



Funambol Developer's Guide

Last revised: July 25, 2008

Table of Contents

1.Introduction	5
1.1. Document structure	5
1.2. Audience.....	5
1.3. Funambol licensing.....	5
1.4. Comments and feedback.....	5
2.Getting started on connector development.....	6
2.1. Introduction.....	6
2.2. Getting started.....	6
2.3. Overview.....	7
2.4. Create the connector project.....	7
2.4.1. MyMergeableSyncSource type.....	9
2.4.2. MySynclet.....	9
2.4.3. MySyncSourceAdminPanel.....	9
2.5. Creating and installing the connector package.....	10
2.6. Creating a SyncSource.....	12
2.7. Testing the connector.....	13
3.Funambol development.....	16
3.1. Data synchronization.....	16
3.1.1. ID handling.....	16
3.1.2. Change detection.....	17
3.1.3. Modification exchange.....	17
3.1.4. Conflict detection.....	17
3.1.5. Conflict resolution.....	18
3.1.6. Full and fast synchronization.....	18
4.Funambol architecture.....	19
4.1. System architecture.....	19
4.1.1. Roles and responsibilities.....	19
4.2. The synchronization engine.....	22
4.3. Execution flow of a request.....	22
5.The synchronization process.....	24
5.1. Preparation.....	25
5.2. Modifications detection.....	25
5.3. Synchronization.....	27
5.4. Finalization.....	27
6.Extending Funambol.....	28
6.1. Building a Funambol module.....	28
6.2. Modules, connectors, listeners and SyncSource types.....	29

6.2.1. Registering modules, connectors and SyncSource types.....	30
7.Developing a SyncSource.....	32
7.1. The SyncSource interface and related classes.....	32
7.1.1. SyncContext.....	33
7.1.2. SyncItem.....	34
7.1.3. Twin items.....	34
7.1.4. The Administration Tool configuration panel.....	35
8.Extending the Funambol Administration Tool.....	36
8.1. Architecture overview.....	36
8.2. ManagementObject and subclasses.....	37
8.2.1. com.funambol.admin.mo.SyncSourceManagementObject.....	38
8.2.2. com.funambol.admin.mo.ConnectorManagementObject.....	38
8.3. ManagementObjectPanel and subclasses.....	39
8.3.1. SourceManagementPanel.....	40
8.3.2. ConnectorManagementPanel.....	40
9.Configuring Funambol components.....	41
9.1. System properties.....	41
9.2. Server JavaBeans.....	41
9.2.1. The configuration path.....	43
9.2.2. Lazy initialization.....	43
9.3. How to configure a standard component.....	43
9.4. How to create a custom configurable object.....	43
9.5. How to get a configured instance.....	45
9.5.1. Tips and tricks.....	45
10.Customizing message processing.....	46
10.1. Overview.....	46
10.2. Preprocessing an incoming message.....	46
10.2.1. Creating an input synclet.....	47
10.2.2. Configuring an input synclet.....	48
10.3. Postprocessing an outgoing message.....	48
10.3.1. Creating an output synclet.....	49
10.3.2. Configuring an output synclet.....	49
10.3.3. The MessageProcessingContext.....	49
10.3.4. How to stop message processing.....	50
11.SyncSource API.....	51
11.1. SyncSource class.....	51
11.1.1. Methods list.....	52
11.2. Mergeable SyncSource methods.....	53
11.2.1. Methods list.....	54

11.3. Filterable SyncSource methods.....	54
11.3.1. Methods list.....	54
12. Officer API.....	55
12.1. Officer class.....	55
12.1.1. Methods list.....	55
13. Web Services API.....	56
13.1. Introduction.....	56
13.1.1. Funambol DS Server Web Services.....	56
14. Funambol Software Development Kit.....	59
14.1. Obtaining and building the source code.....	59
14.2. Developing with a custom environment.....	59
14.3. Developing with Maven.....	60
14.3.1. Maven configuration.....	60
14.3.2. Creating a new module.....	61
14.3.3. Building the module.....	61
14.4. The Funambol Connector Testing Framework.....	61
14.4.1. Usage.....	62
14.4.2. Certifying a connector.....	64
14.4.3. Limitations.....	65
14.4.4. Error codes.....	65
15. Appendix A - Sync4j Interchange Formats.....	67
15.1. SIF-C.....	67
15.2. SIF-E.....	72
15.2.1. Constants.....	76
15.2.2. Recurrent event examples.....	76
15.3. SIF-T.....	77
15.4. SIF-N.....	79
15.4.1. Constants.....	80
16. Appendix B – List of acronyms.....	81
17. Resources.....	82

1. Introduction

This document is intended for developers who aim to develop synchronization services based on the Funambol platform.

This development guide gives a deep insight of the server internals and design, providing guidance to anyone aiming to take advantage of the full range of possibilities that the platform provides.

1.1. Document structure

Chapter 2 is a quick start guide to developing Funambol connectors. Chapters 3-9 present an overview of the Funambol architecture and internal workings. Chapters 11-13 describe in detail the SyncSource, Officer and Web Services APIs. Chapter 14 is a short introduction to the Funambol SDK.

1.2. Audience

This document is addressed to anybody wanting to extend the Funambol platform or simply looking for detailed information on the Funambol architecture.

1.3. Funambol licensing

All Funambol software and software developed using Funambol APIs or SDKs are licensed under AGPL V3 (Afero General Public License) unless separate arrangements have been made with Funambol to reach an explicit commercial agreement.

More information on AGPL V3 can be found here: <http://www.fsf.org/licensing/licenses/agpl-3.0.html>.

1.4. Comments and feedback

The Funambol team wants to hear from you! Please access our community portal at <https://www.forge.funambol.org/participate> and submit your questions, comments, feedbacks or testimonials.

2. Getting started on connector development

2.1. Introduction

This chapter describes how to create a connector that extends the functionality of the Funambol server.

As better described in chapter 6, a Funambol extension is delivered in the form of a module, which consists of a packaged set of files, including classes, configuration files, server beans and initialization SQL scripts. All these contents are deployed into the Funambol server to provide access to a specific back-end (e.g. a database, a REST based service, a web services API, etc.). In general, a module can be viewed as a container for anything related to server extensions. When a module provide access to a specific backend, it is called *connector*.

The following terms and concepts will be used in this document:

Module: a container for anything related to one or more server extensions which are used by the engine to integrate with external systems.

Connector: a particular type of module, with the purpose of connecting to an external data source; in other words, a connector is an extension of the server intended to support the synchronization of a particular data source.

SyncSource: a key component of a connector that defines the way a set of data is made accessible to the Funambol server for synchronization. A SyncSource type represents the template from which an instance of a SyncSource can be created. For example, the *FileSystemSyncSource* type defines how data stored in the file system can be accessed by the Funambol server; however, it does not represent a specific directory to be used for synchronization, and in order to synchronize a specific directory an instance of *FileSystemSyncSource* must be created and configured with the desired directory.

Synclet: a pre or post processing unit that can process a message before it gets into the synchronization engine or just after it is going out from it.

This chapter will guide you through the development, packaging, installation and testing of a module. The module contains a simple SyncSource and Synclet which produce some logging. Once you are familiar with this tutorial you can see real-world examples like the OpenXchange connector [8] or the Exchange connector [9]. Plus, many people have developed many modules and connectors that are available to the public. See [5] for more information.

2.2. Getting started

The following connector development quick-start section assumes a working knowledge of Java, Maven and SQL.

For more detailed information about Funambol development see the next sections of this developer's guide.

In order to follow this guide you need:

- Funambol DS Server installed and running
- Funambol Software Development Kit 7.0.x [7]
- Java 2 SDK version 1.5.x or above
- Apache Maven [4]

Optionally, you may want to download a Maven plug-in for your preferred IDE (see <http://mevenide.codehaus.org>).

Download the software and install it in a directory of your choice; we will assume the following prefix:

`$FUNAMBOL_HOME`: the directory where the bundle has been installed (e.g.: `/opt/Funambol`)

`$FUNAMBOL_SDK_HOME`: the directory where the Funambol SDK has been installed (e.g.: `/opt/Funambol/tools/sdk`)

`$JAVA_HOME`: the directory where Java is installed (e.g.: `/opt/jdk1.5.0_10`)

`$MAVEN_HOME`: the directory where Apache Maven is installed (e.g.: `/opt/apache-maven-2.0.8`)

`$USER_HOME`: the home directory of the operating system user you are on (e.g.: `/home/ste`, `c:\Users\ste`).

Note: Basic knowledge of Apache Maven, its terminology and principles are assumed, as the following sections use terms from the Apache Maven world without explaining them in details.

After installing Maven, you need to configure it so that it points to the Funambol public Maven repository (`m2.funambol.org/repositories`). To do so, copy the file `$FUNAMBOL_SDK_HOME/docs/settings.xml` under `$USER_HOME/.m2`.

2.3. Overview

We will develop the sample module following these steps:

1. Create the connector project
2. Install the module
3. Create a SyncSource instance
4. Test the module with a SyncML client

2.4. Create the connector project

The easiest way to create a connector project is using Maven. Run the following command from a command line shell:

```
mvn archetype:create -DarchetypeGroupId=funambol -DarchetypeArtifactId=funambol-  
module-archetype -DarchetypeVersion=7.0.0 -DgroupId=acme -DartifactId=acmeconnector
```

This will download and create a skeleton application ready to be built which contains:

1. a normal SyncSource
2. a mergeable SyncSource
3. an input/output synclet
4. all configuration and SQL files

Everything is generated in a Maven project located in directory named `acmeconnector`. The content of the directory is as illustrated in Figure 1.

Note: in order to create a Funambol 6.5 project, you just need to use `-DarchetypeVersion=6.5.2` on the command line.

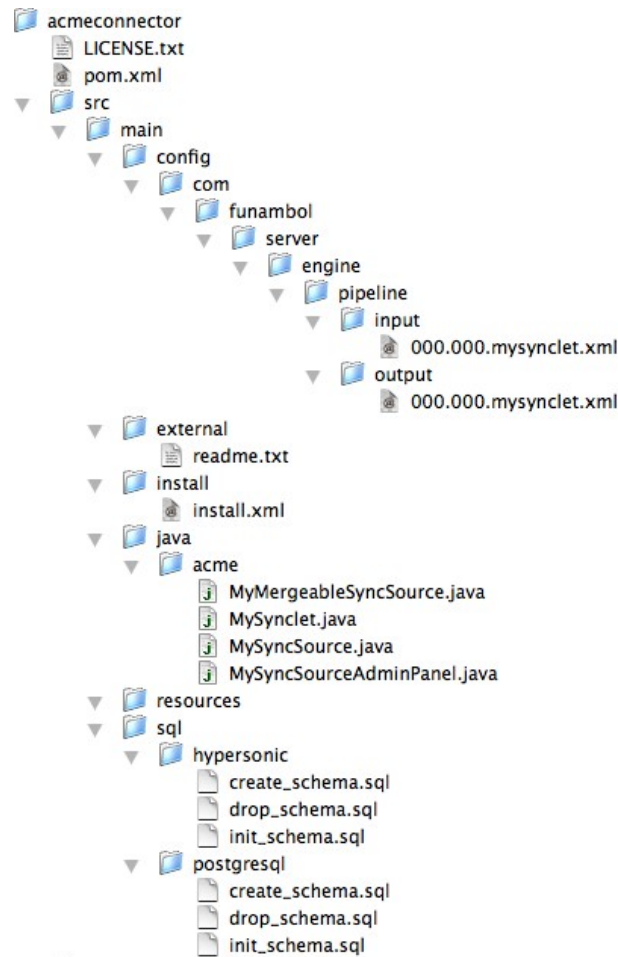


Figure 1: Module source directory structure

The following table explains the function of each file:

<i>File</i>	<i>Description</i>
LICENSE.txt	AGPL license file
pom.xml	Maven project file for the connector
src/main/config/com/funambol/server/engine/pipeline/input/000.000.mysynclet.xml	Sample input synclet configuration
src/main/config/com/funambol/server/engine/pipeline/output/000.000.mysynclet.xml	Sample output synclet configuration
src/main/external/readme.txt	Readme for the content of this directory
src/main/install/install.xml	Module installation file
src/main/java/acme/MyMergeableSyncSource.java	Sample mergeable SyncSource
src/main/java/acme/MySynclet.java	Sample synclet
src/main/java/acme/MySyncSource.java	Sample SyncSource
src/main/java/acme/MySyncSourceAdminPanel.java	Sample administration panel for both MySyncSource and MyMergeableSyncSource
src/main/sql/*/create_schema.sql	SQL scrip to create the database tables required by the module
src/main/sql/*/drop_schema.sql	SQL scrip to drop the database tables required by the module
src/main/sql/*/init_schema.sql	SQL scrip to initialize the database tables required by the module

Take a moment to explore and open each file.

Note: the skeleton project creates both a normal SyncSource and a MergeableSyncSource; in this chapter we will only consider the latter.

2.4.1. MyMergeableSyncSource type

The SyncSource type is the primary component of the connector. The source code created by the artifact is very simple and it only writes some logging info so that we can trace its execution. However in a real case this is where the code necessary to integrate an external data source will go.

MyMergeableSyncSource inherits most of its behavior from *MySyncSource*; open it into an editor or in your IDE and go through it. *MySyncSource* defines three properties that we are going to be able to set through the Administration Tool. These are: *myString*, *myInt* and *myMap*. Getter and setter methods are provided. Note also that the class implements all methods of the SyncSource interface writing a log entry. Each method has also a description of what it does and what the developer should add.

2.4.2. MySynclet

In addition to the SyncSource types described earlier, the archetype project contains also a sample input and output synclet: *MySynclet*. Open it into an editor or in your IDE and go through it.

Note: If you don't need to write a synclet, you can skip this section.

The first thing to note is that it implements both an input and an output synclet, which means that the synclet will be called for both incoming and outgoing messages. The synclet is very simple and once more it just logs the message in the *funambol.myconnector* logger.

2.4.3. MySyncSourceAdminPanel

We want to be able to create a new *MySyncSource* and configure it from the Funambol Administration Tool. This is possible thanks to the class *MySyncSourceAdminPanel*. Open it into an editor or in your IDE and go through it.

MySyncSourceAdminPanel inherits from *SourceManagementPanel*, which is a class of the Admin framework. *SourceManagementPanel* is a JPanel, therefore it has all methods of a swing panel. The *init()* method creates all widgets that we want to display in the Administration Tool and adds them to the panel. These widgets are:

- source name
- data types supported (e.g. text/vcard)
- data type versions supported (e.g. 2.1)
- source URI
- myString
- myInt
- myMap entry

In addition, it adds to the panel a JButton to add a newly created SyncSource or to save the values of an existing SyncSource. An important aspect to note is that here is where the panel interacts with the Administration Tool to persist the changes to the server. The code that performs this task is the following:

```
confirmButton.addActionListener(new ActionListener() {
    public void actionPerformed(ActionEvent event) {
        try {
            validateValues();
            getValues();
            if (getState() == STATE_INSERT) {
                SyncSourceAdminPanel.this.actionPerformed(
                    new ActionEvent(MySyncSourceAdminPanel.this,
                        ACTION_EVENT_INSERT,
                        event.getActionCommand()));
            } else {
                MySyncSourceAdminPanel.this.actionPerformed(
                    new ActionEvent(MySyncSourceAdminPanel.this,
                        ACTION_EVENT_UPDATE,
                        event.getActionCommand()));
            }
        } catch (Exception e) {
            notifyError(new AdminException(e.getMessage()));
        }
    }
});
```

```
    }  
  }  
});
```

The key is that, when needed, the method *actionPerformed()* of the base class is called with a proper event.

The other important method is *updateForm()* where the value of the SyncSource instance are displayed in the proper fields. Again this method is called by the Administration Tool when an existing instance must be displayed.

2.5. Creating and installing the connector package

To insert the created project into a Funambol module, just go into *acmeconnector* and type:

```
mvn package
```

A typical output will be as follows:

```
[INFO] Scanning for projects...  
[INFO] -----  
[INFO] Building acme acmeconnector Module  
[INFO]   task-segment: [package]  
[INFO] -----  
[INFO] artifact org.apache.maven.plugins:maven-resources-plugin: checking for updates  
from artifacts  
[INFO] artifact org.apache.maven.plugins:maven-resources-plugin: checking for updates  
from snapshots  
[INFO] artifact org.apache.maven.plugins:maven-compiler-plugin: checking for updates  
from artifacts  
[INFO] artifact org.apache.maven.plugins:maven-compiler-plugin: checking for updates  
from snapshots  
[INFO] artifact org.apache.maven.plugins:maven-surefire-plugin: checking for updates  
from artifacts  
[INFO] artifact org.apache.maven.plugins:maven-surefire-plugin: checking for updates  
from snapshots  
[INFO] artifact org.apache.maven.plugins:maven-jar-plugin: checking for updates from  
artifacts  
[INFO] artifact org.apache.maven.plugins:maven-jar-plugin: checking for updates from  
snapshots  
[INFO] [resources:resources]  
[INFO] Using default encoding to copy filtered resources.  
[INFO] [compiler:compile]  
[INFO] Compiling 4 source files to /Users/ste/Projects/acmeconnector/target/classes  
[INFO] [resources:testResources]  
[INFO] Using default encoding to copy filtered resources.  
[INFO] [compiler:testCompile]  
[INFO] No sources to compile  
[INFO] [surefire:test]  
[INFO] No tests to run.  
[INFO] [jar:jar]  
[INFO] Building jar: /Users/ste/Projects/acmeconnector/target/acmeconnector-1.0-  
SNAPSHOT.jar  
[INFO] [funambol:s4j]  
[INFO] Exploding Funambol packaging...  
[INFO] Assembling Funambol packaging acmeconnector in  
/Users/ste/Projects/acmeconnector/target/acmeconnector-1.0-SNAPSHOT  
[INFO] Including license file /Users/ste/Projects/acmeconnector/LICENSE.txt  
[INFO]  
[INFO]   Including artifacts:  
[INFO]   -----  
[INFO]     x funambol:server-framework:jar:7.0.3-SNAPSHOT:compile  
[INFO]     x funambol:core-framework:jar:6.5.4:compile  
[INFO]     x funambol:ext:jar:6.5.2:compile  
[INFO]     o org.jibx:jibx-run:jar:1.1.2fun:compile  
[INFO]     o xpp3:xpp3:jar:1.1.2a-fun:compile  
[INFO]     o commons-lang:commons-lang:jar:2.3:compile  
[INFO]     o funambol:admin-framework:jar:6.5.2:compile  
[INFO]  
[INFO]   Excluded artifacts:  
[INFO]   -----  
[INFO]  
[INFO]   Including jar files...  
[INFO] basedir: /Users/ste/Projects/acmeconnector  
[INFO] srcDir: /Users/ste/Projects/acmeconnector/src/main
```

```

[INFO] sqlDirectory: /Users/ste/Projects/acmeconnector/target/acmeconnector-1.0-
SNAPSHOT/sql
[INFO] wsddDirectory: /Users/ste/Projects/acmeconnector/target/acmeconnector-1.0-
SNAPSHOT/wsdd
[INFO]     No config files...
[INFO]     No exclude files...
[INFO]     Including install files...
[INFO]     Including sql files...
[INFO]     No wsdd files...
[INFO]
[INFO] Generating Funambol packaging
/Users/ste/Projects/acmeconnector/target/acmeconnector-1.0-SNAPSHOT.s4j
[INFO] Building jar: /Users/ste/Projects/acmeconnector/target/acmeconnector-1.0-
SNAPSHOT.s4j
[INFO] -----
[INFO] BUILD SUCCESSFUL
[INFO] -----
[INFO] Total time: 12 seconds
[INFO] Finished at: Sat May 17 15:42:08 CEST 2008
[INFO] Final Memory: 9M/16M
[INFO] -----

```

Note: The final message “Generating Funambol packaging /Users/ste/Projects/acmeconnector/target/acmeconnector-1.0-SNAPSHOT.s4j” tells where the package is created. In order to install it into the server, copy *acmeconnector-1.0-SNAPSHOT.s4j* into *\$FUNAMBOL_HOME/ds-server/modules* and follow the steps below.

