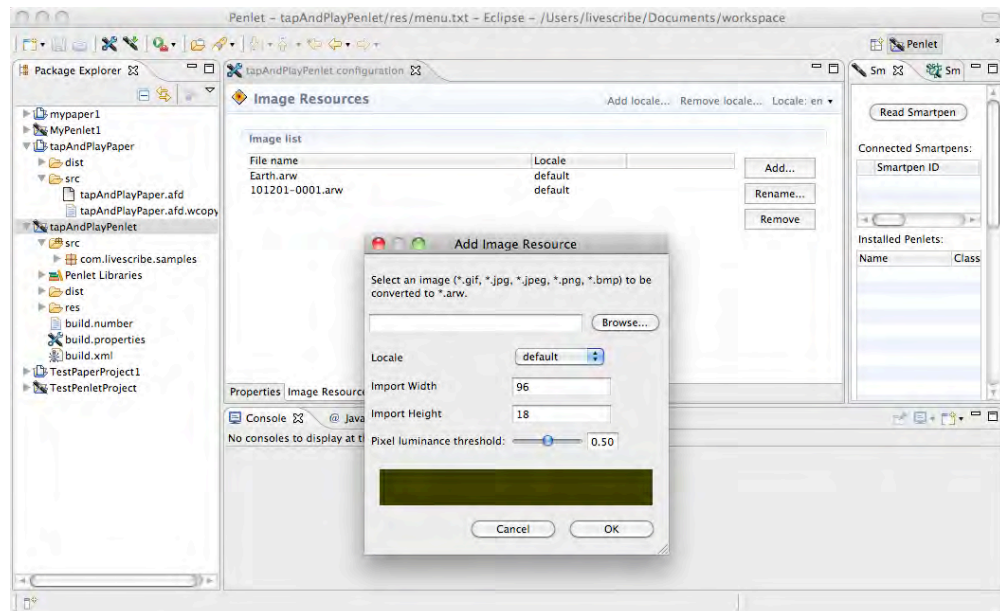
Audio Resources

To define audio resources for your penlet, click the Audio Resources tab in the Penlet Configuration Editor. Use the Add, Rename, or Remove buttons to locate and assign your audio files.

Internationalized audio files are stored in a directory path including `audio` and then the locale name. In the sample `res` tree, you can find English audio files at `res\audio\en_US` and French audio files at `res\audio\fr_FR`.
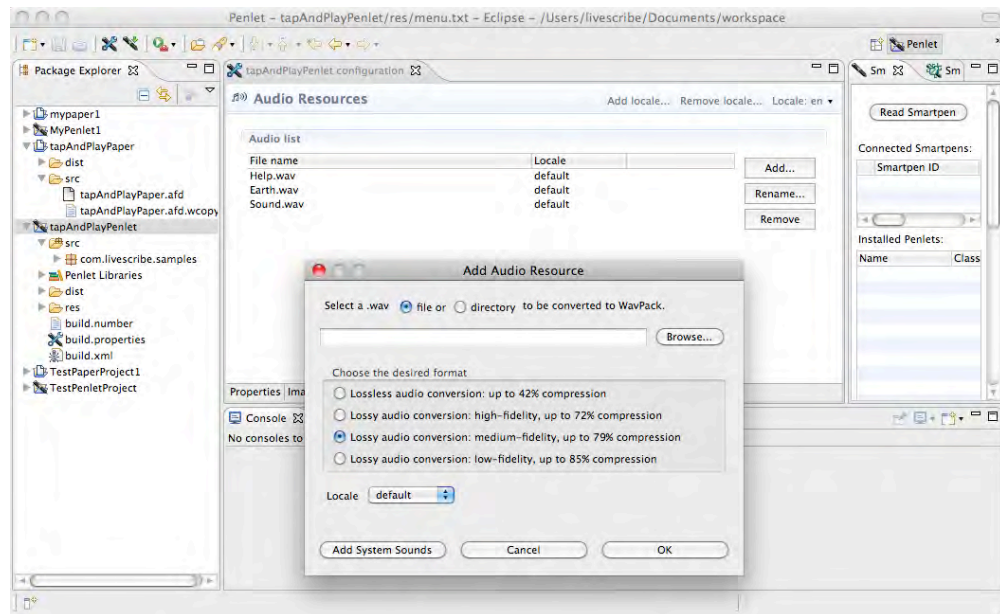
Text Resources

To define text resources for your penlet, click the Text Resources tab in the Penlet Configuration Editor. Use the Add or Remove buttons to create key/value pairs of text strings for your penlet.
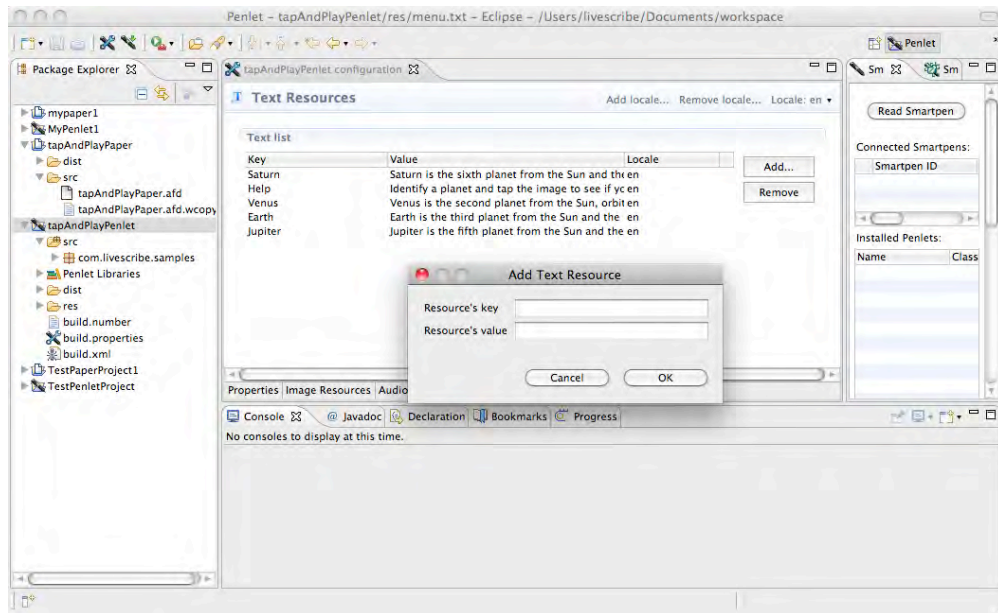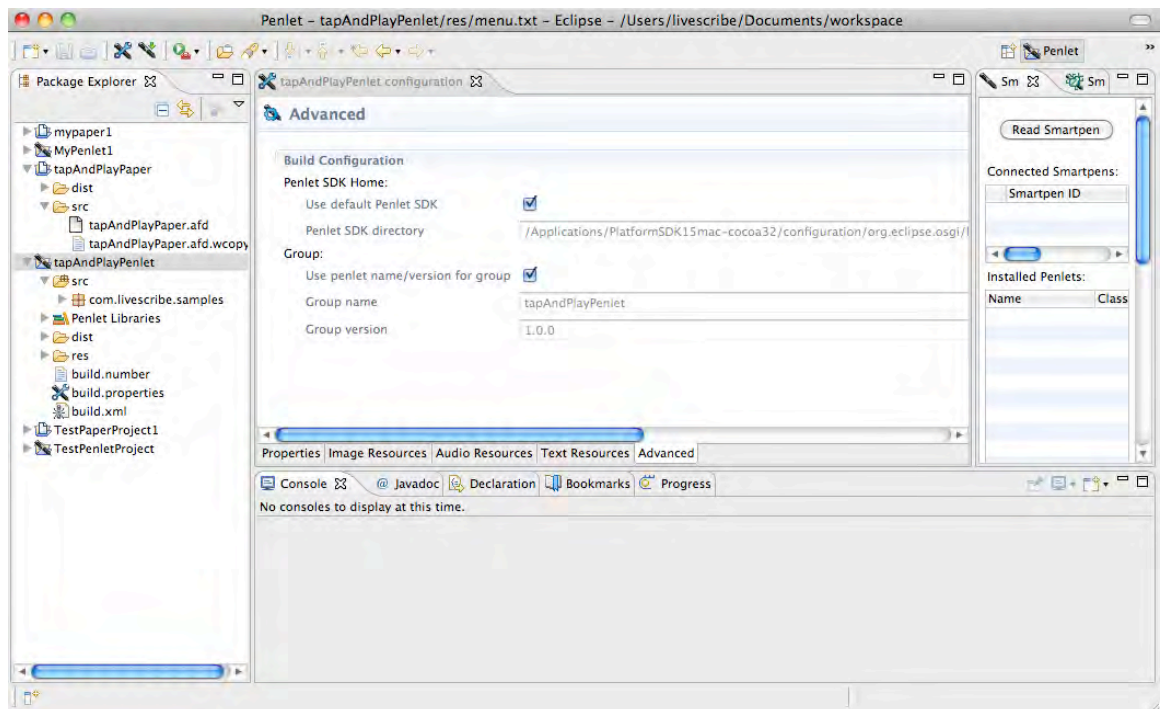
Advanced Settings

Use the Advanced tab to define the penlet group name and version used when deploying applications to end users. You can also configure the Penlet SDK Home folder to set up custom build environments.

About config.txt

The config.txt configuration file specifies application properties that can be read by the penlet at run time. Application properties are similar to resources in that their values are externally specifiable, but they differ in that they are not localizable. The properties in config.txt are not currently writable at run time.

In the penlet's source code project, the config.txt lives in the res folder. It is a plain text file in which you enter properties using name=value syntax. Examples are:

foo=bar
foo2=123

You can access your penlet's configuration properties from your code by calling getAppConfiguration on the PenletContext object. The method returns a Config object. To retrieve a property, call one of the following methods on that Config object:

- getStringValue
- getBooleanValue
- getDoubleValue
- getLongValue

Note: Since property values are specified in the config.txt file as strings, you need to know what type each value should be and call the appropriate method. Here is a code snippet:

Config config = context.getAppConfiguration();
String value = config.getStringValue("foo");
String value = config.getLong("foo2");

Saving Data to the Smartpen

You can save runtime data from your penlet to the smartpen in two ways:

- Serializing data using the `PropertyCollection` class.
- Saving data directly to the file system of the Livescribe smartpen.

Serializing via the PropertyCollection Class

The `PropertyCollection` class allows you to create properties at runtime. The properties for your penlet are stored in a properties file on the smartpen's file system. Since you can set property values to any Java object, this is a convenient way to achieve object serialization and deserialization in your penlet. (Of course, very large objects might degrade your penlet's performance.)

The steps to use a property collection are:

1. Call the static method `PropertyCollection.getInstance`, passing in the `PenletContext` object.
2. Set properties by calling the `setProperty` method.
3. Get properties by calling the `getProperty` method.

You can read more about the `PropertyCollection` class in the `com.livescribe.afp` package.

Saving to the Smartpen File System

You can save data to the file system of the smartpen by using the `com.livescribe.storage` package.

Internationalization

Livescribe smartpens support different locales, allowing the user to select from a pre-determined set. To configure your penlet for multiple locales, follow these guidelines.

1. In your source code, place internationalized resources in the correct directory trees underneath the res directory.

2. Access these resources using methods in the com.livescribe.i18n package.

Following is a sample res directory tree. This example shows a locale-specific directory fr_FR, as well as a general language directory fr. This directory is useful for content that should be the same for all countries using a given language. For example, if you want text displayed the same in fr_FR, fr_CA, you would put it in the fr language directory.

The full search order for the example of fr_FR is:

1. fr_FR
2. fr
3. en_US (the default locale for the smartpen)
4. en
5. default locale (resource root directory - /res, /res/audio, or /res/images depending on the resource type.)

```
res
  └─── en_US
              ├───── messages.properties
              ├───── menu.text
          fr
              ├───── messages.properties
              ├───── menu.text
       fr_FR
              ├───── messages.properties
              ├───── menu.text
     audio
              ├───── en_US
                           ├───── welcome.wav
                           ├───── goodbye.wav
                        fr_FR
                           ├───── welcome.wav
                           ├───── goodbye.wav
     image
              └───── en_US
                           ├───── image.arw
                           ├───── anotherImage.arw
                        fr_FR
                           ├───── image.arw
                           └───── anotherImage.arw
```
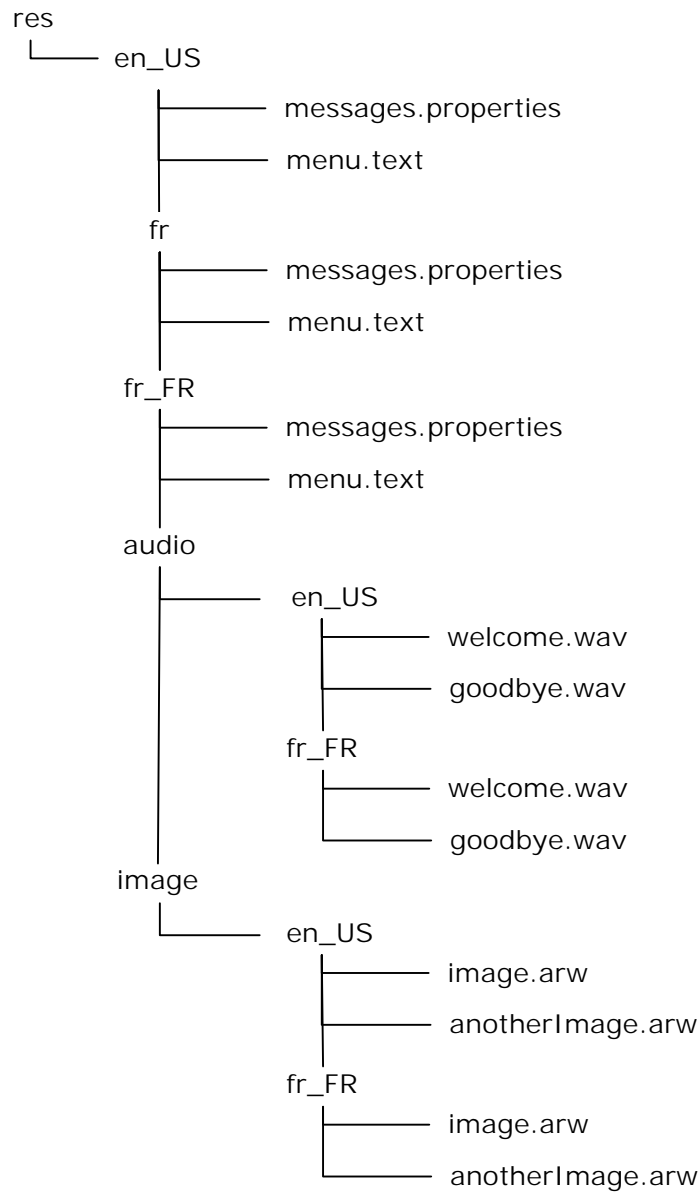
Configuring Penlets for Different Locales

You can configure your penlets to support different locales using the Penlet Configuration Editor.