1. Make sure Funambol is up and running
2. Using a text editor, open the *\$FUNAMBOL_HOME/ds-server/install.properties* file.
3. Find the line that begins *modules-to-install=* in the Module definitions section. This line specifies, in a comma-separated list, the modules to install during installation.
4. Add *acmeconnector-1.0-SNAPSHOT* to the comma-separated list (note that you do not have to specify the *.s4j* filename extension).
5. Save and close *install.properties*.
6. On Windows, open a command prompt window and run the server installation script by typing the following at the prompt:

```

cd $FUNAMBOL_HOME/ds-server
bin\install-modules

```

7. On Unix/Linux, use the command:

```

cd $FUNAMBOL_HOME/ds-server
bin/install-modules

```

Press 'y' when asked to install the database for *acmeconnector* and 'n' for all others.

```

[echo] Funambol Data Synchronization Server will be installed on the Tomcat 5.5.x
application server
[echo] Undeploying funambol...
[echo] Pre installation for modules foundation-7.0.1,acmeconnector-1.0-SNAPSHOT
[echo] foundation-7.0.1 pre-installation...
[echo] foundation-7.0.1 pre-installation successfully completed
[echo] acmeconnector-1.0-SNAPSHOT pre-installation...
[echo] acmeconnector-1.0-SNAPSHOT pre-installation successfully completed
[echo] Copying configuration files
[echo] Post installation for modules foundation-7.0.1,acmeconnector-1.0-SNAPSHOT
[echo] has.install: true
[echo] Starting custom installation...
[echo] Foundation Installation
[echo] Foundation installation successfully completed
[echo] foundation-7.0.1 installation...
[echo] Database installation for module foundation-7.0.1 on hypersonic (/opt/Funambol/
ds-server)

[iterate] The Funambol Data Synchronization Server installation program can now create
[iterate] the database required by the module foundation-7.0.1 (if any is needed).
[iterate] You can skip this step if you have already a valid database created
[iterate] or the module does not require a database.

```

```
[iterate] If you choose 'y' your existing data will be deleted.
[iterate] Do you want to recreate the database?
[iterate]          (y,n)
n
```

and

```
[echo] foundation-7.0.1 installation successfully completed
[echo] has.install: true
[echo] Starting custom installation...
[echo] acmeconnector installation
[echo] acmeconnector installation successfully completed
[echo] acmeconnector-1.0-SNAPSHOT installation...
[echo] Database installation for module acmeconnector-1.0-SNAPSHOT on hypersonic
(/opt/Funambol/ds-server)

[iterate] The Funambol Data Synchronization Server installation program can now create
[iterate] the database required by the module acmeconnector-1.0-SNAPSHOT (if any is
needed) .
[iterate] You can skip this step if you have already a valid database created
[iterate] or the module does not require a database.
[iterate] If you choose 'y' your existing data will be deleted.
[iterate] Do you want to recreate the database?
[iterate]          (y,n)
y
[echo] acmeconnector-1.0-SNAPSHOT installation successfully completed
[war] Warning: selected war files include a WEB-INF/web.xml which will be ignored
(please use webxml attribute to war task)
[echo] Remove output dir

BUILD SUCCESSFUL
Total time: 12 seconds
```

Now the connector is installed.

2.6. Creating a SyncSource

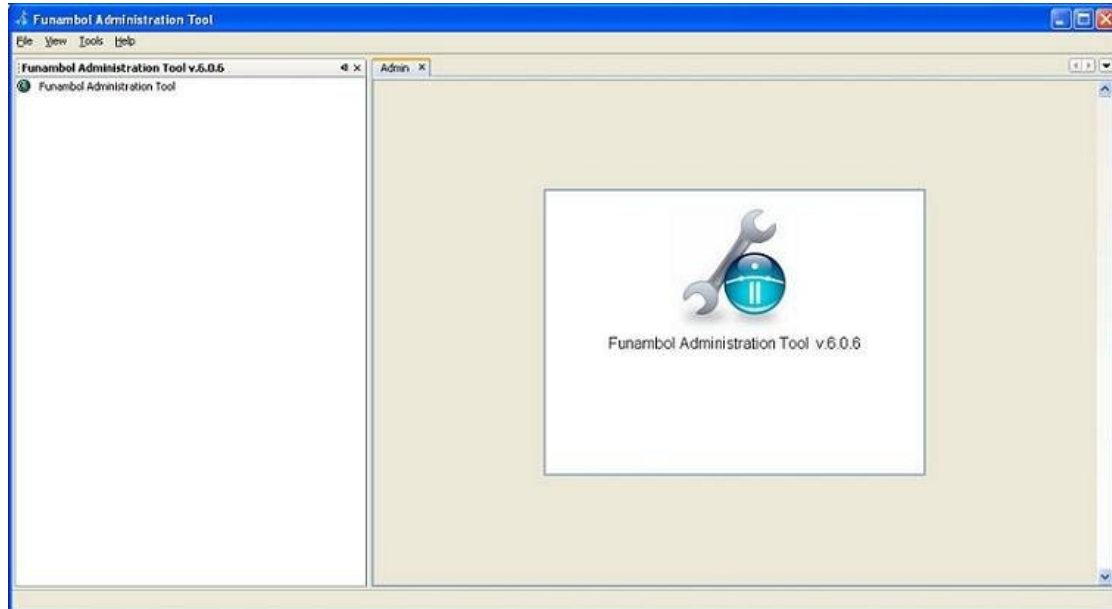


Figure 2: Funambol Administration Tool

Now that the connector is installed, you can see it in the Administration Tool. To access it, start the Funambol Administration Tool by selecting Start | All Programs | Funambol | Administration Tool. The Funambol Administration Tool window appears (see Figure 2).

First of all, you will need to log into the server: on the Main Menu bar, select File | Login. The Login window displays. Verify that the fields are populated as follows, or specify the following values:

```
Hostname/IP: <localhost> (should be your machine name)
Port: 8080
User name: admin
Password: sa
```

Click Login; the Output Window in the lower right pane should display "connected." In the left panel, expand the localhost tree as follows: localhost | Modules | acme | acmeconnector, then select MyMergeableSyncSource. The Edit My SyncSource screen appears in the upper right pane, as shown in Figure 3.

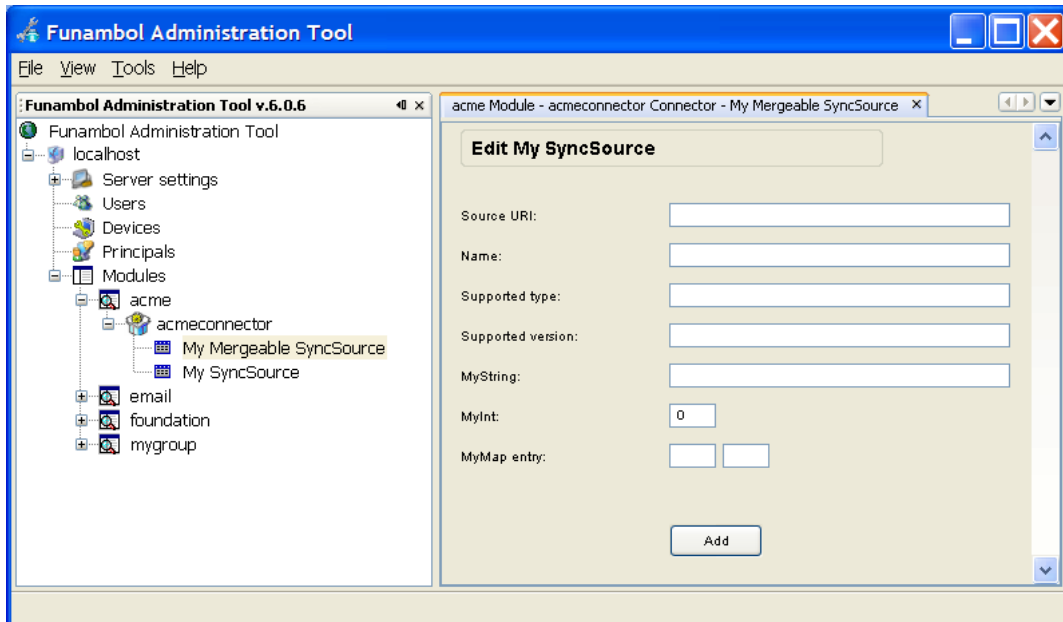


Figure 3: Edit File System SyncSource screen

This window is used to specify configuration values. Insert the following values and press Add:

```
Source URI: acme
Name: Acme
Supported type: text/plain
Supported version: 1.0
MyString: acme connector!
MyInt: 10
MyMap entry: <acme, connector>
```

2.7. Testing the connector

To test the Acme connector we will use a simple command line SyncML client which is distributed in the Funambol SDK under `$FUNAMBOL_SDK_HOME/plugin-cl`. Perform the following:

1. Edit the file `config/spds/sources/briefcase.properties` and set the following values:

```
name=acme
sourceClass=com.funambol.syncclient.spds.source.FileSystemSyncSource
sourceDirectory=db/briefcase
type=text/plain
sync=two-way
encode=true
sourceURI=acme
```

2. Run the command `$FUNAMBOL_SDK_HOME\plugin-cl\binrun.cmd` (or `$FUNAMBOL_SDK_HOME/plugin-cl/run.sh` if using Linux)

You will see an output similar to the following:

Funambol Command Line Tool ver. 7.0.1

```
-----
2008-04-17 16:28:10:969 - # SyncClient API J2SE Log
16:28:10:970 [INFO] - Initializing
16:28:10:973 [INFO] - Sending initialization commands
16:28:11:716 [INFO] - The server alert code for acme is 201
16:28:11:718 [INFO] - Synchronizing acme
16:28:11:745 [INFO] - exchange modifications started
16:28:11:746 [INFO] - Preparing slow sync for acme
16:28:11:747 [INFO] - Detected 0 items
16:28:11:748 [INFO] - Sending modifications
16:28:11:837 [INFO] - Returned 0 new items, 0 updated items, 0 deleted items for acme
16:28:11:838 [INFO] - Mapping started
16:28:11:841 [INFO] - Sending mapping
16:28:11:853 [INFO] - Sending mapping
16:28:11:874 [INFO] - Mapping done
16:28:11:874 [INFO] - Synchronization done
-----
```

In the server log you will be able to see your connector at work. Filtering out the lines that are not of interest for our connector, the log entries will be similar to the following text:

```
[2008-05-17 16:31:56,656] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [] Initializing
acme.MyMergeableSyncSource

[2008-05-17 16:31:56,657] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [] myString: acme connector!

[2008-05-17 16:31:56,657] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [] myInt: 10

[2008-05-17 16:31:56,657] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [] myMap: {acme=connector}

[2008-05-17 16:31:56,703] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [acme] Starting
synchronization: com.funambol.framework.engine.source.SyncContext@e85079

[2008-05-17 16:31:56,703] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [acme] getNewSyncItemKeys ()

[2008-05-17 16:31:56,703] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [acme]
getUpdatedSyncItemKeys ()

[2008-05-17 16:31:56,703] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [acme]
getDeletedSyncItemKeys ()

[2008-05-17 16:31:56,703] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [acme] Committing
synchronization

[2008-05-17 16:31:56,782] [funambol.myconnector] [INFO]
[399F31F06404433DE69A05F71D562ACD] [sc-api-j2se] [guest] [acme] Ending synchronization
-----
```

Note: in the first entries the values inserted earlier in the administration panel are displayed, meaning that the SyncSources were properly configured.

Similarly, we can see the effect of *MySynclet*. Each SyncML message has been logged; for example the first input message is something like the following text:

```
[2008-05-17 16:59:28,576] [funambol.myconnector] [INFO]
[FF4B335BA02F0581DAFF016319EDD263] [] [] []
-----
[2008-05-17 16:59:28,576] [funambol.myconnector] [INFO]
[FF4B335BA02F0581DAFF016319EDD263] [] [] [] Input message[2008-05-17 16:59:28,576]
[funambol.myconnector] [INFO] [FF4B335BA02F0581DAFF016319EDD263] [] [] []
[2008-05-17 16:59:28,585] [funambol.myconnector] [INFO]
[FF4B335BA02F0581DAFF016319EDD263] [] [] [] <?xml version="1.0" encoding="UTF-8"?
><SyncML><SyncHdr><VerDTD>1.1</VerDTD><VerProto>SyncML/1.1</VerProto><SessionID>1234567
8</SessionID><MsgID>1</MsgID><Target><LocURI>http://localhost:8080/funambol/ds</LocURI>
</Target><Source><LocURI>sc-api-j2se</LocURI></Source><Cred><Meta><Type>syncml:auth-
basic</Type></Meta><Data>Z3Vlc3Q6Z3Vlc3Q=</Data></Cred><Meta><MaxMsgSize>250000</MaxMsg
Size><MaxObjSize>4000000</MaxObjSize></Meta></SyncHdr><SyncBody><Alert><CmdID>1</CmdID>
<Data>200</Data><Item><Target><LocURI>acme</LocURI></Target><Source><LocURI>acme</LocUR
I></Source><Meta><Anchor><Last>1211034716596</Last><Next>1211036368401</Next></Anchor><
```

```
/Meta></Item></Alert><Final/></SyncBody></SyncML>[2008-05-17 16:59:28,585]
[funambol.myconnector] [INFO] [FF4B335BA02F0581DAFF016319EDD263] [] [] []
```

And the last output message is something like the following:

```
[2008-05-17 16:59:28,876] [funambol.myconnector] [INFO]
[FF4B335BA02F0581DAFF016319EDD263] [sc-api-j2se] [guest] []
```

```
[2008-05-17 16:59:28,877] [funambol.myconnector] [INFO]
[FF4B335BA02F0581DAFF016319EDD263] [sc-api-j2se] [guest] [] Output message[2008-05-17
16:59:28,877] [funambol.myconnector] [INFO] [FF4B335BA02F0581DAFF016319EDD263] [sc-api-
j2se] [guest] []
[2008-05-17
16:59:28,877] [funambol.myconnector] [INFO] [FF4B335BA02F0581DAFF016319EDD263] [sc-api-
j2se] [guest] [] <?xml version="1.0" encoding="UTF-8"?
><SyncML><SyncHdr><VerDTD>1.1</VerDTD><VerProto>SyncML/1.1</VerProto><SessionID>1234567
8</SessionID><MsgID>5</MsgID><Target><LocURI>sc-api-
j2se</LocURI></Target><Source><LocURI>http://localhost:8080/funambol/ds</LocURI></Sourc
e><RespURI>http://localhost:8080/funambol/ds;jsessionid=FF4B335BA02F0581DAFF016319EDD26
3</RespURI></SyncHdr><SyncBody><Status><CmdID>1</CmdID><MsgRef>5</MsgRef><CmdRef>0</Cmd
Ref><Cmd>SyncHdr</Cmd><TargetRef>http://localhost:8080/funambol/ds</TargetRef><SourceRe
f>sc-api-
j2se</SourceRef><Data>200</Data></Status><Final/></SyncBody></SyncML>[2008-05-17
16:59:28,877] [funambol.myconnector] [INFO] [FF4B335BA02F0581DAFF016319EDD263] [sc-api-
j2se] [guest] []
```

If you have reached this point, you have probably successfully developed a Funambol connector: congratulations! You can now go in more details in the *acmeconnector* project or start your own.

3. Funambol development

The following sections present several concepts related to how Funambol can be used to develop mobile applications. Before digging into the details of Funambol development, it is useful to describe some basic concepts of data synchronization since this is the basis for many Funambol applications and services.

3.1. Data synchronization

All mobile devices – handheld computers, mobile phones, pagers, laptops – need to synchronize their data with the server where the information is stored. This ability to access and update information on the fly is key to the pervasive nature of mobile computing. Yet, today, almost every device uses a different technology for performing data synchronization.

Data synchronization is helpful in respect to many areas:

- Propagating updates between a growing number of applications
- Overcoming the limitations of mobile devices and wireless connections
- Maximizing user experience minimizing data access latency
- Keeping scalability of the infrastructure in an environment where the number of devices (clients) and connections tends to increase considerably
- Understanding the requirements of mobile applications, providing a user experience that helps and is not an obstacle for mobile tasks

Data synchronization is the process of making two sets of data look identical (Figure 4). This involves many concepts, the most important are:

- ID handling
- Change detection
- Modification exchange
- Conflict detection
- Conflict resolution
- Slow and fast synchronization

3.1.1. ID handling

At a first look, ID handling seems a pretty straightforward process and of no interest. Instead, ID handling is an important aspect of the synchronization process and it is not trivial. Each piece of data is usually uniquely identifiable by a subset of its content fields; for example, in the case of a contact entry, the concatenation of first name and last name uniquely selects an entry in your directory. In other cases, the ID is represented by a particular field specifically introduced for that purpose. This may be the case, for example, of a Sales Force Automation mobile application, where an order is identified by an order number or reference. The way an item ID is generated is not determinable a priori and it is application and device specific.