To open this editor, select your penlet project in the Project Explorer and right-click to open a list of available menu items. Select Penlet

Configuration Editor to open it. The editor has tabbed views for Properties, Image Resources, Audio Resources, and Text Resources.
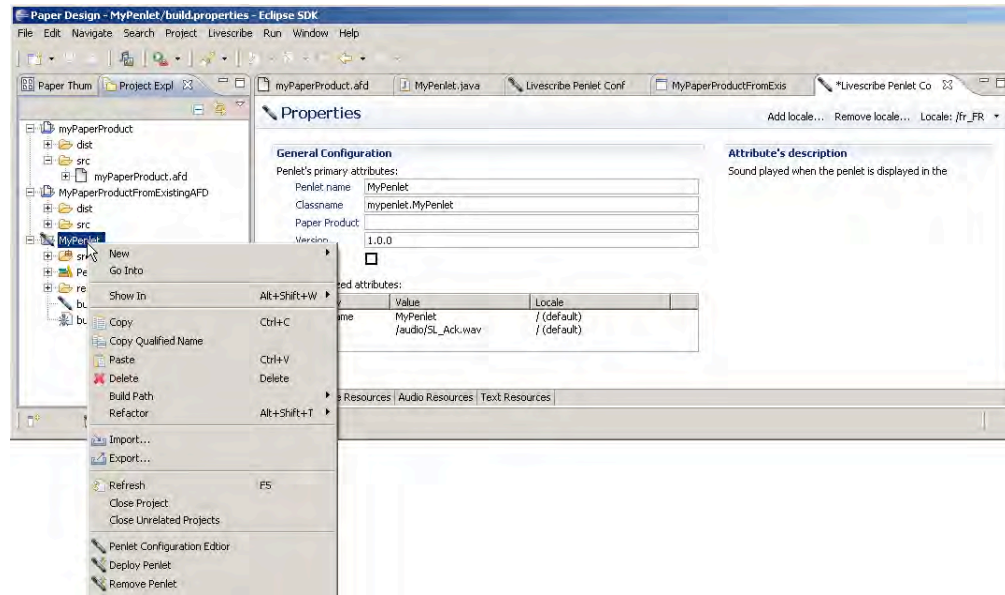
To create a new locale for your penlet, click the Add locale button on any of the configuration editor tabs. Type in one of the supported locales from the list below, and click Okay.

English - US	en_US (default)
French	fr_FR
German	de_DE
Italian	it_IT
Spanish	es_ES
Korean	ko_KR
Simplified Chinese	zh_CN

You can also remove one or more locales by clicking the Remove locale button, choosing one or more of the locales, and clicking Okay.

After creating a locale, you can configure your penlet's properties, images, audio, and text strings for that locale. To do this, click the Locale: drop-down on any of the configuration editor tabs and choose

the locale to be configured. Then enter the locale-specific information, file
paths, and content.

Localizing Penlet Properties

To define basic properties of your penlet, click the Properties tab in the
Penlet Configuration Editor. Here you can change the menu name, launch
sound, penlet name, penlet classname, associated paper product,
version, and toggle the penlet to be locked (removable) or not.

To change a localized attribute:

1. Click the Locale: drop-down menu and select the desired locale.

   For example, you could select the `fr_FR` locale.

2. In the Penlet's localized attributes table, click the cell under
   Value that you wish to change.

   An insertion point will appear and you can edit the value in place.
   For example, you can change the Menu name from HelloWorld to
   BonjourMonde.

3. Click the column head Value.

   The new locale (if you changed it) and the new value will appear in
   the table.
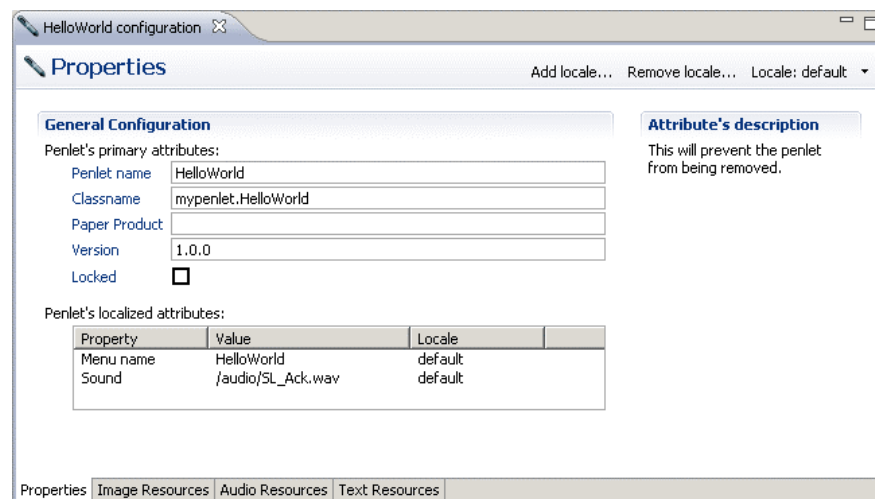
Localizing Image Resources

To define image resources for your penlet, click the Image Resources tab in the Penlet Configuration Editor. Use the Add or Remove buttons to locate and assign your localized images.
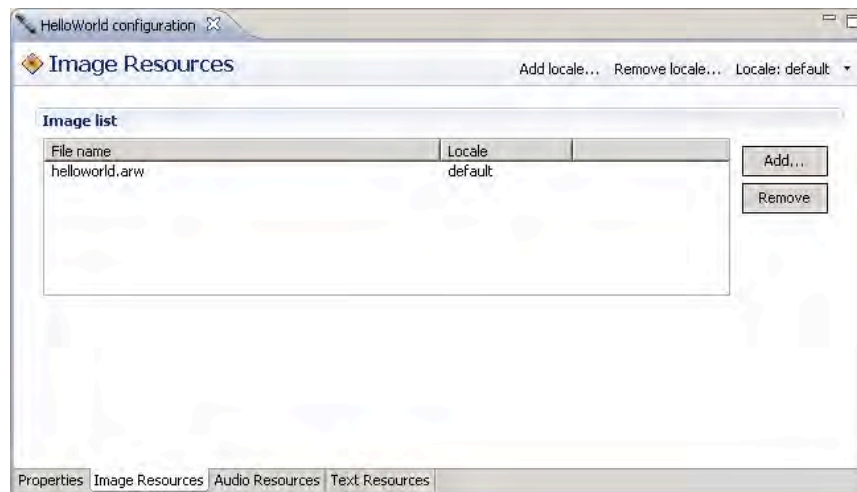
Internationalized image files are stored in a directory path including `images` and then the locale name. In the sample `res` tree, you can find English image files at `res\images\en_US` and French image files at `res\images\fr_FR`.