In an enterprise system, however, data is stored in a centralized database, shared by all users; each single item is known by the system with a unique global ID. In some cases, two sets of data (i.e. the order on the client and the order on the server) represent the same information (*the “order”* made by the customer) but they differ. What could be done to reconcile client and server IDs in order to make the information consistent? Many approaches can be chosen:




Figure 4 - Data synchronization process

- Clients and server agree on a *ID scheme* (a convention on how to generate IDs must be defined and used);
- Each client generates globally unique IDs (GUIDs) and the server accepts client-generated IDs;
- The server generates globally unique IDs (GUIDs) and each client accepts those IDs;
- Client and server generate their own IDs and a mapping is kept between the two. Client side IDs are called Local Unique IDentifiers (LUID) and server side IDs are called Global Unique IDentifiers (GUID). The mapping between local and global identifiers is referred as LUID-GUID mapping.

3.1.2. Change detection

Change detection is the process of identifying the data that was modified after a particular point in time (i.e. the last synchronization). This is usually achieved making use of additional information such as timestamps and state information. For example, a possible database enabled for an efficient change detection is the one shown in Table 1.

<i>ID</i>	<i>First name</i>	<i>Last name</i>	<i>Telephone</i>	<i>State</i>	<i>Last_update</i>
12	John	Doe	+1 650 5050403	N	2008-04-02 13:22
13	Mike	Smith	+1 469 4322045	D	2008-04-01 17:32
14	Vincent	Brown	+1 329 2662203	U	2008-03-21 17:29

Table 1 - A database enabled for efficient change detection

However, sometimes legacy databases do not provide the information needed to accomplish an efficient change detection. Therefore, the matter becomes more complicated and alternative methods must be adopted (based on content comparison, for instance).

3.1.3. Modification exchange

A key component of a data synchronization infrastructure is the way modifications are exchanged between client and server. This involves the definition of a synchronization protocol that client and server have to use to initiate and carry on a synchronization session. In addition to the exchange modification method, a synchronization protocol must also define a set of supported modification commands. The minimal set of modification commands is represented by the following:

- Add
- Replace
- Delete

3.1.4. Conflict detection

Let's suppose two users synchronize their local contacts database with a central server in the morning, before going to the office. After syncing, they have exactly the same contacts on their PDAs. Let's now suppose that they change the telephone number of the same "John Doe" entry, but for some reason with a different number (maybe, one of the two made a mistake). What will happen when the next morning they will synchronize again? Which one of the two new versions of the John Doe record should be taken and stored into the server? This condition is called "conflict" and the server has the duty of identifying and resolving it.

The simplest way to do detect a conflict is by the means of a synchronization matrix (see Table 2).

<i>Database A</i> → ↓ <i>Database B</i>	<i>New</i>	<i>Deleted</i>	<i>Updated</i>	<i>Synchronized/Unchanged</i>	<i>Not Existing</i>
<i>New</i>	C	C	C	C	B
<i>Deleted</i>	C	X	C	D	X
<i>Updated</i>	C	C	C	B	B
<i>Synchronized/Unchanged</i>	C	D	A	=	B
<i>Not Existing</i>	A	X	A	A	X

Table 2 - The synchronization matrix

Because both users synchronize with the central database, we can consider what happens between the server database and one of the client databases at a time: let's call Database A the client database and Database B the server database. The symbols in the synchronization matrix have the following meaning:

- **X** : nothing to do
- **A** : item A replaces item B
- **B** : item B replaces item A
- **C** : conflict
- **D** : delete the item from the source(s) containing it

3.1.5. Conflict resolution

Once a conflict arises and it is detected, a proper action must be taken. Different policies can be applied:

- User decides: the user is notified of the conflict condition and decides what to do; this strategy, like the following "Client wins" is a bit problematic in a server centric synchronization solution: each user may have the same right to modify an item and one user could not be able to decide whether his/her modification should win over the other ones
- Client wins: the server silently replaces conflicting items with the ones sent by the client
- Server wins: the client has to replace conflicting items with the ones from the server
- Timestamp based: the last modified (in time) item wins
- Last/first in wins: the last/first arrived item wins
- Do not resolve

3.1.6. Full and fast synchronization

There are many modes to carry on the synchronization process. The main distinction is between fast and full synchronization. Fast synchronization involves only the items changed since the last synchronization between two devices. Of course, this is an optimized process that relies on the fact that, some time in the past, the devices were fully synchronized; this way, the state at the beginning of the sync operation is well known and sound. When this requisite is not met (because, for instance, the mobile device has been reset and lost the timestamp of the last synchronization), a *full synchronization* must be performed. In this case, the client sends its entire database to the server, which compares it with its local database and returns the modifications that must be applied to be up to date again.

Both fast and slow synchronization modes can be performed in one of the following manners:

- Client to server: the server updates its database with client modifications, but sends no server-side modifications.
- Server to client: the client updates its database with server modifications, but sends no client-side modifications.
- Two-way: client and server exchange their modifications and both databases are updated accordingly.

4. Funambol architecture

4.1. System architecture

The system architecture of a Funambol deployment is logically made of the components illustrated in Figure 5. For the sake of clarity, in the figure each component is represented by a single box, but as explained later, all systems can be redundant to increase availability and support more load.




Figure 5: Funambol system architecture

4.1.1. Roles and responsibilities

This section describes the role and the main responsibilities of the components illustrated in Figure 5.

Device

A device can be any physical device or client software application that can communicate with the Funambol server via SyncML for PIM synchronization or push email. Examples of such devices are:

- SyncML mobile phones (PIM and/or email)
- Supported J2ME phones (email)
- Funambol Outlook, iPod plug-ins
- Community plug-ins (Yahoo, Google, Evolution, ...)

Devices are the main interface through which users access Funambol. Responsibilities of the device/client software include:

- providing the user UI
- initiating the communication with the server
- hosting the local data (PIM/Email, ...)
- collecting/detecting the change log

The communication between the device and the Funambol server is based on the TCP/IP protocol.

Web load balancer

SyncML is an application protocol transported over HTTP. This means that SyncML requests can be considered common HTTP traffic. A web load balancer can therefore be used to balance the incoming load amongst different nodes of a Funambol server cluster.

The main responsibilities of the HTTP load balancer include:

- providing the front-end of the Funambol system
- distributing the SyncML requests amongst the nodes of the Funambol server cluster
- redirecting traffic to the active nodes if one of the nodes fails

Note that the HTTP load balancer is not provided as part of the default installation or deployment. Many different solutions, both hardware and software, can be adopted and any organization may have different best practices already in place. A common solution is to use Apache + mod_jk.

IP load balancer

Unlike SyncML, the protocol used by CTP (Client TCP Push) is not transported over HTTP but it is a pure TCP/IP communication. Since the number of connections that CTP must support is generally very high, an IP load balancer is usually needed. This component works similarly to the web load balancer, but at the Transport level (Layer-4 load balancing).

The main responsibilities of the IP load balancer include:

- providing the front-end of the Funambol system for CTP requests
- distributing CTP requests amongst the nodes of the Push server cluster
- detecting failures on the nodes of the cluster

Note that the IP load balancer is not provided as part of the default installation or deployment. Many different solutions, both hardware and software, can be adopted and any organization may have different best practices already in place. A common solution is to use Linux Virtual Server.

Funambol server

The Funambol server is the core of Funambol push-email. As illustrated in the figure, it is built of three main components: the DS server, the PIM connector and the email connector.

Role of the DS server is to provide the synchronization services and to communicate directly to the devices using SyncML. The main responsibilities of the DS server are:

- hosting the synchronization engine
- accepting and serving synchronization requests
- handling low level device information
- hosting and launching synclets
- providing an interface towards the back-ends
- providing a remote administration interface
- providing server TCP push (STP)

Email connector

The email connector plays a key role in the Funambol architecture, since it allows the server to communicate with different email servers (note that each user can be attached to a different email server). The email connector consists of two main components: the connector itself and the Inbox Listener.

The Inbox Listener is a separate process from the server and portal process; it has the following responsibilities:

- creating and maintaining the user inbox cache
- polling the user inbox regularly to check for new emails
- updating the user inbox cache
- triggering an action to the DS server if the user has new email

Note that the user inbox cache does not contain any sensitive information, but only information about when a message has been received. This is necessary in order to filter out the less recent emails during email download.

The connector module is deployed together with the synchronization engine. It has the following responsibilities:

- searching the user email cache for which messages should be downloaded on the client
- filtering out unwanted messages or content (e.g., retrieve headers only)
- processing emails
- filtering emails
- sending outgoing messages

PIM connector

This connector is the counterpart of the email connector for PIM synchronization. The PIM connector consists of two main components: the connector itself and the PIM Listener.

The PIM Listener is a separate process from the server and portal process; it has the following responsibilities:

- polling the user PIM database regularly to check for updates
- triggering an action to the DS server if the user has PIM changes

The connector module is deployed together with the synchronization engine. It has the following responsibilities:

- searching the user email cache to determine the messages to be downloaded onto the client
- using various techniques to enable the user to filter out unwanted emails (e.g. Retrieve and sync headers only)
- processing emails
- sending outgoing emails

Push server

The push server is a separate process from the synchronization server and portal process and is responsible for the implementation of the client TCP push technology. It is therefore also known as CTP server.

The main responsibilities of the push server are:

- accepting and keeping open CTP connections from devices
- delivering push notifications to the attached devices

SMTP server

This is the server used by Funambol Carrier Edition to send emails to external recipients.

Note: this server is used for service related emails (e.g. invitation or activation emails) and to send user emails for users who are not using a public email service.

Database

This is the database server. Funambol Carrier Edition supports the following database systems:

- PostgreSQL
- MySQL (requires Funambol v7.0 or later)
- Hypersonic

SMS Service

This is the service used to send SMS messages for the following purposes:

- providing the download link for the Windows Mobile and JavaME plug-ins
- OTA configuration of SyncML settings
- SMS push

Email server

This component is not really part of the Funambol software; it represents the users' email servers. Currently, the supported mail protocols are:

- IMAP
- POP
- Google IMAP
- AOL IMAP

4.2. The synchronization engine

The following sections provide more details on a particular component, the DS Server, and describe the architecture of the synchronization engine, which is a key component that can be extended to implement specific needs.

The synchronization engine is the component that implements the synchronization logic; this means:

- identify the sources and the destinations of the data sets to be synchronized
- identify the data that needs to be updated/added/deleted
- determine how updates must be applied
- detect conflicts
- resolve conflicts

In other words, the synchronization engine is the core of a data synchronization server.

The basic framework interfaces and classes are grouped in the package `sync4j.framework.engine`.

4.3. Execution flow of a request

The execution flow of an OMA DS request is illustrated in Figure 6.

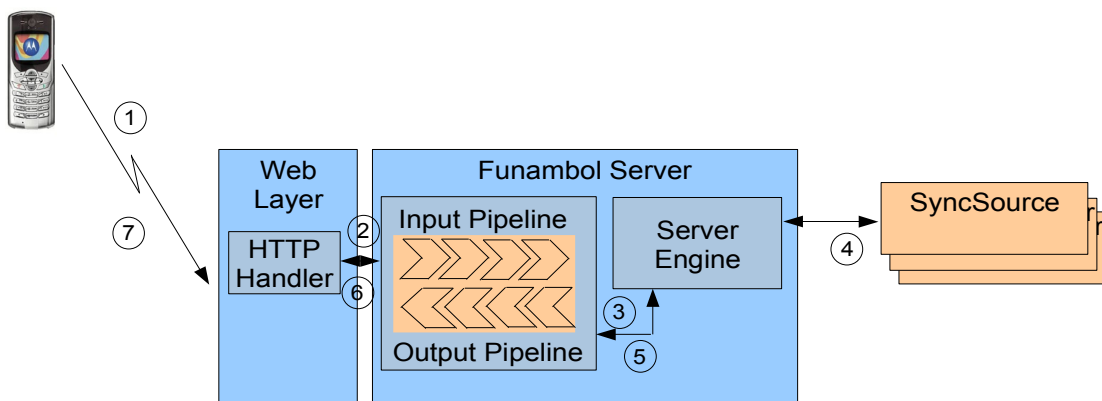


Figure 6 - Execution flow of an OMA DS request

A synchronization session starts with the client device sending a first SyncML message to the server. The request then follows the flow described below:

1. When a new request comes from the client, the HTTP handler takes care of it. After some processing, for example the transformation of the binary message into a more manageable form or the association of the incoming message to an existing synchronization session, the HTTP handler passes the request to the synchronization server.
2. The message first goes through the input message processing pipeline (see later) according to the application needs.
3. The manipulated message comes out from the input pipeline and goes into the server engine for the synchronization processing.
4. When needed, the server engine calls the services of the external (and custom) SyncSources in order to access the real data stores.
5. After processing the incoming message, the server engine builds the response message, which goes through the output message processing pipeline for post-processing.
6. The response message payload is then returned to the HTTP handler, which packs the SyncML message into a HTTP response and sends it back to the device.

5. The synchronization process

The synchronization process is logically accomplished in three steps:

1. Preparation
2. Synchronization
3. Finalization

The Funambol engine goes through these steps coordinating their execution, but delegates most of the synchronization logic to an auxiliary class, implementation of the *SyncStrategy* interface.

There are many types of synchronization; the ones specified by the SyncML protocol are:

<i>Sync Type</i>	<i>Description</i>
Two-way sync (fast)	A normal sync type in which the client and the server exchange information about modified data in these devices. The client sends the modifications first.
Slow sync	A form of two-way sync in which all items are compared with each other on a field-by-field basis. In practice, this means that the client sends all its data from a database to the server and the server does the sync analysis (field-by-field) for this data and the data in the server.
One-way sync from client only	A sync type in which the client sends its modifications to the server but the server does not send its modifications back to the client.
Refresh sync from client only	A sync type in which the client sends all its data from a database to the server (i.e., exports). The server is expected to replace all data in the target database with the data sent by the client.
One-way sync from server only	A sync type in which the client gets all modifications from the server but the client does not send its modifications to the server.
Refresh sync from server only	A sync type in which the server sends all its data from a database to the client. The client is expected to replace all data in the target database with the data sent by the server.
Server Alerted Sync ¹	A sync type in which the server alerts the client to perform sync. That is, the server informs the client to start a specific type of sync with the server.

Table 3 - Sync modes defined by SyncML

The first two are the most important, since the others are derivation of slow and fast sync modes.

In a slow synchronization, the client sends all its items to server, which compare them with the server database and then it sends back the modification that the client has to apply in order to be in sync again. In the case of slow sync, the sources to be synchronized must be fully compared in order to reconstruct the right image of the data on both synchronization endpoints. The way the sets of items are compared is implementation specific and can vary from comparing just the item keys or the entire content of an item.

In a two-way fast synchronization, a data source is queried only for new, deleted or updated items since a given point in time (and for a given user). In this case, the status (deleted/updated/new) and the modification timestamp of the items can be checked in order to decide when a deeper comparison is necessary.

The following sections describe in more detail each phase of the synchronization process and other key aspects of the synchronization engine architecture.

¹ The SyncML specification does not tell anything about how server alerted sync should be achieved, therefore each product can implement it in a different and not interoperable way. As per nowadays, only few devices are known to support this feature. Sync4j does not currently implement server alerted sync.

5.1. Preparation

The preparation phase is the process of analyzing the differences between two or more sources of data (called *SyncSources*) with the goal of obtaining a list of sync operations that applied to the sources involved in the synchronization, will make the databases look identical (Figure 7).




Figure 7- Preparation phase

5.2. Modifications detection

Modifications detection is based on the sets of items represented in Figure 8, applying the modifications matrix of Table 4. You can think of A as the client data source and B as the server data source.




Figure 8 - Synchronization items sets

- **A** – Items belonging to source A (as known via LUID-GUID mapping)
- **B** – Items belonging to source B
- **Am** – Modified items belonging to source A
- **Bm** – Modified items belonging to source B
- **AmBm** – Items modified both in source A and B (intersection between Am and Bm)
- **(A-Am)Bm** – Items unmodified in A, but modified in B
- **Am(B-Bm)** – items unmodified in B, but modified in A

Note that A is the server view of the A source: it contains the items mapped in the server as they are defined in the LUID-GUID mapping. If, for example, the client sends a new item that has never been mapped, this item will be in Am, but not in A. In order to be sure that the new item is not equal to some existing item in B, it must be looked up in B. If an item in B represents the same item as in Am, A is virtually augmented of such item, so that at the end, Am will be a sub-set of A.

Another important aspect to point out is that the entire data sets A and B can be considerably big. Therefore, when possible, it is important to deal with the smallest possible sets of items instead of doing a full item-per-item comparison.

The preparation phase is slightly different depending on the type of the synchronization. In the case of a slow synchronization, all items in the sources must be compared looking for differences that will be translated into synchronization operations. This type of process does not depend on previous synchronizations and, in fact, it is used to fully recreate a database as if no synchronizations have ever taken place. This is achieved resetting the LUID-GUID mapping before starting the modification detection process.

On the contrary, when a fast synchronization is performed, it is assumed that the involved sources rely on a previous data synchronization, so that only the changes since the time of the last synchronization need to be considered.

The algorithm used in the preparation phase is as follows:

Given the sources to be compared, suppose A and B, the goal of the algorithm is to produce an array of operations, in which each element represents a particular synchronization action, i.e. create the item X in the source A, delete the item Y from the source B, etc. Sometimes, it is not possible to decide the action to perform, thus a conflict operation is generated. A conflict might be solved by something external the synchronization process, for instance by a user action. In order to create the operation set, each item in the source A is compared with each item in the source B (to be intended as the selected items depending on the synchronization type).

To determine which operation should be generated the synchronization matrix defined in Table 4 is used.

<i>Database A</i> → ↓ <i>Database B</i>	<i>New</i>	<i>Deleted</i>	<i>Updated</i>	<i>Synchronized/Unchanged</i>	<i>Not Existing</i>
<i>New</i>	C	C	C	C	B
<i>Deleted</i>	C	X	C	D	X
<i>Updated</i>	C	C	C	B	B
<i>Synchronized/Unchanged</i>	C	D	A	=	B
<i>Not Existing</i>	A	X	A	A	X

Table 4 - Synchronization matrix

Where:

- **A** : item A replaces item B
- **B** : item B replaces item A
- **C** : conflict
- **D** : delete the target item
- **X** : do nothing

When a client item should be updated or inserted on the server, but the server does not have a mapping for it, a deeper comparison must be performed. In fact, the new/updated item could have the same content of an existing item on the server; this could even turn into a conflict. For example, suppose that a client tries to insert a new appointment with ID "xyz" at 20041029T1400Z in the meeting room "OceanSide". Even if there is no matching item on the server, if "OceanSide" is already busy at the same time, this could be considered a conflict.