Using Internationalized Image Resources

Using internationalized audio resources in your penlet involves these steps:

1. Place your internationalized image resources in the appropriate directory under `res\images`.

2. Get a reference to the `ResourceBundle` object. You'll need the penlet context for this step:

    ```
    ResourceBundle bundle = this.context.getResourceBundle()
    ```

3. Call the `getImageResource` method on the `ResourceBundle` object.

4. Call the `getImage` method on the `ImageResource` object returned by the previous call.

Converting Localized Images to ARW

To create localized ARW images for your localized penlet project, do the following:

1. Create the localized images in BMP, GIF, JPEG, JPG or PNG format with a depth of 1 bit. (Each pixel is either on or off.)

2. Restrict your image size to 96 x 18 pixels, which is the size of the smartpen display. For more details on size requirements for images, see [Using Bitmap Images](#).

3. Place the images in `src\images\<locale name>`. Thus, place English BMP images in `src\images\en_US` and French BMP images in `src\images\fr_FR`.

4. When you build your penlet project, the images will be automatically converted and placed in the JAR at the following path: `res\images\en_US` or `res\images\fr_FR`, etc.

Note: You can convert an image to ARW format *manually* by using the Livescribe Image Converter in the Eclipse IDE with Livescribe plugins:

1. In the Package Explorer, select the top node of your penlet project.

2. Select Livescribe > Penlet Configuration Editor.

3. On the Image Resources tab, click Add.

4. Click **Browse** to select your image file.

5. Click OK to **Add** the converted image.

   The image will be converted to ARW format and stored in the `res\images` folder of your penlet project. For example, move an English image to : `res\images\en_US or` a French image to `res\images\fr_FR`.

Localizing Audio Resources

To define audio `resources for your` penlet, click the Aud**io** **`Resources`** `tab` in the Penlet Configuration Editor. Use the Add or Remove buttons to locate and assign your localized audio files.

Internationalized audio files are stored in a directory path including `audio` and then the locale name. In the sample `res` tree, you can find English audio files at `res\audio\en_US` and French audio files at `res\audio\fr_FR`.

Using Internationalized Audio Resources

Using internationalized audio resources in your penlet involves these steps:

1. Place your internationalized audio resources in the appropriate locale directory under `res\audio`.
2. Get a reference to the ResourceBundle object. You'll need the penlet context for this step:
    ResourceBundle bundle = this.context.getResourceBundle()
3. Call the getSoundResource method on the ResourceBundle object.
4. Call the getInputStream method on the SoundResource object returned by the previous call.

Localizing Text Resources

To define text resources for your penlet, click the Text Resources tab in the Penlet Configuration Editor. Use the Add or Remove buttons to create key/value pairs of text strings for your penlet's locales.
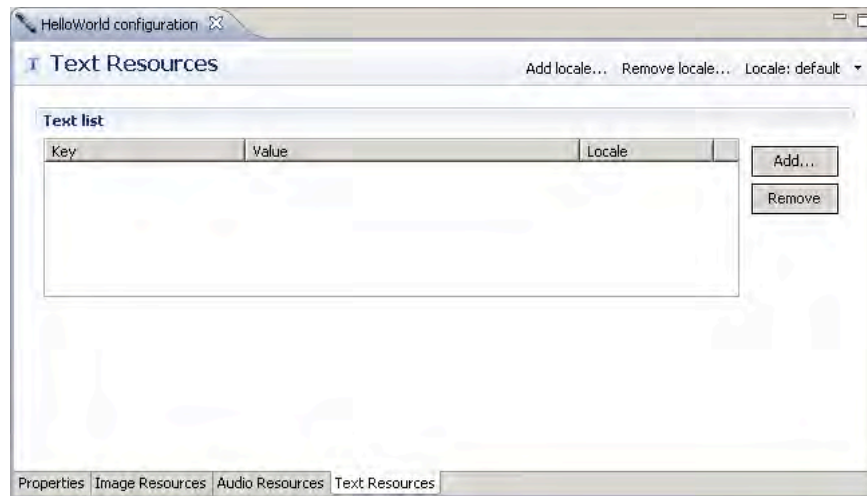
Note: Livescribe recommends that you add all keys to be localized in a default locale, such as the resource root. This location is identified as "default" in the config editor. Either en_US or en can also serve as the default locale. Localize the keys you want to have different values in each locale that you want to support.

Internationalized strings are stored in a file called `message.properties`, which is located in a directory bearing the name of a locale. Thus, in the sample `res` tree, English strings are in `res\en_US\message.properties` and French strings are in `res\fr_FR\message.properties`.

Using Internationalized Text Resources

Using internationalized text resources in your penlet involves these steps:

1. Enter the strings in the `message.properties` file in the following format: `property.name=value`

2. Get a reference to the `ResourceBundle` object. You'll need the penlet context for this step.

3. Call the `getTextResource` method on the `ResourceBundle` object, passing the property name of the desired string.

4. Call the `getText` method on the `TextResource` object returned by the previous call.

In the above description of the property name format:

- By "`property.name`," we mean the dotted property names common to Ant and Java. For example, you might use `greeting.formal` as the name of a property that holds the text of a formal greeting (such as "Good morning, sir or madam."). And you might use `greeting.friendly` as the name of a property that holds the text of a friendly greeting (such as "Hi, there.").

- By "value," we mean the string in the language you wish displayed on the smartpen. In the above example: "Good morning, sir" or "Bonjour, monsieur" And "Hi, there!" or "Salut!".

Here's a very basic code snippet:

```
ResourceBundle bundle = this.context.getResourceBundle();
String strGreeting = bundle.getTextResource(greeting.formal).getText()
```

Assigning Property Names to Constants

In accordance with good coding practice, you'll probably decide to assign the property names to constants and then pass the constant to `getText` call. In that case, you would define constants such as:

```
public static final String I18N_STR_GREET_ADULTS = "greeting.formal"
public static final String I18N_STR_GREET_KIDS = "greeting.friendly"
```

And the code would look like this instead:

```
ResourceBundle bundle = this.context.getResourceBundle();
String strGreeting = bundle.getTextResource(I18N_STR_GREET_FORMAL).getText()
```

Handwriting Recognition

The Livescribe Platform provides support for recognizing users' handwriting. As a user writes, a handwriting recognition engine embedded in the smartpen firmware analyzes strokes and delivers characters, words, and phrases to the current smartpen application.

Paper-Based Input Recognition

Computer recognition of paper-based input is a fascinating and complicated field, full of acronyms such as OCR, HWR, and ICR. This section briefly describes the differences in these technologies as they apply to the Livescribe platform. If you are content with knowing that the Livescribe smartpen performs true online handwriting recognition, you may skip this section. Otherwise, please read on.

ICR and HWR

HWR comes in two varieties, depending on whether written characters are considered "offline" or "online" data. We will start with the official names of these technologies, and then say a word about their informal use.

ICR (Intelligent Character Recognition) is a technology that analyzes *offline* data. It recognizes hand-printed characters that have been captured by a scanner or camera. ICR can take into account the shapes and proximity of characters, but has little other information to aid its interpretation effort. The Livescribe smartpen does not use ICR, in the strict sense of the term.

Online HWR analyzes *online* written input. It incorporates some ICR techniques, but solves an additional challenge—input in real time. For Livescribe, Online HWR is a real-time technology that accepts strokes from digital pens and determines the characters being written. Recognition of written input on the Livescribe platform is true Online HWR.

For Livescribe, strokes are paths traced on dot paper by the smartpen. The paths are captured and stored as a series of points (x,y pairs) in 2-dimensional space, ordered by time. A single handwritten character is

made up of one or more strokes. A stroke starts when the pen tip touches down on the paper and ends when the pen tip lifts up again. Stroke characteristics include:

- Stroke's start time

- Stroke's beginning and end points

- Sequence of points—sampled at equal intervals—which compose the stroke's path.

The engine receives strokes made by the smartpen on dot paper, and delivers digital text that a penlet can use.

Used Interchangeably

As often happens, the sharp distinctions between *ICR* and *HWR* are rarely observed in everyday usage. ICR is the older technology, and, indeed, some ICR techniques persist in HWR. For these reasons, the term *ICR* is sometimes applied to any analysis of written input, whether online or offline. Thus, the smartpen's handwriting recognition engine is called the *ICR engine*, and the Java API for handwriting recognition is found in the `com.livescribe.icr` package.

In the Livescribe API and documentation, you can treat *ICR* and *HWR* as synonyms.

Digital Text and Digital Ink

With the foregoing explanation in mind, we can clarify two other terms that you may see in the Livescribe platform: *digital text* vs. *digital ink*.

Digital text refers to electronic sequences of characters that are digitally-encoded according to an industry standard, such as ASCII. The ICR engine takes written input and delivers digital text to the current penlet.

Digital ink is a term that refers collectively to the strokes captured by a digital smartpen. Livescribe handwriting recognition engine can translate digital ink into digital text. Not all penlets use HWR.

Overview of Handwriting Recognition Process

The handwriting recognition process can be described as follows:

- Once the digital ink is captured by the smartpen and sent to the recognition engine, the recognition cycle begins.

- The recognition engine uses files called resources that give it information about alphabets, segmentation, character subsets, language contents and so on.

- The recognition result is returned to the smartpen application as digital text.

The handwriting recognition engine used by the Livescribe Platform is the central element in the recognition process: it uses powerful handwriting recognition algorithms combined with information about each language being recognized to yield optimal recognition results.

Tuning for Performance

Your smartpen applications will achieve the best performance if you follow these guidelines:

- Limit the number of files that your application creates on the file system. This number should not exceed 500 files.

- Limit the number of files in the application's JAR. This number should not exceed 1000 files. The JAR includes class files and resource files. Audio resource files, in particular, can become quite numerous.

If your application needs more than 1000 individual audio resources, you should consider bundling them into a few, larger files that will reside either in the JAR or on the file system. You can then access individual audio resources by calling API methods that enable direct seeking into the files.

Note: Files in the JAR must NOT be stored with zip compression.

Sample Translator

Many developers learn most effectively from working samples. This section examines the Sample Translator penlet, an Open Paper penlet. It exercises the fundamental functionality of the Livescribe Java API.

This overview will cover the highlights only. If you wish more detail, please consult the Javadoc in the Livescribe Platform SDK.

Please browse to `livescribe.com` and log on to the Developer Forums. Download the `com.livescribe.samples.translator` project from the Developer Forum titled Sample Code & Docs & FAQs. Open the `SampleTranslator.java` file and follow along as you read. To further assist you, the source code is amply commented throughout.

We approach the topic of programming penlets in two stages:

1. Sample Translator: User's Perspective will describe the Translator penlet in operation, from a user's perspective. Since the Livescribe Livescribe smartpen is a new platform and its input and output models not yet widely known, we will spend a little time examining the user features of this penlet.

2. Sample Translator: Developer's Perspective looks at the source code for Sample Translator. It very briefly describes the methods that provide the domain-specific functionality of Sample Translator (that is, how this penlet matches an English source word to the written word and spokeju4c 9in audio in the target language). We then jump into the heart of the sample, examining how it exercises the Livescribe Java API.

Sample Translator: User's Perspective

The Sample Translator penlet translates individual words from English to one of four target languages: Spanish, Swedish, Mandarin, and Arabic. The equivalent word appears on the OLED display of the Livescribe smartpen at the same time an audio clip plays, demonstrating the pronunciation of the word by a native speaker.

Following is a quick examination of the Sample Translator penlet from the user's perspective. You may wish to follow along on your smartpen by building the `com.livescribe.samples.translator` project and deploying (installing) on your

Livescribe smartpen. For information on using the Livescribe IDE, consult the manual titled *Getting Started with the Livescribe Platform SDK*.

Launching the Sample Translator Penlet

The user launches the application from the smartpen's Main Menu as follows:

1. The user taps on any Nav Plus on Livescribe dot paper, which launches the smartpen's Main Menu. The words "Main Menu" appear on the Livescribe smartpen display and the corresponding audio plays.

2. The user taps on the down arrow of Nav Plus to view the list of available applications, displayed one at a time.

3. When "Sample Translator" appears on the Livescribe smartpen display, the user taps on the right arrow of Nav Plus to launch the Translator application.

Translating a Source Word

1. As soon as the application starts, a menu list of available target languages appears on the Livescribe smartpen display. This is the application menu for Translator.
    Note: Not every penlet will have an application menu. Upon starting, some penlets will display a message or communicate with the user by playing audio.

2. The user taps the up or down arrow of any Nav Plus until the desired language appears.
    On the Livescribe smartpen display, there may be a small triangle at the upper, lower, or right edge. These triangles mean that a user can "move" in that direction. When the up triangle is visible, the user can tap the up arrow on a Nav Plus to see another menu item in the browse list. A down triangle indicates that the user can tap the down arrow. A right triangle means that the user can select the current menu item and an action will occur. Usually, this will have one of the following results:

    *   Some text and/or an image will display.

    *   An audio will play.

    *   Display and audio will occur simultaneously.

3. The user taps on the Nav Plus right arrow to start the corresponding translation engine. The smartpen produces a message, instructing the user to write a word. This message is multimedia: the text displays and the appropriate audio plays.

4. On a page of Livescribe Open Paper, the user writes a word from the list of English source words.

5. The translation engine finds the match in the target language, and the Livescribe smartpen displays the target word and plays an audio recording of the word, pronounced by a native speaker.

   The format of the display is: `source word – target word`

   For instance, if the user writes the word `one` while the application is in Spanish mode, the display shows: `one – uno`. For Spanish and Swedish, the output consists of Latin characters. In the case of Arabic and Mandarin, the target word is displayed as an image that represents the appropriate non-Latin characters.

Tapping a Previously Written Word

After users write word on the page, they can later return to the word and tap it. The correct translation will be displayed and the accompanying audio played. The application will use the current target language when performing a translation.

Consider our previous example in which the user wrote `one`:

1. The user returns to the application's menu list of target languages and selects Arabic.

2. The user finds the same page of dot paper and clicks the word `one` that they wrote earlier.

3. This time the Arabic word will be displayed and the Arabic audio played.

4. The user taps on a different word or write a new word.

5. Please remember that this is a sample application, and the word list is short, consisting of the numbers 0-10, hello, goodbye, please, thank you, thanks, coffee, chocolate, banana, beer, and water. If the user writes a word *not* on that list, the application displays a message: "Try writing:" followed by the list of English source words.

Returning to Application Menu List

When users get the "Write or Tap a Word" message, they can tap the up, left, down arrow or center of a Nav Plus. This action causes the application menu list to re-appear. In the case of Sample Translator, that is the target language menu list. Similarly, if a translation is still visible on the display and the user taps up, left, down or center on Nav Plus, the target language menu list re-appears.