In order to detect such situations, the synchronization engine will ask for items similar to the one that is trying to add/update. Those similar items are called twins. Note that we used by choice similar and not

equal. This is because how much an item should look like an existing item in order to be considered a twin may be implementation specific. Each source should be able to find twin items accordingly to its own logic.

5.3. Synchronization

The synchronization step is the phase where the operations prepared in the previous step are executed. Executing an operation means applying the required modification to the involved SyncSources. This is done by the synchronization engine.

5.4. Finalization

The third and last step is intended for cleaning up purposes. In addition, usually in this phase the client sends the LUID-GUID mapping resulted in the synchronization just performed.

6. Extending Funambol

The Funambol platform can be extended in many areas in order to integrate Funambol into existing systems and environments. Figure 9 illustrates the most common integration user cases, and the Funambol modules involved:

- Officer: integrating with an external authentication and authorization service;
- SyncSource: integrating with an external data source
- Synclet: adding pre or post processing to a SyncML message
- Admin WS: integrating with an external management tool

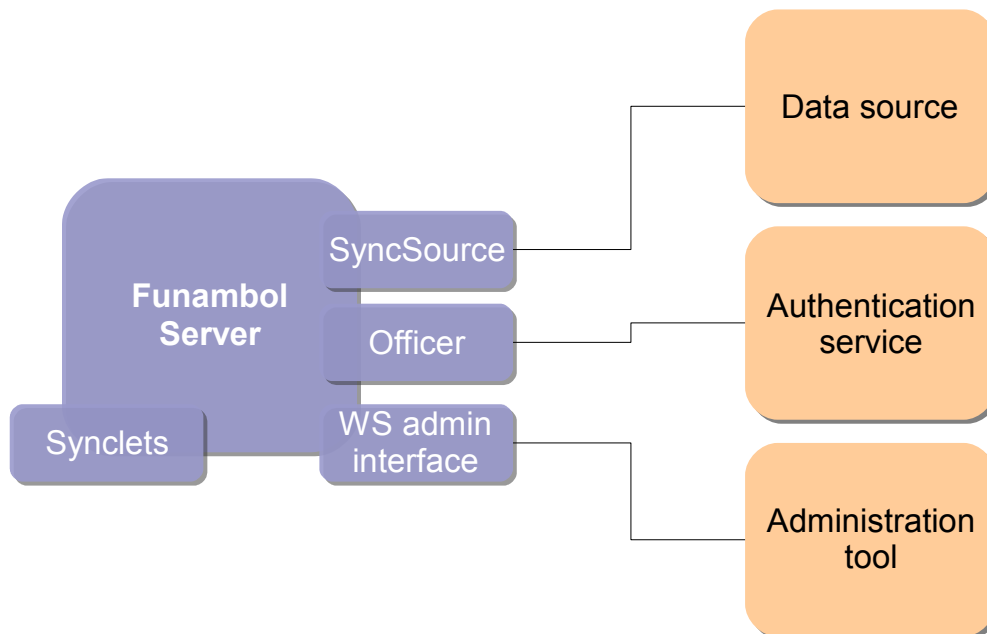


Figure 9: Integration components

Funambol extensions are distributed and deployed as Funambol modules. This section describes the structure of a Funambol module, while the following sections detail each of the scenarios above.

A Funambol module represents the means by which developers can extend the Funambol server. A module is a packaged set of files containing classes, installation scripts, configuration files, initialization SQL scripts, components and so on, used by the installation procedure to embed an extension into the server core.

For more information on how to install Funambol modules see [3].

6.1. Building a Funambol module

A Funambol module is a zip package named following the convention:

```
<module-name>-<major-version>.<minor-version>.s4j
```

Where *<module-name>* is the name of the module without spaces and with small caps only and *<major/minor-version>* are the major and minor version numbers. A new version of a module with a minor version number change must be backward-compatible, while changes in the major version number imply that a migration may be required.

The package must have the structure illustrated in Figure 10.




Figure 10 - Module package structure

In the following, entries ending with a '/' represent directories and filenames in *italic* are given just as examples (in a real package they will be replaced with real filenames).

The module classes are packaged in a main jar file called *<modulename>.jar*.

Configuration files are stored under the package directory *config*, creating subdirectories as needed.

Note: even if it is not mandatory, usually SyncSource instance configuration files are stored under a subtree in the form *<module-id>/<connector-id>/<sourcetype-id>*, which is the convention used by the Administration Tool when creating a new SyncSource instance.

The directory *install* contains *install.xml*, an Ant script called when the module is installed; this is the hook where a module developer can insert module specific installation tasks. Installation specific files can be organized in subdirectories under *install*. If the module requires a custom database schema, the scripts to create, drop and initialize the database are stored under the *sql/<database>* directory, where *<database>* is the name of the DBMS as listed in the *install.properties* file. Finally, the *exclude* directory is used to store files that will be temporarily used by the installation procedure, but that will not be (automatically) copied. These can be used by the module installation script for any purpose.

6.2. Modules, connectors, listeners and SyncSource types

As already mentioned, a module is a container for anything related to one or more server extensions which are used by the engine to communicate and integrate with external systems. These extensions are usually specific to the backend that must be integrated. A specific case of such extension is when the main purpose is to connect to an external data source, in which case the module is called *connector*. In other words, a connector is an extension of the server, intended to support the synchronization of a particular data source.

In order to access the data source, the connector must provide a so called "SyncSource type". A SyncSource type represents the template from which an instance of a SyncSource can be created. For example, the *FileSystemSyncSource* type made available by the Funambol is the means used by the server to synchronize data stored in the file system. However, it does not represent a particular directory to synchronize; to synchronize a specific directory (for instance */data/contacts*) a real SyncSource instance must be created and configured with the desired directory. You can think of a SyncSource type as a class and of a SyncSource as an instance.

An additional (but optional) component that a connector can provide is called *listener*. This component is in charge to detect changes in the backend so that the server can trigger a device to synchronize the changes.

Note: Funambol provides out-of-the-box a module called Foundation that contains basic functionalities, libraries and classes that all custom modules require. This module must not be removed by any Funambol installation.

6.2.1. Registering modules, connectors and SyncSource types

Modules, connectors and SyncSource types are registered filling the following database tables:

- *fnbl_module* for module information
- *fnbl_connector* for connector information
- *fnbl_sync_source_type* for SyncSource type information
- *fnbl_connector_source_type* for connector-SyncSource type associations
- *fnbl_module_connector* for module-connector association

Note: The last two tables are used to create the hierarchy module-connector-SyncSource type that you can see in the Administration Tool.

As an example, let's consider the foundation module registration. When Funambol is installed, the foundation module is installed too. It brings a connector called *FunambolFoundationConnector*, which, in turn, contains the SyncSource types *PIM Contact SyncSource*, *PIM Calendar SyncSource*, *FileSystem SyncSource* and *SIF SyncSource* (Figure 11).

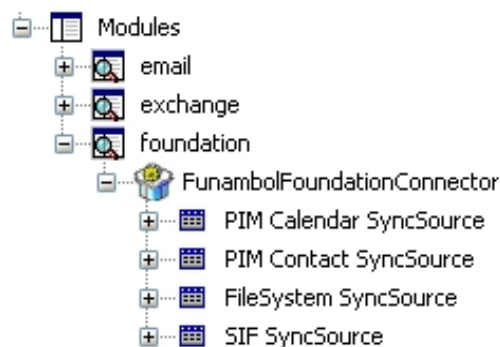


Figure 11: Foundation Connector Module in the Administration Tool

This hierarchy is obtained with the following SQL commands:

1. Module registration:

```
insert into fnbl_module (id, name, description)
values('foundation','foundation','Foundation');
```

2. SyncConnector registration:

```
insert into fnbl_connector(id, name, description)
values('foundation','FunambolFoundationConnector','Funambol Foundation Connector');
```

3. The Foundation Connector belongs to the foundation module:

```
insert into fnbl_module_connector(module, connector)
values('foundation','foundation');
```

4. The SyncSource Type registration:

```

insert into fnbl_sync_source_type(id, description, class, admin_class)
values('contact-foundation','PIM                                     Contact
SyncSource','com.funambol.foundation.engine.source.PIMContactSyncSource','com.funambol
.foundation.admin.PIMContactSyncSourceConfigPanel');

insert into fnbl_sync_source_type(id, description, class, admin_class)
values('calendar-foundation','PIM                                     Calendar
SyncSource','com.funambol.foundation.engine.source.PIMCalendarSyncSource','com.funambo
l.foundation.admin.PIMCalendarSyncSourceConfigPanel');

insert into fnbl_sync_source_type(id, description, class, admin_class)
values('fs-foundation','FileSystem
SyncSource','com.funambol.foundation.engine.source.FileSystemSyncSource','com.funambol
.foundation.admin.FileSystemSyncSourceConfigPanel');

insert into fnbl_sync_source_type(id, description, class, admin_class)
values('sif-fs-foundation','SIF
SyncSource','com.funambol.foundation.engine.source.SIFSynchronSource','com.funambol.founda
tion.admin.SIFSynchronSourceConfigPanel');

```

5. Finally, the SyncSource type belongs to the Foundation Connector:

```

insert into fnbl_connector_source_type(connector, sourcetype)
values('foundation','contact-foundation');

insert into fnbl_connector_source_type(connector, sourcetype)
values('foundation','calendar-foundation');

insert into fnbl_connector_source_type(connector, sourcetype)
values('foundation','fs-foundation');

insert into fnbl_connector_source_type(connector, sourcetype)
values('foundation','sif-fs-foundation');

```

Note: two classes are specified for each SyncSource type registration: the class (for instance *com.funambol.foundation.engine.source.FileSystemSyncSource*), which actually implements the SyncSource interface and the *admin_class*, which instead is used to create a new SyncSource instance and to configure it in the Administration tool.

In the following section, we will see how these two classes are developed.

Note: in this guide, SyncSource and SyncSource Type are often treated as synonyms, even if they are in the template-instance relationship seen before.

7. Developing a SyncSource

A SyncSource is the means a set of data is made available to the synchronization engine. Therefore, in order to synchronize any type of data (files, database tables, calendar events and so on), there must be a proper SyncSource able to extract and store the data from and to the real data store.

Goal of the Funambol platform is to provide a collection of SyncSources for the most common uses (i.e. files), but new SyncSources can be independently developed and plugged in the synchronization engine so that the server will be able to process synchronization requests targeted to virtually any data source.

7.1. The SyncSource interface and related classes

The core of the SyncSource architecture is the interface `sync4j.framework.engine.source.SyncSource`. This interface does not make any assumption on the type of data being synchronized, so that its concrete implementations are completely free to access their own underlying storage.

A SyncSource is identified by a *sourceURI* and usually a *name*; the former is the URI that a SyncML client must specify as target in order to synchronize this particular SyncSource; the latter is a human-friendly name used for displaying purposes only. Note that they must be both unique within a Funambol installation.

A SyncSource is also associated to a type, in the form of a MIME type that represents the type of data handled by the source.

The most important methods defined by the SyncSource interface are: (see chapter 11 for details)

<i>Method</i>	<i>Description</i>
<code>beginSync()</code>	This is the first SyncSource method the sync engine calls and it is used to specify who is going to synchronize and which type of synchronization is requested.
<code>endSync()</code>	This is the latest SyncSource method the sync engine calls and it may be used to perform finalization tasks.
<code>getUpdatedSyncItems</code>	Called to retrieve the updated SyncItems for the given principal since the given point in time. For a definition of what a Principal is, please see Section 7.1.1 below.
<code>getUpdatedSyncItemKeys</code>	Called to retrieve the SyncItemKey of the updated items for the given principal since the given point in time.
<code>getNewSyncItems</code>	Called to retrieve the new SyncItems for the given principal since the given point in time.
<code>getNewSyncItemKeys</code>	Called to retrieve the SyncItemKey of the new items for the given principal since the given point in time.
<code>getDeletedSyncItems</code>	Called to retrieve the deleted SyncItems for the given principal since the given point in time.
<code>getDeletedSyncItemKeys</code>	Called to retrieve the SyncItemKey of the deleted items for the given principal since the given point in time.
<code>getAllSyncItems</code>	Called to retrieve all the SyncItems for the given principal since the given point in time.
<code>setSyncItem</code>	Called to insert or update the given item.
<code>setSyncItems</code>	Called to insert or update the given items.
<code>removeSyncItem/s</code>	Called to remove the given item(s).
<code>getSyncItemFromTwin</code>	Called to find items that represent the same information as the given item. It is used in conflict detection and during slow sync to associate a client item with a server item which is similar enough.

Table 5 - SyncSource: most important methods

Funambol provides others two subtypes of the *SyncSource* interface: *MergeableSyncSource* and *FilterableSyncSource*.

com.funambol.framework.engine.source.MergeableSyncSource: this interface should be used when you want that conflicting data are merged when possible. The merge is performed when a conflict is

detected in order to avoid loss of information. For example, let's assume that the same contact has been defined on both the client adding a mobile phone, and the server, adding an email address. At the next sync a conflict will be detected. If the contact is generated by a *MergeableSyncSource*, the sync engine asks the SyncSource to merge the conflicting items. In the example, the result of the merge will be a contact with both the new mobile number and email address. This contact will be stored on the server and sent back to the client.

<i>Method</i>	<i>Description</i>
mergeSyncItem(serverKey: SyncItemKey, clientItem: SyncItem): boolean	Called when a conflict must be resolved merging the items. Server side, the merge result must be persistent in the underlying datastore. If the item on the client must be updated, this method must return true and put the new content in the given clientItem.

com.funambol.framework.engine.source.FilterableSyncSource: this interface should be used when the SyncSource supports SyncML 1.2 filtering.

<i>Method</i>	<i>Description</i>
getSyncItemStateFromId(itemKey: SyncItemKey): char	Called to retrieve the state of an item. This method is required because the server cannot know directly the state of an item not inside the filter criteria.
isSyncItemInFilterClause(item: SyncItem): boolean	Called to know if an item satisfies the filter.
isSyncItemInFilterClause(itemKey: SyncItemKey): boolean	Called to know if the item specified with the given key satisfies the filter.

7.1.1. SyncContext

The *begin()* method of SyncSource is called when a new sync is started. It requires just a *SyncContext* as input parameter. A *SyncContext* contains the following information:

- *principal* (of type *com.funambol.framework.security.Principal*)
- *syncMode* (of type int)
- *filter* (of type *com.funambol.framework.filter.Filter*)
- *since* (of type java.sql.Timestamp)
- *to* (of type java.sql.Timestamp)

A *principal* is composed of a *Sync4jUser* and a *Sync4jDevice* because the same user may use different devices (or the same device could be used by different users).

syncMode represents the synchronization type performed. Can be one of:

<i>syncMode</i>	<i>synchronization type</i>
200	TWO
201	SLOW
202	ONE_WAY_FROM_CLIENT
203	REFRESH_FROM_CLIENT
204	ONE_WAY_FROM_SERVER
205	REFRESH_FROM_SERVER

filter represents the filter required by the client. If a filter is specified and the SyncSource is a *FilterableSyncSource*, the expected behavior of the SyncSource is:

- the methods *getAllSyncItemKey()*, *getNewSyncItemKeys()*, *getDeletedSyncItemKeys()* and *getUpdatedSyncItemKeys()* return only the items inside the filter
- the methods *getSyncItemKeysFromTwin*, *getSyncItemStateFromId* ignore the filter.

Since is the last successfully completed synchronization session start time.

to represents the current synchronization start time.

7.1.2. SyncItem

Items returned by a SyncSource are encapsulated in *funambol.framework.engine.SyncItem* objects. *SyncItem* defines the following methods:

<i>Method</i>	<i>Description</i>
getKey	Returns the item key.
getParentKey	Returns the item's parent key (hierarchic sync)
getState	Returns the item state.
setState	Sets the item state.
getContent	Returns the item content
setContent	Sets the item content
getFormat	Returns the item format
setFormat	Sets the item format
getType	Returns the item type
setType	Sets the item type
getTimestamp	Returns the timestamp of the last change of the item state and it is used in the synchronization process, in order to determine the operation to be performed on the sources.
setTimestamp	Sets the item timestamp
getSyncSource	Returns the SyncSource the item belongs to.

Table 6 - SyncItem methods

Note: when a SyncSource creates a *SyncItem*, it must always provide a value for the content and for the timestamp.

7.1.3. Twin items

An item X is a twin of an item Y when from the SyncSource point of view, X and Y represent the same information.

For example, two event items, both at the same time and in the same place may be considered the same appointment, even if the associated note is different. Or two contacts may be considered the same person if they have the same first, middle and last name.

Because the SyncSource is the only entity with knowledge about how data are stored and represented in the data source, the SyncSource is the only component that can select the twins of a given item. The synchronization engine does it calling *getSyncItemsFromTwin()*.

Twin items are necessary in two circumstances:

- during full sync
- during conflict detection

During full sync, the client sends all its items and the server has to discover which operation the client should apply in order to make its data set look identical to the one stored on the server. Because the two devices are supposed to be out of sync, the server cannot rely on LUID-GUID mappings. In this case, for each item given by the client, the server must search for twin items in the backend data source. If a twin is not found, the new item is added. If a twin is found, the server will consider to have such client item and won't send back a command for it. If the SyncSource being synchronized is mergeable SyncSource, the synchronization engine will ask the SyncSource instance to merge the client and server items and will send back the merged item to the client.

The same process is valid during conflict detection. When during fast sync a client sends a new item, the server should first check if this item is conflicting with something else. Therefore, the server calls

the SyncSource's `getSyncItemsFromTwin()`. If this method returns something, a conflict is detected and handled accordingly. If no twin is found the item can be added.

The SyncSource developer should apply the more appropriated comparison logic according to the type of data the SyncSource manages. Optionally, the SyncSource can disable this type of conflict with no impact on the calling module, by implementing this method as a stub and always returning a failure to find a twin item.

One final note about twin computation is that it should be kept as simple as possible, since it involves database searches that may have a big impact on performance.

7.1.4. The Administration Tool configuration panel

The Administration Tool makes it possible to perform many administration tasks, including the configuration of the Funambol DS Server, loggers and SyncSources (for more information, please see [3]).