Sample Translator: Developer's Perspective

Now that we have examined the Sample Translator application from the end-user's perspective, you are in a better position to understand individual classes and method calls. It's time to look at the actual code.

Domain-Specific Code

The purpose of this penlet is to translate individual English words to one of four target languages: Spanish, Swedish, Arabic, and Mandarin. The domain-specific logic in this penlet is not our primary focus. A few words of explanation, however, will help you orient yourself.

The code creates an array called `ENGLISH_SOURCE_WORDS` to hold the words that the user can write and the penlet will translate. It also creates an array of target words (i.e., translations) for Spanish and Swedish. The target display for Mandarin and Arabic is a little more complicated, since they do not use the Latin alphabet. The characters for those words are stored as images. All four target languages have audio resources that contain the pronunciation of the target word by a native speaker.

The set up work for the translation "lists" depends on the `createDictionary` and `createImages` methods:

- The `createDictionary` method creates a hash table for each target language: the English source word is the key and the target word (a string or an image, as appropriate) is the value.

- The `createImages` method locates each image resource from the `/images/ directory` in the penlet JAR, gets a stream, and creates an Image object. It

then returns an array of these images for the target language (Mandarin or Arabic).

The `createEngines` method calls both the `createDictionary` and `createImages` methods to implement a translation "engine" for each target language. In the case of Spanish and Swedish, the images are not required, so `createEngines` just retrieves the target words from the appropriate target word array.

To get everything started, the `SampleTranslator` constructor calls `createEngines`.

User Writes a Word

When the user writes a word, the code (ultimately) calls `processText`, which in turn calls:

- `showTranslation`
- `addDynamicArea`

The `showTranslation` method, as you might guess, displays the English word, a dash, and the translated target word. It also plays the pronunciation audio file. The `addDynamicArea` method creates a region for the word and adds it to the region collection for the page of dot paper that the user tapped on.

User Taps a Written Word

When users tap on a word they wrote, the region for that word already exists. The code calls `processDynamicAreaId`, which retrieves the English source word and the target word or image and then calls `showTranslation`.

And that's it. Now we can look at how to implement the code that uses the Livescribe Smartpen Java API.

Constructor and Life Cycle

The `SampleTranslator` constructor and life cycle methods are pretty straightforward.

The constructor creates the translation engines, as already mentioned. It also creates a `ScrollLabel` object, as well as the `Vector` object that will hold the `BrowseList.Item` objects required for the application menu (that is, the scrolling list

of target languages). In addition, it initializes the `mode` field, which will be used by the Handwriting Recognition engine. The complete code of the constructor looks like this:

```
public SampleTranslator() {
    this.labelUI = new ScrollLabel();
    this.createEngines();
    this.vectorMenuItems = new Vector();
    for (int i = 0; i < ENGINES.length; i++) {
      this.vectorMenuItems.addElement(ENGINES[i]);
    }

    this.mediaPlayerUI = MediaPlayer.newInstance(this);
    this.mode = MODE_READY;
    this.setEngine((byte) ENGLISH_TO_SPANISH);
}
```

initApp method

The `initApp` method gets the `Display` object required for output to the smartpen OLED. It also adds a `PenTipListener` to handle `penDown` events. Since the `SampleTranslator` class implements the `PenTipListener` interface, we pass a *this* reference to the `PenTipListener`. The appropriate snippet is:

```
this.display = this.context.getDisplay();
    this.context.addPenTipListener(this);
```

Note: The other event handlers in the `PenTipListener` interface (`penUp`, `singleTap`, and `doubleTap`) are of no interest to this penlet and are implemented as no-ops.

activateApp

Of the various things that happen in `activateApp`, the most important to us are the following three:

```
this.menuView = new BrowseList(this.vectorMenuItems, null);
```

This line creates a `BrowseList` object, passing in the `vectorMenuItems` created in the constructor. The application menu of target languages is a `BrowseList` object. Each entry in the menu is a `BrowseList.Item` object.

```
this.display.setCurrent(this.menuView);
```

This line sets the current `Display` object to the `BrowseList` object called `menuView`, which is the application menu. The current `Display` object must be set, or no display to the OLED will occur. We will be resetting this object when we are finished with the application menu and wish to display other objects, such as a `ScrollLabel` object.

```
this.switchToMode(MODE_TEXT_INPUT);
```

This calls a non-API method—that is, one particular to this penlet and not part of the Livescribe Java API. The `switchToMode` method initializes the Handwriting Recognition engine for use in our penlet. It also registers the `StrokeListener`, passing in a reference to the `SampleTranslator` instance, which implements the `StrokeListener` interface.

deactivateApp

The `deactivateApp` consists of one line:

```
this.switchToMode(MODE_READY);
```

This is the same method we show in `activateApp`, except that this time it is called with the `MODE_READY` constant. Ultimately, this code clears the Handwriting Recognition engine of strokes, disposes of its associated resources, and sets the `ICRContext` to `null`. It also removes the `StrokeListener`.

Note the symmetry between `activateApp` and `deactivateApp`: first we add the `StrokeListener` and then we remove it. If the system switches away from the Sample Translator penlet and later switches back to it, the `activateApp` will be called again and the `StrokeListener` added.

destroyApp

This penlet does nothing in `destroyApp`. Large resources, such as the Handwriting Recognition engine resources, were already released in `deactivateApp`.

canProcessOpenPaperEvents

This method is inherited from the penlet class, where its implementation returns false. The `SampleTranslator` class overrides it and returns true. This is a simple but *essential* step for any penlet that wishes to receive events such as `penDown` when the user taps on Open Paper.

```
public boolean canProcessOpenPaperEvents() {
    return true;
}
```

Displaying a BrowseList

This penlet creates `BrowseList.Item` objects by implementing the `BrowseList.Item` interface in the static member class `Engine`. The Engine class serves a double purpose:

- It provides `getTargetLangContent` and `getTargetLangAudio` methods to return the appropriate target word and audio. They are called every time a source word must be translated.

- As implementer of `BrowseList.Item`, it also provides the language name (English, Spanish, Mandarin, and Arabic) and accompanying menu audio for each item on the application menu. It does this by implementing `getAudioMimeType`, `getAudioStream`, `getText`, and `isSelectable`.

The application menu is affected only by the `BrowseList.Item` section of the `Engine` class. Let's review the code involved:

- In its constructor, the penlet instantiates the `Engine` class once for each target language and assigns the engines to `vectorMenuItems`.

- In activateApp, it creates a `BrowseList` object, passes in the vector, and returns `menuView`.

- Also in `activateApp`, it calls `this.display.setCurrent(this.menuView)`

This last call hands the initialized `BrowseList` object to the system, which uses it to display the current item of the application menu. (We must, however, handle changing the focus of the `BrowseList` object, as you will see in the `handleMenuEvent` discussion in a moment.)

isSelectable

This method is specified in the `BrowseList.Item` interface. The `Engine` class implements it and returns true for each engine created, ensuring that on the OLED, that item in the menu application has a small triangle displayed to the right. The triangle means that the user can tap the right arrow of the Nav Plus and get a response from the penlet. In the case of Sample Translator, the response is the visual and audio message urging the user to write a word.

In the `menuHandleEvent` method, a penlet can call `isSelectable` when it receives a `MENU_RIGHT` event. If the return is false, the penlet can choose not to respond and let the system process the event. The Sample Translator penlet does not make this test, however, since it knows that each language on the application menu should have a triangle pointing to the right.

Displaying a ScrollLabel

`BrowseList` objects are not the only kind of display available. This penlet uses `ScrollLabel` objects to display messages to the user. For instance, the "Write a word" message and the results of the Handwriting Recognition engine are both displayed to the smartpen OLED by using a `ScrollLabel` object.

You may remember that in the penlet constructor, we created a `ScrollLabel` object called `LabelUI`. This object is used throughout the code for displaying messages and results. The required sequence of calls is:

- `labelUI.draw`

- `display.setCurrent`

You can see actual calls, for example, in the section of `handleMenuEvent` that begins with: `if (event == MenuEvent.MENU_RIGHT)`. The lines are:

```
this.labelUI.draw(INPUT_PROMPT, true);
this.display.setCurrent(this.labelUI);
```

Registering Listeners

Sample Translator implements four listeners: `StrokeListener`, `PenTipListener`, `HWRListener`, and `MenuEventListener`. Listeners fall into two categories when it comes to registration. Some must be registered with the `PenletContext` object and some do not. Of these listeners, the two that do not need to be registered are:

- `MenuEventListener`

- `HWRListener`

In this penlet, the listeners that must be registered and unregistered are:

- `StrokeListener`

- PenTipListener

This last group of listeners must be explicitly added by calling
`context.addStrokeListener(this)` and `context.addPenTipListener(this)`. The
this reference, of course, is the penlet class that implements the respective
interfaces. Also, these listeners should generally be unregistered by calling
`context.removeStrokeListener(this)` and
`context.removePenTipListener(this)`.

The Sample Translator source follows this model in the case of the `StrokeListener`,
which is added by `activateApp` and removed (indirectly) by `deactivateApp`. In both
cases, the `switchToMode` method is called directly, and it calls the appropriate add or
remove method. We could have done the same for the `PenTipListener`.

The Handwriting Recognition Engine

This penlet uses the Handwriting Recognition engine to analyze users' handwriting
and return a best-guess at the word written. It uses intelligent character recognition
to accomplish this feat. In fact, we use the terms HWR (Handwriting Recognition)
and ICR (Intelligent Character Recognition) interchangeably for the present version
of the Java API.

This code creates the ICR engine context when `activateApp` is called and destroys it
when `deactivateApp` is called. Let's look at the method that `activateApp` calls to
perform the initialization of the ICR engine: `switchToTextInputMode`. Note that this
is not an API call. The code gets an `ICRContext` object, specifying what timeout
determines the end of a word. When this time has passed with no more user input,
the ICR engine returns its best-guess to the penlet. A usable timeout is 1000
milliseconds.

The code then proceeds to add language and handwriting resources needed by the
ICR engine. These are part of the Livescribe Smartpen Java API., Finally, it registers
the `StrokeListener`.

```
private void switchToTextInputMode() {
    if (this.hwrEngine == null) {
// Obtain an ICR Engine with a (1) second user pause timeout
    this.hwrEngine = this.context.getICRContext(1000, this);
    this.hwrEngine.addResource(ICRContext.HPR_AK_DEFAULT);
    this.hwrEngine.addResource(ICRContext.SK_ALPHA);
    this.hwrEngine.addResource("lex_translatordemo.res");

    this.hwrEngine.addResource(ICRContext.LK_WORDLIST_30K);
    }
// Enable Penlet to obtain Stroke events from the system
```

```
    this.context.addStrokeListener(this);
}
```

The `deactivateApp` releases these same resources and destroys the ICR engine when it calls the `switchToReadyMode` method (also not an API call). It unregisters the `StrokeListener` at the same time. The calls are:

```
this.context.removeStrokeListener(this);
this.hwrEngine.clearStrokes();
this.hwrEngine.dispose();
this.hwrEngine = null;
```

Event Handling

There are five types of events that `SampleTranslator` handles. They provide the heart of its functionality. The event handlers are:

- `handleMenuEvent`

- `strokeCreated`

- ICR engine events, which includes these event handlers:

    a. `hwrResult`

    b. `hwrUserPause`

- `penDown`

handleMenuEvent

The system calls this event handler whenever it has a new menu event. The penlet can handle the event and return true—to indicate that the event is fully handled and need not be further processed by the system. If the penlet returns false, then the system handles the menu event in a generic way.

This penlet handles menu events in two ways:

- It sets the current focus of the `BrowseList` object (the application men) and passes it to the `setCurrent` method on the Display object.

OR

- It displays a user input prompt such as "Write a word." or "Try writing: (Hello Goodbye Please Thank you Thanks Zero One Two Three Four Five Six Seven Eight Nine Ten Coffee Chocolate Banana Beer Water)"

The `CurrentMenu` field is not an API field. It simply keeps track of what "mode" the application menu is in: `TRANS_MENU_LANGUAGE`, `TRANS_MENU_WRITE_TAP`, and `TRANS_MENU_ACTIVE`. You need not worry about the details, unless they interest you.

The event-handling logic that exercises the API involves the `MenuEvent` class constants: `MENU_UP`, `MENU_DOWN`, `MENU_SELECT`, `MENU_LEFT`, and `MENU_RIGHT`. Recall that `MENU_SELECT` stands for tapping on the center of the Nav Plus.

Up, Down, Center, and Left Menu Events

The up, down, center, and left menu events may arrive when the Sample Translator user has written or tapped or word or has just received the "Write a Word" message. In that case, we set the focus of the `BrowseList` object (i.e., `menuView`) and pass it to the `Display` object for display on the Livescribe smartpen OLED. The code is:

```
if ((event==MenuEvent.MENU_UP)||
   (event==MenuEvent.MENU_DOWN)||
   (event==MenuEvent.MENU_SELECT)||
   (event==MenuEvent.MENU_LEFT)) {

   this.menuView.setFocusItem(this.currentEngine);
   this.display.setCurrent(this.menuView);
   setCurrentMenuMarker(TRANS_MENU_LANGUAGE);
   return true;
}
```

Note that we set the focus of the `menuView` object to a number representing the current target language (`currentEngine` is a byte that stands for the current translation engine). Then we pass `this.menuView` to `display.setCurrent` so that `menuView` will be displayed.

Right Menu Event

The `RIGHT_MENU` event is key to any penlet. This is the event that the system sends when the user taps on the right arrow of a Nav Plus. It means "select the current item" or "display a submenu." In Sample Translator, it displays the "Write a word" message.

The relevant code is:

```
if (event == MenuEvent.MENU_RIGHT) {

// Obtain the single prompt string and play associated APM
   this.playCommandAPM();

// Draw the single prompt string
   this.labelUI.draw(INPUT_PROMPT,true);
   this.display.setCurrent(this.labelUI);
```

```
    setCurrentMenuMarker(TRANS_MENU_WRITE_TAP);
    return true;
}
```

Note that this case handles the event by playing a sound *and* displaying text. The term APM means "Audio Punctuation Mark". In this case, it is the audio command to the user: "Write a word." When users tap the right arrow, they are leaving the application menu. In order to display to the OLED, we need a `ScrollLabel` object such as `labelUI`. We call `draw` on that object and then pass it to `display.setCurrent`. Don't' forget this last step; if you do, your `ScrollLabel` object will not be displayed.