After logging into the Administration Tool and connecting to the server, users will see something similar to the window in Figure 12. Selecting an existing SyncSource, a configuration panel is displayed on the right side of the window.

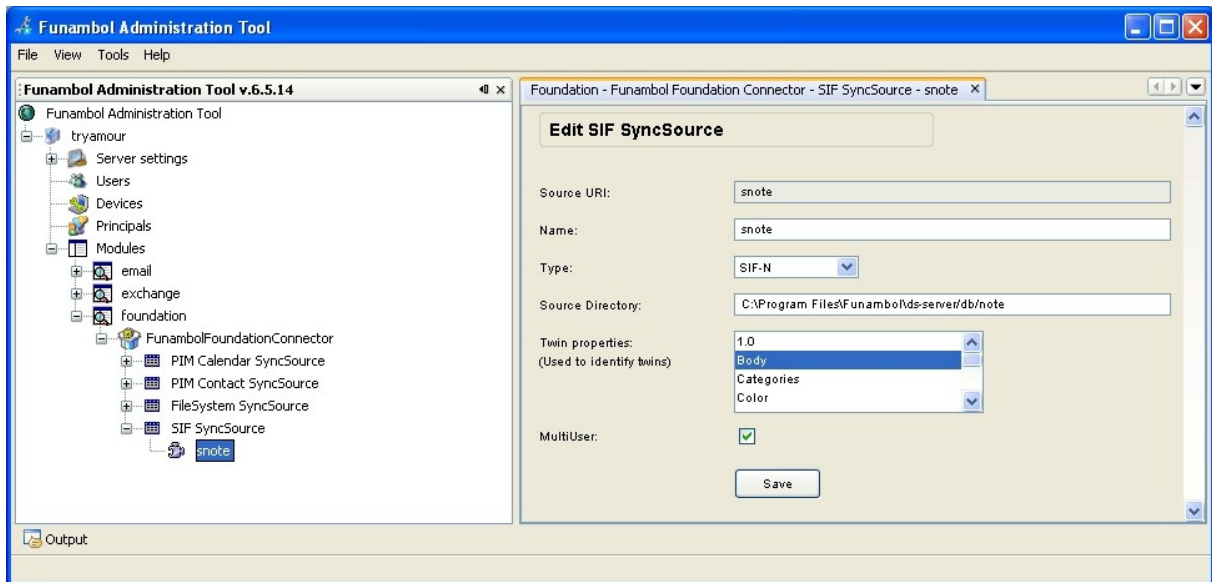


Figure 12: SyncAdmin showing connectors, modules and SyncSources

The Administration Tool is designed to be extended by developers so that a new custom SyncSource type can be configured into the tool through a custom management panel. The Administration Tool extension mechanism is described in the following chapter.

8. Extending the Funambol Administration Tool

The Funambol Administration Tool can be extended so that any kind of object can be configured through the UI. The extension mechanism is based on the concept that a module developer should be able to provide custom panels that an administrator can access in order to configure a specific server bean. Currently, a developer can provide management panels for the following objects:

- SyncSource types
- Connectors

Note: the classes implementing such management panels will be installed on the server as part of the module installation.

8.1. Architecture overview

The Administration Tool extension mechanism is based on the concept of *management object*. A management object is potentially anything that can be configured; however, in this guide the focus is on providing a way for a module developer to provide custom configuration panels for connectors and SyncSource types.

A management object is represented by the *com.funambol.syncadmin.mo.ManagementObject* class, which wraps a generic Object instance together with a management path (i.e. the path of the server bean instance in the configuration path). “Configuring a management object” means mainly setting its properties based on the input from a user. To allow this, the abstract class *com.funambol.syncadmin.ui.ManagementObjectPanel* abstracts a JPanel whose role is to configure a given management object.

Two specialized subclasses of *ManagementObjectPanel* are available to developers to configure SyncSource types and connectors. These are *SourceManagementPanel* and *ConnectorManagementPanel*. These two classes are still meant not to be directly instantiated and are therefore abstract.

A module developer can provide a configuration panel to a SyncSource type or a connector just extending *SourceManagementPanel* or *ConnectorManagementPanel* with a class showing a suitable data entry form for the object. When the module is loaded by the admin, those panel classes are downloaded from the server and instantiated into the Administration Tool when the user selects a SyncSource type or a connector.

The user can perform the following high-level actions on a custom panel:

- Adding a new object
- Updating an object
- Deleting an object

Such actions are notified to the interested classes in the same way swing events are commonly notified. Figure 13 illustrates the sequence diagram of the calls between different components of the Administration Tool or server when a user performs different operations on the connector / SyncSource type nodes.

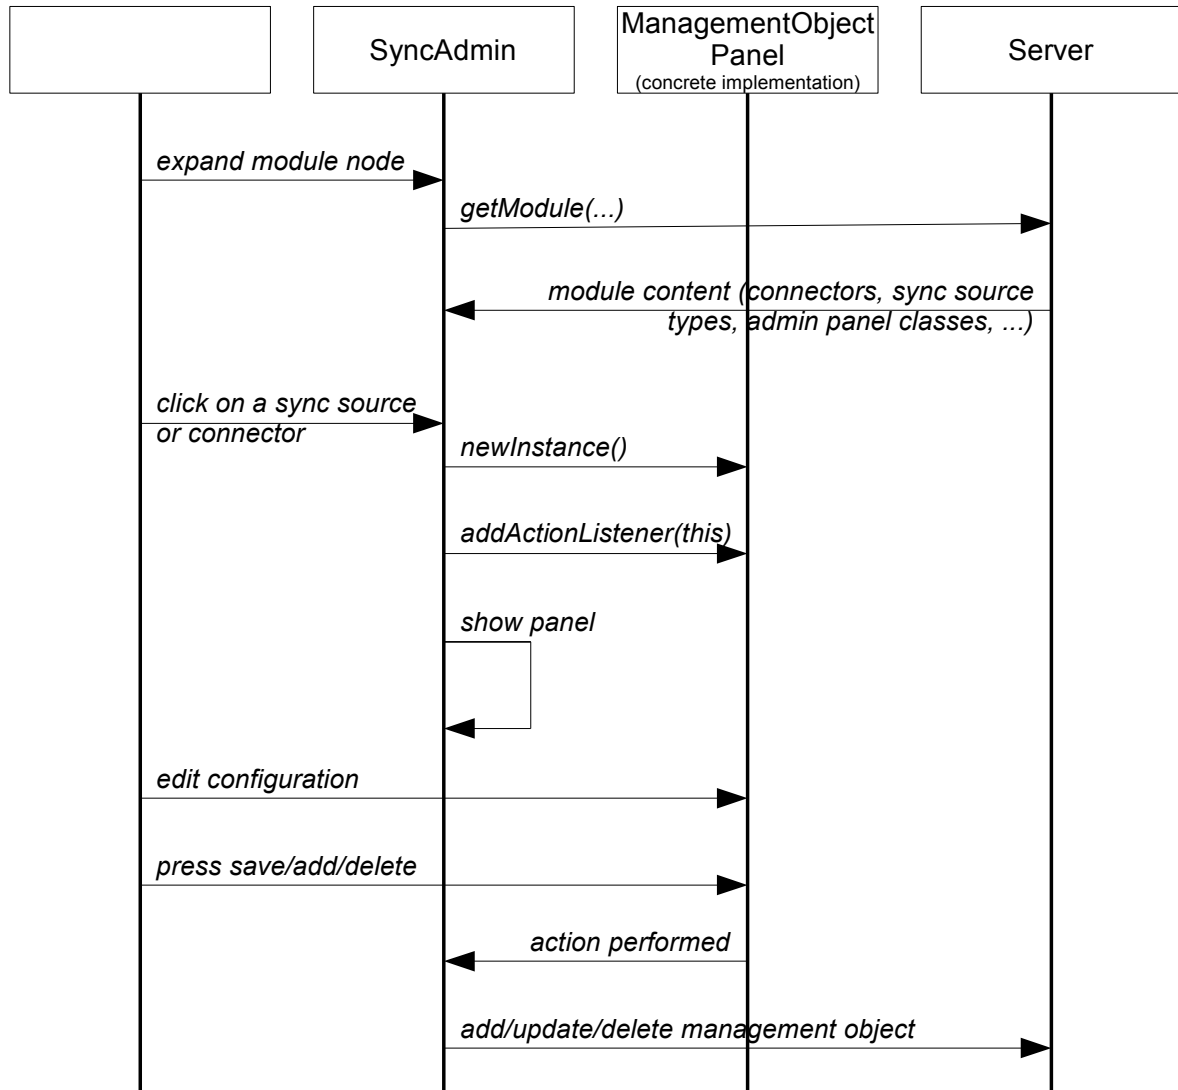


Figure 13 - Admin-ManagementObjectPanel collaboration diagram

8.2. ManagementObject and subclasses

The key concept around extending the Funambol Administration Tool with custom configuration panels is the *management object*. A management object can be any object we want to be able to manipulate through of the Administration Tool. In particular, any server beans can be considered a management object.

A management object is an association between a Java object instance and a management path; the management path represents the name of such instance in the context of a tree-like configuration repository.

In the Admin framework a management object is represented by the class `com.funambol.admin.mo.ManagementObject`. This has the interface of the table below.

Method	Description
Constructors	
ManagementObject (Object object, String path)	Creates a new management object from the instance and the instance pathname
Public methods	
Object getObject()	Returns the object instance

<i>Method</i>	<i>Description</i>
void setObject(Object o)	Sets the object instance
String getPath()	Returns the management object pathname
void setPath(String path)	Sets the management object pathname

Two specializations of *ManagementObject* are provided as described in the following: *SyncSourceManagementObject* and *ConnectorManagementObject*. These are better described in the following sections.

8.2.1. com.funambol.admin.mo.SyncSourceManagementObject

This management object represents a SyncSource instance. It associates the instance to the module / connector / SyncSourceType IDs.

The interface of this class is shown in the table below.

<i>Method</i>	<i>Description</i>
Constructors	
SyncSourceManagementObject(SyncSource source, String moduleId, String connectorId, String sourceTypeId)	Creates a new SyncSource management object from the instance and its module/connector/type IDs
Public methods	
String getModuleId()	Returns the module ID
void setModuleId(String ID)	Sets the module ID
String getConnectorId()	Returns the connector ID
void setConnectorId(String ID)	Sets the connector ID
String getSourceTypeId()	Returns the source type ID
void setSourceTypeId(String ID)	Sets the source type ID
String getTransformationsRequired	Returns the engine transformations required (for instance 'b64')
setTransformationsRequired(String transformations)	Sets the required transformations

8.2.2. com.funambol.admin.mo.ConnectorManagementObject

This management object represents a connector configuration object. It associates the instance to the module / connector IDs. Plus, the management path is created as:

```
moduleId + '/' + connectorId + '/' + connectorName + ".xml"
```

The interface of this class is shown in the table below.

<i>Method</i>	<i>Description</i>
Constructors	
ConnectorManagementObject(Object obj, String moduleId, String connectorId, String connectorName)	Creates a new connector management object from the instance and its module/connector/type IDs. obj represents the configuration object, retrieved from the server as a server bean with the pathname generated as above.
Public methods	
String getModuleId()	Returns the module ID
void setModuleId(String ID)	Sets the module ID
String getConnectorId()	Returns the connector ID

<i>Method</i>	<i>Description</i>
void setConnectorId(String ID)	Sets the connector ID

8.3. ManagementObjectPanel and subclasses

ManagementObjectPanel is an abstract class that provides no user interaction widgets. It must be extended by concrete implementations in order to do something useful.

To notify actions to the framework, *ManagementObjectPanel* is a *java.awt.event.Action* subject (referring to the *Observer* pattern). This means that any class interested in knowing about (observing) what is happening in a management panel shall implement the *java.awt.ActionListener* interface and register itself to the panel, calling *addActionListener()*. In the Admin framework such interested classes are represented by the controllers (*SyncSourcesController*, *ConnectorController*, and so on).

The *ActionEvent* generated and notified by *actionPerformed()* to all listeners shall have the following characteristics:

- the panel instance as source
- one of the following event IDs:
 1. ACTION_EVENT_INSERT
 2. ACTION_EVENT_UPDATE
 3. ACTION_EVENT_DELETE
- an arbitrary value as command string

For example, a custom management panel to configure a connector may have a “save” button, through which the user can persist changes. When the button is pressed, the panel should notify the action to the listeners. This can be done with a code similar to the following:

```

public class ConfigPanel
extends ManagementObjectPanel {
...

public ConfigPanel() {
... widgets creation and initialization ...

saveButton.addActionListener(new ActionListener() {
public void actionPerformed(ActionEvent event) {
try {
... values validation ...
... management object update ....

ConfigPanel.this.actionPerformed(
new ActionEvent(ConfigPanel.this, ACTION_EVENT_UPDATE,
event.getActionCommand())
);
} catch (Exception e) {
notifyError(new AdminException(e.getMessage()));
}
}
});
...
}

```

ManagementObjectPanel has the following interface.

<i>Method</i>	<i>Description</i>
Abstract methods	
void updateForm()	Tells the panel that it has to update the UI with the values in the management object.
Public methods	
void setManagementObject(ManagementObject)	Sets the object under editing by this management panel

<i>Method</i>	<i>Description</i>
ManagementObject getManagementObject()	Returns the edited object
void addActionListener(ActionListener)	Registers the given action listener
void removeActionListener(ActionListener)	Unregisters the given action listener
void notifyError(AdminException)	Called when an error related to the panel should be displayed
<i>Protected methods</i>	
void actionPerformed(ActionEvent)	Notifies an action event to all listeners.

Two subclasses of a *ManagementObjectPanel* are provided by the Funambol administration framework. They are described in the following sections.

8.3.1. SourceManagementPanel

This class specializes a *ManagementObjectPanel* to a particular type of management objects: SyncSources. This base class is designed to be extended by concrete implementations to configure new and existing SyncSources. The additions to the base class are:

- it handles management objects of type *com.funambol.framework.engine.source.SyncSource*
- it defines two states for the management panel: *insert* and *update*

The former is achieved providing a *getSyncSource()* methods that returns the SyncSource instance under editing. The latter is used to differentiate the case where the user is creating a new SyncSource (therefore the SyncSource instance under editing does not exist yet on the server) from the case where the user is editing an existing SyncSource. The administration framework will take care of changing the status from INSERT to UPDATE.

SourceManagementPanel has the following interface.

<i>Method</i>	<i>Description</i>
<i>Public methods</i>	
void setState(int)	Sets the panel state.
int getState()	Returns the panel state
SyncSource getSyncSource()	Returns the SyncSource instance under editing

8.3.2. ConnectorManagementPanel

This class represents the base class for the management panel of a connector. It does not adds much the its base class, but it just provides a shortcut to easily gather the server bean instance used to configure the connector.

ConnectorManagementPanel has the following interface.

<i>Method</i>	<i>Description</i>
<i>Public methods</i>	
Object getConfiguration()	This is a shortcut to getManagementObject().getObject().

9. Configuring Funambol components

One of the Funambol design goal is to provide a framework that can be used to implement any kind of synchronization service, extending existing modules or plugging in new modules. All this require a lot of configuration information and possibly an easy way to add configuration for new extensions. Configuration files should be easy to understand, access and change.

Funambol uses mainly two configuration techniques:

- System properties
- Server JavaBeans

In the following sections these two types of configuration are described in details.

9.1. System properties

The only system property directory used by the Funambol server is *funambol.home* which must point to the directory where Funambol is installed (commonly referenced as *\$FUNAMBOL_HOME*).

This property is specified at JVM invocation time using the *-D* option. On many systems, it is sufficient to set the *JAVA_OPTS* environment variable in order to get it included into the JVM launching command.

9.2. Server JavaBeans

Many Funambol components are configured as *server JavaBeans*. Server JavaBeans are JavaBeans used server-side. The idea is to store a bean configuration as the serialized form of the bean itself. This way, a bean can be instantiated, configured and serialized to persist its configuration. Later, the bean can be deserialized in memory as a properly configured instance.

Funambol makes use of the standard Java facility to serialize objects into XML (and to deserialize them from XML). This is achieved by using the classes *java.beans.XMLEncoder* and *java.beans.XMLDecoder*. Since configuration files created with such encoder/decoder are easy to use, read and write, they can be created and modified manually with a simple text editor, without the need of a dedicated GUI. An additional advantage of this approach is that server JavaBeans are not requested to implement *java.io.Serializable* because *XMLEncoder* does not require it.

This is an example of a server JavaBean:

```
<?xml version="1.0" encoding="UTF-8"?>
<java version="1.4.1_01" class="java.beans.XMLDecoder">
  <object class="sync4j.framework.server.store.PersistentStoreManager">
    <void property="jndiDataSourceName">
      <string>java:/jdbc/sync4j</string>
    </void>
    <void property="stores">
      <array class="java.lang.String" length="2">
        <void index="0">
          <string>sync4j.server.store.SyncPersistentStore</string>
        </void>
        <void index="1">
          <string>sync4j.server.store.EnginePersistentStore</string>
        </void>
      </array>
    </object>
  </java>
```

```
</void>
</object>
</java>
```

In order to help server JavaBeans handling, Funambol uses the factory class *com.funambol.framework.tools.beans.BeanFactory*, which in turn makes use of a customized class loader; the class loader handles configuration files in a so called *config path*, in the same way a common class loader handles classes in the classpath.

The XML syntax uses the following conventions:

Each element represents a method call.

- The "object" tag denotes an expression whose value is to be used as the argument to the enclosing element.
- The "void" tag denotes a statement which will be executed, but whose result will not be used as an argument to the enclosing method.
- Elements which contain elements use those elements as arguments, unless they have the tag: "void".
- The name of the method is denoted by the "method" attribute.
- XML's standard "id" and "idref" attributes are used to make references to previous expressions - so as to deal with circularities in the object graph.
- The "class" attribute is used to specify the target of a static method or constructor explicitly; its value being the fully qualified name of the class.
- Elements with the "void" tag are executed using the outer context as the target if no target is defined by a "class" attribute.
- Java's String class is treated specially and is written `<string>Hello, world</string>` where the characters of the string are converted to bytes using the UTF-8 character encoding.