Navigating Up and Down in a BrowseList

Consider the state of Sample Translator when the OLED is currently displaying the application menu and the user taps the up and down arrows to sequentially access the four choices for target language. In that situation, `handleMenuEvent` must implement the browse up and browse down response of the `BrowseList` object. Here's the code:

```
int selection
if (event == MenuEvent.MENU_DOWN) {
    selection = this.menuView.focusToNext();
}
else if (event == MenuEvent.MENU_UP) {
    selection = this.menuView.focusToPrevious();
}
else {
    selection = this.menuview.getfocusIndex();
}
this.setEngine((byte)selection);
return true;
}
```

The `focusToNext` and `focusToPrevious` methods are defined in the `BrowseList` class. Note `setEngine`, which is not part of the API. It updates the `currentEngine` field of Sample Translator, based on the whether the user browsed up or down. Thus, the `currentEngine` field always reflects the current focus item (or, translation engine) of the `BrowseList`. Keep this fact in mind for the next section.

Tapping Back From a Right Menu Event

When the user taps the right arrow of a Nav Plus, Sample Translator abandons the application menu and displays the "Write a word" message, followed by the English word and translation (if the user's writing is successfully read by the ICR engine and the source word is contained in the English source word array.)

In that state, the penlet may wish to respond to an up, down, left, or center tap As we saw above, Sample Translator causes the application menu to be displayed. But at that point, how does the penlet know which language on the menu to display. If you look at the code in the section titled "Up, Down, Center, and Left Menu Events" you will quickly identify the line of code:

```
this.menuView.setFocusItem(this.currentEngine);
```

We set the current focus of `menuView` by passing the number that represents the current translation engine. When the user was tapping up and down, we preserved that number by calling

```
this.setEngine((byte)selection)
```

Remember: If "tapping back" after a right menu event makes sense in your penlet, you must preserve the current focus of the application menu's `BrowseList` object when you handle the `MENU_RIGHT` event.

strokeCreated

This event handler is called by the system when the user makes a stroke on dot paper. The system passes as parameters the start time of the stroke, the region on which it occurred, and the page of dot paper (i.e., the `PageInstance` object).

The system knows which penlet owns the region by calling the `getInstance` method on the region. The value returned is the penlet instance ID, which is assigned by the system. If the stroke occurs on Open Paper, the region ID is 0. Consequently, the instance ID is also 0, since the area ID is a 16-bit subset of the region ID.

Sample Translator tests for the instance ID associated with the region that `strokeCreated` passes in. If it is 0, then the stroke occurred on Open Paper and the penlet calls `addStroke` to send the stroke to the current ICR engine for analysis into a character. Otherwise, the stroke should be ignored. The code looks like this:

```
if (OPEN_PAPER_INSTANCE_ID==region.getInstance()){
   if (this.engine != null) {
    this.hwrEngine.addStroke(pageInstance, startTime);
   }
}
```

For the sake of convenience, this penlet defined the constant `OPEN_PAPER_INSTANCE_ID = 0`.

HWR Events: hwrUserPause and hwrResult

When the ICR engine (also known as: the HWR engine) receives strokes via the `addStroke` method call, the engine tries to assemble strokes into likely characters. It then compares the growing character string with the words in its lexicon.

hwrUserPause

When the user stops writing for 1000 milliseconds, the ICR engine posts an `hwrUserPause` event, which causes the system to call the `hwrUserPause` event handler for the current penlet. It passes as parameters the time the word was written and the result that the ICR engine produced. The result is a String.

In Sample Translator, the `hwrUserPause` handler calls the non-API method `processText`, which gets the translated word and audio and then calls `showTranslation` to display the translation and play the audio. Then `processText` proceeds to:

1. Get the bounding box of the ICR result by calling getTextBoundingBox on the ICRContext object.
2. Determine an areaId, based on the string passed by the result parameter. Sample Translator simply uses that string to find the index of the ENGLISH_SOURCE_WORDS array and uses that index as the areaId. (See the non-API method getAreaId for details.)
3. Get the region collection for the current page instance and call the non-API method addDynamicArea to perform these tasks:

    *   create a new Region object with the areaId that you found n the previous step.
    *   add the region to the region collection.

4. Finally, processText calls clearStrokes on the ICRContext object.

The code in `processText` that pertains to creating a new region includes the following snippets. Please note that `hwrEngine` is an `ICRContext` object:

```
...
Rectangle wordBox = this.hwrEngine.getTetBoundingBox()
...
RegionCollection rc=this.context.getCurrentRegionCollection();
Rectangle wordbox = this.hwrEngine.getTextBoundingBox();
```

```
...
if (!rc.isOverlappingEsxistingArea(wordBox) && wordAID >=0){
   addDynamicArea (wordBox, wordAID, ac);
```

The code in `addDynamicArea` actually creates the region and adds it to the region collection:

```
Region tempRegion = new Region (areaID, false, false);
ac.addRegion(rect, tempRegion, false);
```

When `processText` has completed, an area ID is now associated with the new region. When a user taps on that region, `areaId` can be used to determine what behavior the penlet should exhibit. For details, see penDown.

Always call clearStrokes

At the very end of `processText` you see this call:

```
this.hwrEngine.clearStrokes();
```

Remember to clear the strokes from the ICR engine when you have finished processing `hwrUserPause` event handler. If you do not, the ICR engine will give unpredictable results.

hwrResult

The `hwrResult` event handler is called whenever the ICR engine analyzes a character and then tries out various words in its lexicon that would fit. A penlet can display these intermediate steps, process them in some way, or ignore them. If the penlet ignores them, users will see no feedback on the smartpen OLED while they are writing a word.

Sample Translator chooses to display each intermediate "guess" of the ICR engine as comforting feedback to the user that the penlet is still operating. In addition, an inaccurate result lets the user know that they will have to re-write the current word.

The code for displaying the ICR engine's results in real time is:

```
this.labelUI.draw(result();
if (this.display.getCurrent() != this.labelUI) {
   this.display.setCurrent(this.labelUI);

}
```

penDown

The system calls the `penDown` handler when a user taps the smartpen tip down on dot paper. Like `strokeCreated`, the system passes as parameters the time of the event, the region it occurred on, and the page instance.

In Sample Translator, the code first checks to see if the `penDown` is on Open Paper. It does this by calling `getInstance` on the region. This returns the `instanceId` of the penlet. An `instanceId` of 0 indicates that the event occurred on no region—that is, on Open Paper. Our code simply returns:

```
if (OPEN_PAPER_INSTANCE_ID==region.getInstance() {
   return;
}
```

If the `penDown` was on a region, then we know that it belongs to Sample Translator. When an event occurs on a region, the system sends the event just to the *owner* of the region. The pertinent code is:

```
int areaID = region.getAreaId();

// Log AreaIDs that translator handles

// If AreaID is between 0 and source word array length

if((areaID >= 0) && (areaID < (ENGLISH_SOURCE_WORDS.length)) )
{
// If the source word is supported, request engine processing
   processDynamicAreaId(areaID);
}
```

The non-API method `processDynamicAreaId` uses the `areaID` to retrieve the English word from the `ENGLISH_SOURCE_WORDS` array, the word or image from the correct target language hash table, and then call the non-API method `showTranslation`, which displays the translation and plays the pronunciation.