Although all object graphs may be written using just these three tags, the following definitions are included so that common data structures can be expressed more concisely:

- The default method name is "new".
- A reference to a Java class is written in the form `<class>javax.swing.JButton</class>`.
- Instances of the wrapper classes for Java's primitive types are written using the name of the primitive type as the tag. For example, an instance of the Integer class could be written: `<int>123</int>`. Java's reflection is internally used for the conversion between Java's primitive types and their associated "wrapper classes".
- In an element representing a nullary method whose name starts with "get", the "method" attribute is replaced with a "property" attribute whose value is given by removing the "get" prefix and decapitalizing the result.
- In an element representing a monadic method whose name starts with "set", the "method" attribute is replaced with a "property" attribute whose value is given by removing the "set" prefix and decapitalizing the result.
- In an element representing a method named "get" taking one integer argument, the "method" attribute is replaced with an "index" attribute whose value the value of the first argument.
- In an element representing a method named "set" taking two arguments, the first of which is an integer, the "method" attribute is replaced with an "index" attribute whose value the value of the first argument.
- A reference to an array is written using the "array" tag. The "class" and "length" attributes specify the sub-type of the array and its length respectively.

9.2.1. The configuration path

Server JavaBeans are looked for in the configuration path, which is analogous to the class path for classes lookup. This is implemented reading the serialization files from a custom class loader, *com.funambol.framework.config.ConfigClassLoader*. This class loader (which is instead our server beans loader), is configured to read objects from the configuration path. The config path is built appending */config* to the *funambol.home* system property value. For example, if *funambol.home* is set to */opt/Funambol/ds-server*, the config path would be */opt/Funambol/ds-server/config*.

9.2.2. Lazy initialization

When a server bean is deserialized from its XML form, the classloader that loads the serialization file calls the empty constructor first and then it sets the bean property values using the *setXXX()* methods provided by the class. However, some classes may need additional operations to be performed in order to properly initialize (after *setXXX()* methods are called). To support this *lazy initialization* approach, these classes can implement *com.funambol.framework.tools.beans.LazyInitBean*, which defines a separate *init()* method. When Funambol loads a *LazyInitBean*, after bean instantiation (or deserialization) and configuration (calling the setter methods), it calls the bean's *init()* method, giving the bean the opportunity to complete its initialization.

9.3. How to configure a standard component

Making a change to a configuration bean is as easy as editing a text file. Let's take as example the configuration file for the *DBOfficer* component. The configuration bean full path is *com/funambol/server/security/DBOfficer.xml* (please note that this path is relative to the config path) and its content is below:

```
<?xml version="1.0" encoding="UTF-8"?>
<java version="1.4.0" class="java.beans.XMLDecoder">
  <object class="sync4j.server.security.DBOfficer">
    <void property="clientAuth">
      <string>syncml:auth-basic</string>
    </void>
    <void property="serverAuth">
      <string>none</string>
    </void>
  </object>
</java>
```

The object element specifies which Java class will be instantiated and the property element sets the corresponding instance property. Therefore, to change the preferred client authentication type, it is sufficient to edit the file, change the *clientAuth* property and save. The next time this bean will be used, the new configuration value will be picked up.

9.4. How to create a custom configurable object

With this technique, any Java object can be configured, from a simple Java class to a very complex Java object tree. For example, this configures a String object:

```
<?xml version="1.0" encoding="UTF-8"?>
<java version="1.4.2" class="java.beans.XMLDecoder">
  <string>This is a String!</string>
</java>
```

A more interesting example is given, for instance, by the class *com.funambol.framework.config.LoggingConfiguration*. The class looks like the following:

```
public class LoggingConfiguration {
    // ----- Private data
    private ArrayList loggers;
    // ----- Constructors

    /** Creates a new instance of LoggingConfiguration */
    public LoggingConfiguration() {
    }

    /**
     * Getter for property loggers.
     */
}
```

```

    *
    * @return Value of property loggers.
    */
    public ArrayList getLoggers() {
        return loggers;
    }

    /**
     * Setter for property loggers.
     *
     * @param loggers New value of property loggers.
     */
    public void setLoggers(ArrayList loggers) {
        this.loggers = loggers;
    }
}

```

A possible configuration file for such a class could be:

```

<?xml version="1.0" encoding="UTF-8"?>
<java version="1.4.2_04" class="java.beans.XMLDecoder">
  <object class="sync4j.framework.config.LoggingConfiguration">
    <void property="loggers">
      <object class="java.util.ArrayList">
        <!--
          funambol
        -->
        <void method="add">
          <object class="com.funambol.framework.config.LoggerConfiguration">
            <void property="append">
              <boolean>true</boolean>
            </void>
            <void property="count">
              <int>1</int>
            </void>
            <void property="fileOutput">
              <boolean>true</boolean>
            </void>
            <void property="level">
              <string>INFO</string>
            </void>
            <void property="limit">
              <int>100</int>
            </void>
            <void property="name">
              <string>funambol</string>
            </void>
            <void property="pattern">
              <string>logs/ds-server.log</string>
            </void>
          </object>
        </void>
      </object>
    </void>

    <!--
      funambol.engine
    -->
    <void method="add">
      <object class="com.funambol.framework.config.LoggerConfiguration">
        <void property="append">
          <boolean>true</boolean>
        </void>
        <void property="count">
          <int>1</int>
        </void>
        <void property="inherit">
          <boolean>true</boolean>
        </void>
        <void property="level">
          <string>INFO</string>
        </void>
        <void property="limit">
          <int>100</int>
        </void>
        <void property="name">
          <string>funambol.engine</string>
        </void>
        <void property="pattern">
          <string>logs/ds-server.engine.log</string>
        </void>
      </object>
    </void>
  </object>

```

```
</void>
</object>
</void>
</object>
</java>
```

Note: see later how to automatically create such a file.

9.5. How to get a configured instance

Configuration beans are accessed through the singleton *com.funambol.framework.config.Configuration* object. For example, to instantiate a configured *LoggingConfiguration* instance, use the code below.

```
Configuration c = Configuration.getConfiguration();
LoggingConfiguration logging =
c.getBeanInstanceByName("sync4j/server/logging/logging.xml");
```

9.5.1. Tips and tricks

It is not necessary to write a configuration file by hand from scratch. To write a bean instance for the first time use the *com.funambol.framework.tools.beans.BeanFactory*'s *saveBeanInstance()* method to save a configured instance into a file. For example:

```
JButton b = new JButton("press me");
BeanFactory.saveBeanInstance(b, new File("button.xml");
```

The result is the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<java version="1.4.2_04" class="java.beans.XMLDecoder">
  <object class="javax.swing.JButton">
    <string>press me</string>
  </object>
</java>
```

10. Customizing message processing

This section explains how to extend Funambol customizing the processing of incoming and outgoing messages.

10.1. Overview

The OMA DS protocol is an XML-based client-server protocol. This means that each OMA DS message is an XML document and a response message always follows a request message.

In the Funambol implementation, an OMA DS message reaching the server goes through some transformations. These may be:

- XML level transformations, processing the message in its native XML representation, and
- message transformations, acting on a Java representation of the message

Figure 14 helps explaining this process.

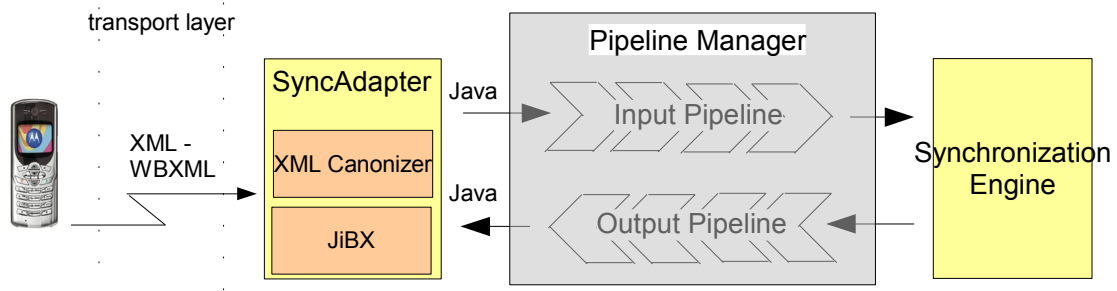


Figure 14 - Message processing architecture

In order to save bandwidth and processing power, OMA DS messages can be also WBXML encoded. Regardless of how the message is encoded, its content is first delivered to a *SyncAdapter* component by the transport layer. The *SyncAdapter* first transforms the message into XML if it was WBXML encoded and then the XML message is reduced to a “canonical” form in order to get rid of device specific singularities. XML canonization is the standard XML level transformation executed by the system.

Even when in the canonical XML form, the message is still hard to manipulate, since XML needs to be parsed. Plus, each component that needs to access any of the OMA DM message elements would have to parse the XML again, with a big impact on performance. For these reasons, the canonic XML message is transformed into an object tree that represents exactly the message.

After an incoming message has been transformed into an object tree, it passes through the input message processing pipeline before it gets to the synchronization engine. This gives the opportunity of further processing the message when it is in a more manageable representation. In a similar way, a response message going out from the engine, passes through the output message processing pipeline before getting transformed to its XML (and then WBXML) representation.

The input and the output pipelines are completely customizable, so that custom message pre and post processing can be easily added to the system.

Input and output message processing components are also called *synclets*.

10.2. Preprocessing an incoming message

To preprocess an incoming message we have to create an input processor component (input synclet) and to configure the pipeline manager accordingly. This is described below..

10.2.1. Creating an input synclet

An input synclet is a class that implements the *com.funambol.framework.engine.pipeline.InputMessageProcessor* interface. This interface defines just one method: *preProcessMessage(MessageProcessingContext context, SyncML msg)*. context is a parameter that is shared amongst all the synclets (both input and output) involved in the message processing. msg is the object tree representing the SyncML message. The object tree is composed of instances of classes in the *com.funambol.framework.core* packages and represents a hierarchical view of the message.

For example, the synchronization message below will be transformed in the object hierarchy of Figure 15.

```

<SyncML>
<SyncHdr>
<VerDTD>1.1</VerDTD>
<VerProto>SyncML/1.1</VerProto>
<SessionID>12345678</SessionID>
<MsgID>2</MsgID>
<Target><LocURI>http://localhost</LocURI></Target>
<Source><LocURI>syncml-phone</LocURI></Source>
<Cred>
  <Meta><Type>syncml:auth-basic</Type></Meta>
  <Data>Z3Vlc3Q6Z3Vlc3Q=</Data>
</Cred>
</SyncHdr>
<SyncBody>
<Alert>
<CmdID>1</CmdID>
<Data>200</Data>
<Item>
<Target><LocURI>test</LocURI></Target>
<Source><LocURI>test</LocURI></Source>
<Meta>
<Anchor>
<Last>234</Last>
<Next>276</Next>
</Anchor>
</Meta>
</Item>
</Alert>
<Final/>
</SyncBody>
</SyncML>

```




Figure 15 - com.funambol.framework.core object tree example

An example of an input synclet is the following.

```

package com.foo.synclet;

import com.funambol.framework.logging.*;
import com.funambol.framework.core.*;
import com.funambol.framework.engine.pipeline.*;

```

```

import com.funambol.framework.tools.SyncMLUtil;

public class LoggingSynclet
implements InputMessageProcessor {
    // ----- Private data -----
    private static final FunambolLogger log =
FunambolLoggerFactory.getLogger("engine");

    // ----- Public methods -----

    /**
     * Logs the input message and context
     *
     * @param processingContext the message processing context
     * @param message the message to be processed
     *
     * @throws Sync4jException
     */
    public void preprocessMessage(MessageProcessingContext processingContext,
                                 SyncML message)
    throws Sync4jException {
        if (log.isLoggable(Level.INFO)) {
            log.info("-----");
            log.info("Input message processing context");
            log.info("");
            log.info(processingContext.toString());
            log.info("-----");

            log.info("Input message");
            log.info("");
            log.info(Util.toXML(message));
            log.info("-----");

            //
            // Sets the device ID to foo
            //
            message.getSyncHdr().getSource().setLocURI("foo");
        }
    }
}

```

This example synclet just logs to the *funambol.engine* logger the input message; also, it modifies the message setting the device ID to "foo".

One very important thing to keep in mind when developing a synclet (input or output) it is that in the current DS Server implementation just one instance is used to serve all the requests and so no instance/class variables must be used to store its status. Indeed the *MessageProcessingContext* must be used.

10.2.2. Configuring an input synclet

The input synclet so created is configured and deployed in the DS server adding a server side JavaBeans like the following in the *ds-server/config/com/funambol/server/engine/pipeline/input* directory.

```

<?xml version="1.0" encoding="UTF-8"?>
<java version="1.5.0" class="java.beans.XMLDecoder">
  <object class="com.foo.synclet.LoggingSynclet" />
</object>
</java>

```

Note: any XML file in the *ds-server/config/com/funambol/server/engine/pipeline/input* directory is used to create an input synclet, so this directory should be under strict control by the system administrator. The synclet loading and execution is done in alphabetic order, so if more than one synclet is used, a naming convention is an important part of processing control.

10.3. Postprocessing an outgoing message

To postprocess an outgoing message we have to create an output processor component (output synclet) and to configure it accordingly. This is described below.

10.3.1. Creating an output synclet

An output synclet is a class that implements the *com.funambol.framework.engine.pipeline.OutputMessageProcessor* interface. This interface defines just one method: *postProcessMessage(MessageProcessingContext context, SyncML msg)*.

The concepts behind the output message processing are the same as per input processing.

An example of an output synclet is the class shown below. This synclet injects a Get command into the outgoing message to request client capabilities:

```
package com.foo.synclet;

public class AddGetSynclet
implements OutputMessageProcessor {
    // ----- Constants

    // ----- OutputMessageProcessor

    public void postProcessMessage(MessageProcessingContext processingContext,
                                   SyncML message)
    throws Sync4jException {
        AbstractCommand[] commands = message.getBody().getCommands();

        AbstractCommand[] newCommands = new AbstractCommand[commands.length+1];

        Meta meta = new Meta();
        meta.setType("application/vnd.syncml-devinf+xml");

        Item item = new Item(
            new Target("/devinf11"),
            null,
            null,
            null,
            false
        );

        Get get = new Get(
            new CmdID(newCommands.length),
            false,
            null,
            null,
            meta,
            new Item[] { item }
        );

        System.arraycopy(commands, 0, newCommands, 0, commands.length);
        newCommands[commands.length] = get;
    }
}
```

10.3.2. Configuring an output synclet

As for the input synclet, the output synclet so created is configured and deployed in the DS server adding a server side JavaBeans like the following in the *ds-server/config/com/funambol/server/engine/pipeline/output* directory.

```
<?xml version="1.0" encoding="UTF-8"?>
<java version="1.5.0" class="java.beans.XMLDecoder">
  <object class="com.foo.synclet.AddGetSynclet" />
</object>
</java>
```

Note: any XML file in the *ds-server/config/com/funambol/server/engine/pipeline/output* directory is used to create an input synclet, so this directory should be under strict control by the system administrator. The synclet loading and execution is done in alphabetic order, so if more than one synclet is used, a naming convention is an important part of processing control.

10.3.3. The MessageProcessingContext

Funambol uses the *MessageProcessingContext* mentioned above to store the following properties:

Property name	Type	Scope	Description
funambol.session.id	java.lang.String	Session	Session ID of the current session
funambol.session.messageType	java.lang.String	Session	'WBXML' or 'XML' accordingly to the type of the messages in the current session
funambol.request.parameters	java.util.Map	Request	Parameters passed on the query string of the URL used at the transport level
funambol.request.headers	java.util.Map	Request	Headers passed in the request at the transport level.

10.3.4. How to stop message processing

If needed, an input or output pipeline component can ask the manager to stop further processing of the message. The `FunambolSynclet` does this by throwing a `com.funambol.framework.pipeline.StopProcessingException` exception.

11. SyncSource API

The SyncSource identifies the remote database that will be synchronized with the client database; it wraps a set of items to be synchronized. Any type of data (files, database tables, calendar events and so on) can be synchronized, but there must be a proper SyncSource for each type, capable of extracting and storing the data for a real data store.

The goal of Funambol is to provide a collection of SyncSources for the most common uses (for instance, files), although SyncSources can be independently developed and plugged in the synchronization engine, allowing synchronization requests targeted to virtually any database or type of data.

11.1. SyncSource class

The SyncSource class provides an implementation of the main methods; then, the derived classes *MergeableSyncSource* and the *FilterableSyncSource* implement additional specific methods.

Figure 16 shows the class diagram of a *DummySyncSource* that implements *MergeableSyncSource*.




Figure 16: *DummySyncSource* implements *MergeableSyncSource*.

Figure 17 shows the class diagram of a *DummySyncSource* that implements *FilterableSyncSource*.




Figure 17: *DummySyncSource* implements *FilterableSyncSource*

11.1.1. Methods list

Note: All of these methods throw an exception: *SyncSourceException*.

<i>Method</i>	<i>Return</i>	<i>Description</i>
beginSync(SyncContext context)	void	Init method of the synchronization process
endSync()	void	final method of the synchronization process
getAllSyncItemKeys()	SyncItemKey[]	Get all the IDs of the items in the back end
getSyncItemFromId(SyncItemKey syncItemKey)	SyncItem	Get the item associated to the given key from the Back end system
getNewSyncItemKeys(Timestamp since, Timestamp to)	SyncItemKey[]	Get the IDs of the new items from the given since timestamp to the given to timestamp
getUpdatedSyncItemKeys(Timestamp since, Timestamp to)	SyncItemKey[]	Get the IDs of the updated items from the given since timestamp to the given to timestamp
getDeletedSyncItemKeys(Timestamp since, Timestamp to)	SyncItemKey[]	Get the IDs of the deleted items from the given since timestamp to the given to timestamp
getSyncItemKeysFromTwin(SyncItem syncItem)	SyncItemKey[]	Get the IDs of the items that match with the given item
addSyncItem(SyncItem syncItem)	SyncItem	Called by the synchronization engine to add the given item to the back end. syncItem contains the item to add. The SyncSource developer can use SyncSource methods to extract the real item content and perform the appropriate actions on the back end.
updateSyncItem(SyncItem syncItem)	SyncItem	Called by the synchronization engine to update the given item to the back end. syncItem contains the item to add. The SyncSource developer can use SyncSource methods to extract the real item content and perform the appropriate actions on the back end.
removeSyncItem(SyncItemKey syncItemKey, Timestamp ts, boolean softDelete)	void	Called by the synchronization engine to delete the given item from the back end. syncItem contains the key of the item to delete.
setOperationStatus(String operation, int statusCode, SyncItemKey[] keys)	void	This call notifies the status code of the given operation performed on the items with the given keys.

SyncItem is a class in the package: import *com.funambol.framework.engine.SyncItem*. The methods available in this interface are reported in the following table:

<i>Method</i>	<i>Description</i>
getKey	Returns the item key.
getParentKey	Returns the item's parent key (hierarchic sync)
getState	Returns the item state.
setState	Sets the item state.
getContent	Returns the item content
setContent	Sets the item content
getFormat	Returns the item format
setFormat	Sets the item format
getType	Returns the item type
setType	Sets the item type
getTimestamp	Returns the timestamp of the last change of the item state and it is used in the synchronization process, in order to determine the operation to be performed on the sources.
setTimestamp	Sets the item timestamp
getSyncSource	Returns the SyncSource the item belongs to.

The *Key* and *ParentKey* properties are *com.funambol.framework.engine.SyncItemKey* objects identifying the item or the parent item. The parent item is used in the case of hierarchical synchronization where a parent/child structure is synchronized (e.g. email synchronization). In many cases a *SyncItemKey* just encapsulates a string.

The *State* (*getState*, *setState*) property represents the state of the item and can be one of the following values:

<i>Value</i>	<i>Description</i>
SyncItemState.NEW	The item is new
SyncItemState.UPDATED	The item has been updated
SyncItemState.DELETED	The item has been deleted
SyncItemState.SYNCHRONIZED	The items has been already synchronized
SyncItemState.UNKNOWN	The item is in a unknown state
SyncItemState.NOT_EXISTING	The item is not existing yet
SyncItemState.CONFLICT	The item is in a conflict state
SyncItemState.PARTIAL	The item contains only a portion of the content

Note that only NEW, UPDATED and DELETED items are available for use by a SyncSource. The other states are only used and processed by the synchronization engine. The associated Timestamp property is the timestamp of the last change.

Content is the item content returned encapsulated into a `object`. The MIME type of the content is represented by the property *Type* (*getType*, *setType*).

As a developer you do not usually need to develop your own implementation of *SyncItem*; Funambol provide a default implementation that should be sufficient in most cases. This is *com.funambol.framework.engine.SyncItemImpl*. See the source code of one of the Funambol connector for examples of how to use this class.

11.2. Mergeable SyncSource methods

The *MergeableSyncSource* class is an extension of a SyncSource that allows to merge the content of two items when a conflict is detected in order to avoid a loss of whatever information. For instance, while synchronizing the contacts, the server could find out that the same contact was updated on the server and on the client. An example could be that a new email address has been added in the client, while a telephone number was added in the server image of the record. In this case, the server can merge the server contact information and the client contact information, keeping both the new email address and phone number.

11.2.1. Methods list

Note: all the methods throw a exception: *SyncSourceException*

<i>Method</i>	<i>Return</i>	<i>Description</i>
mergeSyncItems(SyncItemKey syncItemKey, SyncItem syncItem)	boolean	Called by the synchronization engine to merge the server item identified by syncItemKey with the content obtained from the client and stored in syncItem. This methods must return true if the content has been changed so that the item resulting from the merge will be sent back to the client.

11.3. Filterable SyncSource methods

The *FilterableSyncSource* class is an extension of a *SyncSource* that allows the handling of filters during the sync process. Two type of filters can be used by a device:

- *Record filter*: permits to specify the records to sync (for example, only today's emails)
- *Field filter*: permits to specify criteria on the fields to synchronize (for example, skip the contact's photo)

11.3.1. Methods list

Note: all the methods throw a exception: *SyncSourceException*

<i>Method</i>	<i>Return</i>	<i>Description</i>
isSyncItemInFilterClause(SyncItem syncItem)	boolean	Detect if the given item matches the filter clause. Called to know if an item satisfies the filter
isSyncItemInFilterClause(SyncItemKey syncItemKey)	boolean	Detect if the item linked to the ID matches the filter clause. Called to know if the item specified with the given key satisfies the filter
getSyncItemStateFromId(SyncItemKey syncItemKey)	char	Get the "status" of the given item. Called to retrieve the status of an item. This method is required because the server isn't able to know the status of an item not inside the filter criteria.

12. Officer API

The Funambol security architecture is designed to be pluggable and is based on a very simple concept: authentication and authorization are centralized in a single dedicate component called officer.

An officer is a Java class that implements a specific interface. Concrete implementations provide adapters for external security services. For instance, a common external security service could be a database storing user profile information; a DBOfficer can be plugged into Funambol in order to perform authentication and authorization through the user information stored into the database.

12.1. Officer class




Figure 18: DummyOfficer diagram class

Figure 18 shows the class diagram of a *DummyOfficer*.

12.1.1. Methods list

Method	Return	Description
authenticateUser(Cred credential)	Sync4jUser	This method verifies the user on EIS
authorize(Principal principal, String resource)	Officer.AuthStatus	If required this method authorizes the user for a specific purpose.

Cred is a class in the package: *com.funambol.framework.core*

Principal is a class in the package: *java.security.Principal*

The possible statuses of authorization are:

- RESOURCE_NOT_AVAILABLE
- INVALID_RESOURCE
- NOT_AUTHORIZED
- AUTHORIZED
- PAYMENT_REQUIRED

13. Web Services API

In order to solve the problem of application-to-application communication, Funambol provides a web services layer that allows for applications to be integrated with the Funambol Platform and to access the resources provided by the system.

The application integration becomes much more flexible because web services provide a form of communication that is not tied to any particular platform or programming language. Thanks to the web service, Funambol makes resources available over the networks in a standardized fashion.

13.1. Introduction

The Funambol DS Server exposes the web services listed in the next paragraphs. The endpoint of the web service is: <http://localhost:8080/funambol/services/admin>.

In order to access the web service, the client must be authenticated. The default credentials are:

user = **admin**

password = **sa**

13.1.1. Funambol DS Server Web Services

The following table lists the web services exposed by the Funambol DS Server:

<i>Method</i>	<i>Parameters</i>	<i>Return</i>	<i>Description</i>
getServerVersion	String version	String	Returns the server version.
setServerConfiguration	ServerConfiguration config	void	Saves and apply the new server configuration
getServerConfiguration		ServerConfiguration	Returns the server configuration
getServerBean	String bean	String	Returns the server bean with the given name. If the bean does not exist, an AdminException is thrown
setServerBean	String bean, String obj	void	Sets the server bean with the given name.
login	String username	void	This method is only used to check credentials. Being this WS stateless, its methods are always called on request and each of them has to pass authentication. However, for instance the SyncAdmin, has a login panel that is displayed before any access to one of the other WS methods. In order to provide to the user an immediate feedback, the SyncAdmin (or any other client can just call login())
authorizeCredential	Credential authCred, String resource	AuthorizationResponse	Checks if user is authorized to use the given resource.
getUsers	String clause	Sync4jUser[]	Read all users that satisfy the parameter of search.
countUsers	String clause	int	Count the number of users that satisfy the specif clauses.
addUser	Sync4jUser user	void	Adds a new user and the assigned role to it
setUser	Sync4jUser user	void	Update the information of the specific user
deleteUser	String username	void	Delete the user
getDevices	String clause	Sync4jDevice[]	Read a list of device that satisfy the specific clause.

countDevices	String clause	int	Count the number of device that satisfy the specific clauses.
addDevice	Sync4jDevice d	String	Insert a new device and return the device ID
setDevice	Sync4jDevice d	void	Update the information of specific device.
deleteDevice	String deviceID	void	Delete the device with the given ID
getDevice	String deviceID	Sync4jDevice	Retrieves the device with the given ID.
getDeviceCapabilities	String deviceID	String	Read a device's capabilities for the specific device.
setDeviceCapabilities	String caps, String deviceID	Long	Set a device's capabilities for a specific device and return the identifier of the inserted capabilities
getPrincipals	String clause	Sync4jPrincipal[]	Read all principal that satisfy the clauses
countPrincipals	String clause	int	Count the number of principal that satisfy the specific clauses.
addPrincipal	Sync4jPrincipal p	long	Insert a new principal and return the principal ID
deletePrincipal	long principalID	void	Delete the specific principal
getRoles		String[]	Read the list of available roles
getLoggers		LoggerConfiguration[]	Returns the available LoggerConfiguration
getAppenders		String	Returns an XML serialization of a map with all available Appender objects on the server.
setAppender	String appenderBean	void	Sets the given appender
getAppenderManagementPanel	String appenderClassName	String	Returns the class name of the management panel to configure an appender with the given class name.
setLoggerConfiguration	String loggerBean	void	Saves and apply the new logger configuration
setLoggingConfiguration	LoggingConfiguration config	void	Saves and apply the new logging configuration
getLatestDSServerUpdate		Component	Returns the latest DS Server available update
getLastTimestamps	String clause	LastTimestamp[]	Read all last timestamp that satisfy the clauses.
deleteLastTimestamp	long principalId, String sourceId	void	Delete the specific last timestamp.
countLastTimestamps	String clause	int	Count the number of last timestamps that satisfy the specific clauses.
getModulesName		Sync4jModule[]	Read all modules information
getModule	String moduleID	String	Read the information of the specific module
addSource	String moduleId, String connectorId, String sourceTypeId, String source	void	Adds a new source into datastore and create a relative XML file with configuration. The source must have a defined source type. The source type must refer to a connector. The connector must refer to a module.
getSync4jSources	String clause	Sync4jSource[]	Read a list of syncSource that satisfy the specific conditions.
getSyncSourceClasses	String[] sourceTypesID	String[]	Returns an array with the classes relative to the given source types
deleteSource	String sourceUri	void	Delete a specific source and the relative file of configuration.
setSource	String moduleId, String connectorId, String sourceTypeId, String source	void	Update a specific source into datastore and create a relative XML file with configuration. The source must have a defined source type. The source type

			must refer to a connector. The connector must refer to a module.
sendNotificationMessage	String deviceId, String alerts, Integer uimode	void	Sends a notification message to the given device.
sendNotificationMessages	String username, String alerts, Integer uimode	void	Sends a notification message to all devices of the principals with the given username

The installation of the Email Module adds the web services listed in the following table. The endpoint of the email module web service is: <http://localhost:8080/funambol/services/email>

<i>Method</i>	<i>Parameters</i>	<i>Return</i>	<i>Description</i>
getUsers	String clause	MailServerAccount[]	Get the list of users that matches the filter clause
getUser	String username	MailServerAccount	Get the user for the given username
getUserFromID	long accountID	MailServerAccount	Get the user for the given account ID
insertUser	MailServerAccount msa	int	insert user and the mail server configuration
updateUser	MailServerAccount msa	int	update user
disableUser	long accountID	int	Disable the account related to the given ID
enableUser	long accountID	int	Enable the account related to the given ID
markUserAsDelete	long accountID	int	Delete the user from the fnbl_email_account and set as 'D' in the fnbl_email_push_registry
checkAccount	MailServerAccount msa, Integer timeout	MailServerError	check the account on the mail server
insertPubMailServer	MailServer ms	int	insert public mail server configuration
deletePubMailServer	String mailServerID	int	delete mail server related to the mail server ID
updatePubMailServer	MailServer ms	int	update public mail server
getPubMailServers	String clause	MailServer[]	Get the list of the public mail servers
getPubMailServerFromID	String mailServerID	MailServer	get public mail server
getCachedInfo	String username, String protocol	SyncItemInfoAdmin[]	get the info from the local inbox folder (fnbl_email_inbox) for the specified username
getImapFolders	MailServerAccount msa	Object[]	Returns all the folders names for the given IMAP account.

14. Funambol Software Development Kit

The Funambol Software Development Kit (SDK) [7] is the group of tools available to develop Funambol extensions.

To install it, extract the archive *funambol-sdk-<version>.tar.gz* in a directory of choice; the directory structure shown in Figure 19 will be created.

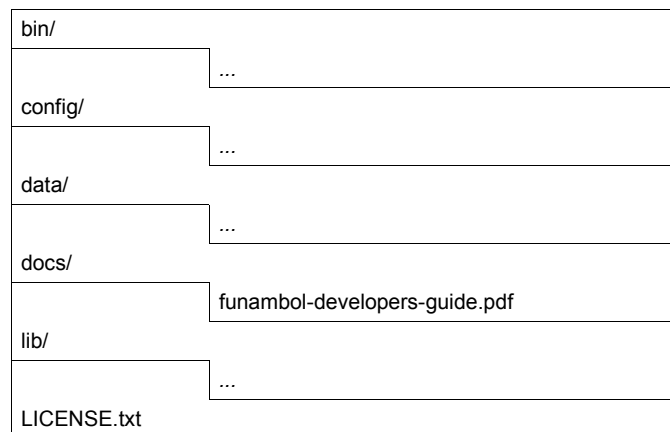


Figure 19: Funambol SDK directory structure

There are two ways to develop Funambol extensions: using your own custom development environment or using Maven [4]. Both methods will be explained in the following sections, but the preferred method (since Funambol v6.5) is using Maven.

In the following sections, it is assumed that the developer is familiar with the following concepts:

- Java development
- Funambol architecture
- Maven (optional)

14.1. Obtaining and building the source code

Funambol is open source! The best source of information and documentation is the source code, which is available on the Funambol public CVS [6].

The instructions on how to download and build Funambol v6.5 can be found here:

<https://wiki.objectweb.org/sync4j/Wiki.jsp?page=BuildingFunambolV6.5>

The instructions on how to download and build Funambol v7 can be found here:

<https://wiki.objectweb.org/sync4j/Wiki.jsp?page=HOWTOBuildCapri>

14.2. Developing with a custom environment

This section describes how to develop a Funambol extension using just an editor and the Java Development Kit (JDK).

When developing in a custom environment there is no any additional requirement other than generating a s4j package as described earlier; how to do it is left to the developer. One important thing to note is that in order to compile your classes you have to make sure to have the following framework jars in the build classpath:

Name	Download url
Funambol v6.5	
core-framework-<version>.jar	http://m2.funambol.com/repositories/artifacts/funambol/core-framework/6.5.4/core-framework-6.5.4.jar
server-framework-<version>.jar	http://m2.funambol.com/repositories/artifacts/funambol/server-framework/6.5.7/server-framework-6.5.8.jar
ds-server-<version>.jar	http://m2.funambol.com/repositories/artifacts/funambol/ds-server/6.5.14/ds-server-6.5.14.jar
admin-framework-<version>.jar	http://m2.funambol.com/repositories/artifacts/funambol/admin-framework/6.5.2/admin-framework-6.5.2.jar
Funambol v7	
core-framework-<version>.jar	http://m2.funambol.com/repositories/artifacts/funambol/core-framework/7.0.0/core-framework-7.0.0.jar
server-framework-<version>.jar	http://m2.funambol.com/repositories/artifacts/funambol/server-framework/7.0.3/server-framework-7.0.3.jar
ds-server-<version>.jar	http://m2.funambol.com/repositories/artifacts/funambol/ds-server/7.0.2/ds-server-7.0.2.jar
admin-framework-<version>.jar	http://m2.funambol.com/repositories/artifacts/funambol/admin-framework/6.5.2/admin-framework-6.5.2.jar

These libraries can be found under the directory lib of the Funambol SDK.

For an example on how to develop a connector, see chapter 2, “Getting started on connector development”.

14.3. Developing with Maven

Maven is a software project management and comprehension tool. Based on the concept of a project object model (POM), Maven can manage a project's build, reporting and documentation from a central piece of information. For more information about Maven, see [4]. In the next section it is assumed that the reader is already familiar with Maven and its concepts.

Even if many Funambol v6.5 components have been put under the Maven build system already, only with Funambol v7 Maven has become the default build tool for Funambol components. All Funambol artifacts are now represented, built and deployed as Maven artifacts. Funambol has also set up a public maven repository where all components are published. Funambol maintains two repositories: *artifacts*, for released artifacts and *snapshots*, for snapshots; the latter keeps the snapshots for the last 10 days. The repository can be accessed through the URL <http://m2.funambol.org/repositories>

14.3.1. Maven configuration

Before using Maven, the repository above must be added to your maven environment. Do to so, edit your *settings.xml* file (either under maven or your user home) and add the following sections:

```

<repositories>
  <repository>
    <id>artifacts</id>
    <url>http://m2.funambol.org/repositories/artifacts</url>
  </repository>
  <repository>
    <id>snapshots</id>
    <url>http://m2.funambol.org/repositories/snapshots</url>
  </repository>
</repositories>
<pluginRepositories>
  <pluginRepository>
    <id>artifacts</id>
    <url>http://m2.funambol.org/repositories/artifacts</url>
  </pluginRepository>
  <pluginRepository>
    <id>snapshots</id>
    <url>http://m2.funambol.org/repositories/snapshots</url>
  </pluginRepository>
</pluginRepositories>

```

14.3.2. Creating a new module

The best and quickest way to create a Funambol module with Maven is using the funambol-module-archetype:

```
mvn archetype:create -DarchetypeGroupId=funambol -DarchetypeArtifactId=funambol-  
module-archetype -DarchetypeVersion=6.5.0 -DgroupId=<your group> -DartifactId=<your  
artifact>
```

For example:

```
mvn archetype:create -DarchetypeGroupId=funambol -DarchetypeArtifactId=funambol-  
module-archetype -DarchetypeVersion=6.5.0 -DgroupId=yourgroup  
-DartifactId=yourconnector
```

This creates a new Maven project for *yourgroup:yourconnector* in a directory called *myconnector*. This projects contain a skeleton of a SyncSource and of a input and output Synclet.

Note: the command line above creates a Funambol v6.5 module specifying all the required dependencies. You don't need to worry about which jars you have to download and add to the classpath, as Maven will do all this for you. To create a Funambol v7 module, just use the version 7.0.0 of funambol-maven-archetype:

```
mvn archetype:create -DarchetypeGroupId=funambol -DarchetypeArtifactId=funambol-  
module-archetype -DarchetypeVersion=7.0.0 -DgroupId=yourgroup  
-DartifactId=yourconnector
```

14.3.3. Building the module

To build the module created as described below, just go inside the newly created directory and run maven. For example:

```
mvn package
```

14.4. The Funambol Connector Testing Framework

The Funambol Connector Testing Framework (FCTF) is a command-line tool that can be used to test and validate any Funambol connector against any set of PIM or e-mail items, without setting up a complete test environment comprising a device, a Data Synchronization server and a back-end system.

Note: you need a running back-end system and the connector for an FCTF instance to work.

Each FCTF instance is defined by its own set of PIM and/or e-mail items and corresponding expected results, in addition to the connector to be used.

Three operation modes are available:

1. Automatic mode (default): only pre-defined sets of tests depending on the MIME type can be performed.
2. Single-test mode: only one test is specified using the *--command* keyword.
3. Batch mode: a batch file containing a series of tests is specified using the *--batch* keyword.

Automatic tests allow for different levels of certification of a connector. A connector that passes the basic automatic tests, for example, can be declared to be FCTF-certified at the basic level. The certification level can be specified on the command line after the *--testset* keyword, the default value being basic.

The batch files (and, therefore, also the items used for the individual tests) may or may not be connector-specific, since the expected results can change depending on the features of the back-end system.

14.4.1. Usage

The Funambol Connector Testing Framework has been designed to be used from the command line; the executables can be found under the *bin* directory of the work environment.

Once the work environment has been generated, the user is expected to change their working directory to *<installation directory>/Funambol/tools/sdk* and run the command *bin/fctf* from there.

An example of a typical command line that will launch the automatic tests at the basic certification level for a vCard sync source belonging to the Foundation standard connector is:

```
bin/fctf -s config/foundation-vcard.xml
```

bin/fctf requires at least one argument (the SyncSource); the following tables details all available options:

Argument	Short form	Parameter(s)	Default value	Usage notes
--source	-s	The path to the XML file defining the SyncSource; it usually begins with <i>config</i> .	N/A	Mandatory. It is needed to load the SyncSource that will be operated by the tool.
--testset	-t	The certification level (it must correspond to an existing subdirectory of data).	basic	Default option. If this option is selected, the FCTF will operate in the automatic-tests mode.
--command	-c	A single command, according to the grammar defined in section 14.4.1, Tests.	N/A	If this option is selected, the FCTF will operate in single-test mode.
--batch	-b	The path to the batch command file containing the tests to be performed (one per line).	N/A	If this option is selected, the FCTF will operate in batch mode.
--datasource	-D	A list of datasource names (usually, 2: the user DB and the core DB). The first part of each name must correspond to an existing subdirectory of <i>config/com/funambol/server/db</i> , the second part to an existing <i>.xml</i> file within that subdirectory.	jdbc/fnblcore jdbc/fnbluser	
--authtype	-a	The type of authentication (<i>basic</i> , <i>md5</i> or <i>MAC</i>).	basic	This is used for connector whose officer uses a special authentication type.
--path	-P	The path to the data item.	data	This is used to load different sets of items without modifying the built-in work environment.
--user	-u	The username.	fctf	
--password	-p	The user's password.	the username	This is useful only when the user needs be authenticated.
--device	-d	The device ID	fctf	
--officer	-o	The path to the XML file defining the officer; it usually starts with <i>config</i> .	config/test-officer.xml	This depends on the connector to be tested.
--verbose	-v	No parameter.	N/A	If selected, in case of error the stack trace of the exception raised will be displayed.

Tests

The FCTF performs a pre-defined series of tests that are a combination of the following basic tests:

- removal of an item from the back-end system (*d*);
- update of an item on the back-end system (*u*);
- retrieval of an item from the back-end system (*g*);
- retrieval of an item's twin from the back-end system (*t*).

All of these tests are preceded by the addition of an item on the back-end system, and followed, if necessary, by the removal of the items. Therefore every test is independent and leaves the back-end server in a clean state, if no concurrent operations are performed in the meanwhile on the same item.

In the list above, the letter in brackets is the one used to call for an individual test. It will be followed by the item(s) file name(s). The letter *x* is also used to indicate the expected result, if applicable. The same grammar is used both in the command-line single-test mode and the batch mode:

- `<SyncSource> d <filename>` - the item is added in the back-end system and then removed;
- `<SyncSource> u <filename1> <filename2> x <filename3>` - the first item is added in the back-end system, it's updated with the modified content of *filename2* and then compared to the expected result in *filename3*;
- `<SyncSource> g <filename1> x <filename2>` - the first item is added in the back-end system, then retrieved and compared to the expected result in *filename2*;
- `<SyncSource> t <filename1> <filename2> ... <filenameN> x <K>` - all items but the first one are added in the back-end system, then the twins of the first one are looked for; *K* (any number from 0 to *N*) twins are expected to be found.

If one such command is called in the single-test mode, it must be specified after the `--test` keyword. If a series of commands is launched in the batch mode, they will be listed, one per line, in the batch file.

In the automatic mode, the tests do not need to be specified using this grammar because the test set is pre-defined. For example, if the items to be tested are *item1.txt*, *item2.txt* and *item3.txt*, the following tests will be automatically generated and launched:

- `g item1.txt x item1.txt`
- `g item2.txt x item2.txt`
- `g item3.txt x item3.txt`
- `u item1.txt item2.txt x item2.txt`
- `u item2.txt item3.txt x item3.txt`

This will be done for every MIME type supported by the SyncSource being tested, starting from the preferred MIME type.

Test items

The items are provided as vCalendar, iCalendar, vCard, SIF-E, SIF-T, SIF-C and RFC882 data. They are in the *data* subdirectory. The MIME type of each item set is inferred from the file extension:

- `.vcs` for vCalendar (1.0)
- `.ics` for iCalendar (vCalendar 2.0)
- `.vcf` for vCard
- `.sife`, `.sift`, `.sifc` for the SIF format
- `.eml` for RFC882 data

It's important to notice that these data are not filtered by the synclets, therefore this tester cannot be used to test the behavior of different devices.

Items used in the automatic tests are grouped in different subdirectories of directory *data* according to the certification level they belong to. Each MIME type within a certification level has its own directory where items of that type will be collected. The directory subtree will reproduce the structure of MIME taxonomy. For example, vCalendar (1.0) items belonging to the Basic certification level will be found under `data/basic/text/x-vcalendar`. The file names are relevant because the items are sorted in alphabetical order when the automatic tests are generated and launched.

Test pass conditions

During automatic tests, server-generated data (in vCalendar, iCalendar, vCard, SIF-E, SIF-T, SIF-C or RFC882 format) are extracted by the back-end response message and compared with the corresponding item in the pre-defined set of the expected results. If the comparison does not outline any data loss, the tests have succeeded.

It is important to underline that the FCTF only checks for the possible loss of data between the expected result and the actual result. If the actual result contains *more* information than the expected one, that is considered as a successful comparison.

The comparison is done on a line-wise basis for vCard, vCalendar and iCalendar items, while for SIF items it is based on (non-empty) end nodes (leaves) of the XML node tree.

In this way, the SyncML message is not compared as such to an expected result, as in the more generic Funambol Test Suite. The FCTF focuses only on the PIM and e-mail data.

The same mechanism is used in the single-test and batch modes for retrieval and update tests. In the removal and twin-search cases, the expected results are just the actual deletion of the item and a positive result of the twin search. If these outcomes occur, the tests have succeeded.

14.4.2. Certifying a connector

The certification of a connector may require a few steps.

Installing the libraries

Note: This step is always necessary.

The libraries needed for the connector's SyncSources to run (for example, the connector's module) must be copied as JARs in *lib/ext*. This includes the libraries for the access to data sources, like the DB.

Installing the SyncSources

Note: This step is always necessary.

The *.xml* files containing the marshalled version of SyncSources must be copied in the *config* subdirectory. These files are usually exactly the same *.xml* files that can be found under the *config* directory in the connector's source, already prepared for usage by the Funambol Administration Tool.

Installing the officer

Note: This step is often, but not always, necessary.

Some connectors require an officer for user authentication. In this case, its marshalled version must also be copied in the *config* subdirectory. If no user authentication is needed on the back-end system, the default dummy officer (*TestOfficer*) can be used.

Setting up the data source

Note: This step is often not necessary.

In the *config/com/funambol/server/db* directory and its subdirectory *jdbc* there are *.xml* files containing the marshalled version of the automatic configurer of the JDBC connections to the datasources. New files can be added or the default ones can be modified to replicate the features of the data source in use.

Setting up the data for individual (or batch) tests

Note: This step is not strictly necessary for certification, only for other tests.

Item files can be added in the *data* subdirectory. It's important to follow the file extension rules as of Test items, section 14.4.1.

Setting up the data for a custom certification

Note: this step is not necessary for pre-defined certification levels, only for custom certifications.

Custom certifications can be added under the *data* directory. In order to do that, a subdirectory must be created with the name of the custom certification level. Under that subdirectory, it is necessary to create a directory tree that follows the MIME types tree, including all MIME types that are wished to be included in the custom certification. The *data/basic* subdirectory can be used as an example. Then,

item files must be created in the correct directories. The file names are relevant because the automatic tests will be generated and launched on the basis of the alphabetical order.

Launching the tests

See section 14.4.1, Usage.

14.4.3. Limitations

A key purpose of the FCTF is to test the mapping between the foundation data model and the back-end data model. This cannot be tested directly if the tool has to be an automated one because it cannot perform any direct test on the status of the back-end system. The back-end status can only be inferred by the data retrieved by the connector, but those data, in turn, are inserted there by the insert method of the connector.

If a property *A* on a client item is *incorrectly* mapped to property *b* on the server and this is “correctly” mapped back to the client as *B*, the difference between *A* and *B* may be used to detect an error in the behavior of the connector. On the contrary, if a property *A* on a client item is *correctly* mapped to property *a* on the server and this is correctly mapped back to the client as *A*, the item will pass the test. But what if a property *A* on a client item is *incorrectly* mapped to property *b* on the server and this is again *incorrectly* mapped back to the client as *A*? There is no way to tell this case from the previous one.

<i>Original value</i>	<i>Value on the server after synchronization with the client</i>	<i>Value on the client after retrieval from the server</i>	<i>Test outcome</i>	<i>Bugs?</i>
A	b	B	Failed	Yes
A	a	A	Passed	No
A	b	A	Passed	Yes

As a consequence, there's a class of bugs that cannot be verified with this tool, like field swaps, some time zone errors etc. Manual tests with human interaction with the back-end system are still needed to investigate these cases.

14.4.4. Error codes

Please note that a connector cannot be considered certified as long as it has not passed the automatic tests for the required certification level with the application successfully exiting without displaying any fatal warning or error.

When the application exits with an error, a brief error message is usually displayed and an error code is provided. The most likely solution to the problem corresponding to each error code is explained in this table:

<i>Code</i>	<i>Mode</i>	<i>Most likely solution</i>
2	any	Check that the SyncSource specified with the <i>-s</i> argument is valid and can be unmarshalled. The required libraries must be under <i>lib/ext</i> .
3	any	Check that the officer specified with the <i>-o</i> argument is valid and can be unmarshalled. The required libraries must be under <i>lib/ext</i> .
4	any	Check that the datasource name specified with the <i>-D</i> argument is correct.
5	any	Check that the datasource is available, its parameters correct and the required libraries under <i>lib/ext</i> .
6	any	Check that the datasource is up and running properly.
7	any	Check that the <i>.xml</i> files used to configure and bind the datasources are correct and in the correct locations.
99	single test	Check the grammar of the test command. It must match one of the cases listed in section 14.4.1, Tests.
102	single test	Check the mapping used by the SyncSource and debug the connector. Be sure that the expected results were correct.
103	single test	Check the mapping and twin search criterion used by the SyncSource and debug the connector. Be sure that the expected twin count were correct.

Code	Mode	Most likely solution
104	single test	Check that the back-end system is up and running. Debug the addition (create) function of the SyncSource.
105	single test	Debug the deletion (remove) function of the SyncSource.
106	single test	Debug the retrieval (get) function of the SyncSource.
107	single test	Debug the update (modify) function of the SyncSource.
108	single test	Debug the twin search (get twins) function of the SyncSource.
110	single test	Debug the SyncSource.
111	single test	Check that all item files mentioned in the test command exist under the default item directory (<i>data</i>) or the directory specified with the <i>-P</i> argument.
112	single test	Check that the extensions of all item files mentioned in the test command are among the supported extensions listed in section 14.4.1, Test items.
113	single test	Check that the twin count command ends with an integer.
202	batch	Check the mapping used by the SyncSource and debug the connector. Be sure that the expected results were correct.
203	batch	Check the mapping and twin search criterion used by the SyncSource and debug the connector. Be sure that the expected twin count were correct.
204	batch	Check that the back-end system is up and running. Debug the addition (create) function of the SyncSource.
205	batch	Debug the deletion (remove) function of the SyncSource.
206	batch	Debug the retrieval (get) function of the SyncSource.
207	batch	Debug the update (modify) function of the SyncSource.
208	batch	Debug the twin search (get twins) function of the SyncSource.
210	batch	Debug the SyncSource.
211	batch	Check that all item files mentioned in the test command exist under the default item directory (<i>data</i>) or the directory specified with the <i>-P</i> argument.
212	batch	Check that the extensions of all item files mentioned in the test command are among the supported extensions listed in section 14.4.1, Test items.
213	batch	Check that the twin count command ends with an integer.
220	batch	Check that the batch file specified with the <i>-b</i> option exists.
302	automatic tests	Check the mapping used by the SyncSource during the addition of a new item and debug the connector.
304	automatic tests	Check that the back-end system is up and running. Debug the addition (create) function of the SyncSource.
305	automatic tests	Debug the deletion (remove) function of the SyncSource.
306	automatic tests	Debug the retrieval (get) function of the SyncSource.
310	automatic tests	Debug the SyncSource.
311	automatic tests	Do not modify the automatic test files while the application is running.
312	automatic tests	Check that the extensions of all item files in the automatic tests directories are among the supported extensions listed in section 14.4.1, Test items.
320	automatic tests	Check that the directory specified with the <i>-a</i> option exists <i>under the default item directory (data)</i> or the directory specified with the <i>-P</i> argument.
352	automatic tests	Check the mapping used by the SyncSource during the update of an item and debug the connector.
354	automatic tests	Debug the addition (create) function of the SyncSource.
355	automatic tests	Debug the deletion (remove) function of the SyncSource.
356	automatic tests	Debug the retrieval (get) function of the SyncSource.
357	automatic tests	Debug the update (modify) function of the SyncSource.
360	automatic tests	Debug the SyncSource.
361	automatic tests	Do not modify the automatic test files while the application is running.

15. Appendix A - Sync4j Interchange Formats

The Sync4j Interchange Format (SIF) is a way to represent PIM data coming from different clients in a common structure to make it easier information exchange.

SIF format is based on a XML representation of PIM data.

The goal we wanted to achieve with this representation (as opposed to other standards) is having a simple XML structure representing any possible PIM properties. The basic idea is that every PIM object (contact, calendar, and so on) has properties in the form *name-value*. Therefore, in a SIF document tags represent properties and tags content represent their values.

We have a SIF representation for each type of PIM data. They are SIF-C for contacts, SIF-E for events, SIF-T for tasks and SIF-N for notes. The Following sections explain each SIF format in detail.

15.1. SIF-C

A contact contains information about a person. Each client (Outlook, Pocket PC, Palm...) can store different kind of personal data. For example, Pocket PC devices use a subset of the information used by Outlook.

An example of a SIF contact is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<contact>
  <Companies>Maga S.p.A.</Companies>
  <CompanyMainTelephoneNumber/>
  <Email3Address>john@yahoo.com</Email3Address>
  <Business2TelephoneNumber>+001 2345776</Business2TelephoneNumber>
  <CarTelephoneNumber/>
  <Email2Address>john2@hotmail.com</Email2Address>
  <OtherAddressCountry/>
  <OtherFaxNumber/>
  <Suffix/>
  <BusinessAddressPostOfficeBox/>
  <FirstName>John</FirstName>
  <Subject/>
  <Hobby>Hockey</Hobby>
  <HomeAddressPostOfficeBox/>
  <OtherTelephoneNumber/>
  <PersonalWebPage>http://www.jhon.com</PersonalWebPage>
  <Department>dep 1</Department>
  <Home2TelephoneNumber>+001 3456 7767</Home2TelephoneNumber>
  <HomeAddressStreet>21th Street</HomeAddressStreet>
  <JobTitle>Programmer</JobTitle>
  <Anniversary>2004-11-23</Anniversary>
  <PrimaryTelephoneNumber>+001 45667 34543</PrimaryTelephoneNumber>
  <MobileTelephoneNumber>3357766689</MobileTelephoneNumber>
  <YomiCompanyName/>
  <BusinessAddressCountry>New York</BusinessAddressCountry>
  <Sensitivity>0</Sensitivity>
  <OtherAddressState/>
  <NickName>JJ</NickName>
  <HomeAddressPostalCode>54065</HomeAddressPostalCode>
  <OrganizationalIDNumber/>
  <ManagerName>Mr.White</ManagerName>
  <BusinessTelephoneNumber>+001 12399 9999</BusinessTelephoneNumber>
  <YomiLastName/>
  <WebPage>http://www.maga.com</WebPage>
  <BusinessAddressCity>New York</BusinessAddressCity>
  <Email2AddressType>SMTP</Email2AddressType>
  <Title>Ing</Title>
  <FileAs>Doe, John</FileAs>
  <MiddleName>Patrick</MiddleName>
  <HomeAddressCountry>home country</HomeAddressCountry>
  <Birthday>1973-10-03</Birthday>
  <RadioTelephoneNumber/>
  <OtherAddressPostalCode/>
  <BusinessAddressPostalCode>54667</BusinessAddressPostalCode>
  <BusinessAddressStreet>45th street</BusinessAddressStreet>
```

```

<AssistantTelephoneNumber/>
<PagerNumber/>
<HomeAddressCity>home city</HomeAddressCity>
<Profession>IT developer</Profession>
<HomeAddressState>home state</HomeAddressState>
<YomiFirstName/>
<OtherAddressStreet/>
<OtherAddressCity/>
<CallbackTelephoneNumber/>
<OtherAddressPostOfficeBox/>
<Initials>J.D.</Initials>
<Mileage/>
<Language/>
<EmailAddress>john.doe@funambol.com</EmailAddress>
<Children/>
<BusinessFaxNumber>+001 456 65 5556456</BusinessFaxNumber>
<Email3AddressType>SMTP</Email3AddressType>
<Importance>0</Importance>
<EmailAddressType>SMTP</EmailAddressType>
<Body>a little note...</Body>
<TelexNumber/>
<OfficeLocation/>
<AssistantName/>
<Spouse/>
<Categories/>
<HomeTelephoneNumber>+001 456 65 5454 </HomeTelephoneNumber>
<BusinessAddressState>USA</BusinessAddressState>
<ComputerNetworkName/>
<CompanyName>Sync4j</CompanyName>
<HomeFaxNumber/>
<LastName>Doe</LastName>
</contact>

```

The SIF-C fields defined by default in Sync4j are listed in the table below.

Property	Description
Anniversary	Returns or sets the anniversary. It is in format YYYY-MM-DD
AssistantName	Returns or sets the name of the person who is the assistant for the contact. Corresponds to the Assistant's name: box on the Details page of a ContactItem.
AssistantTelephoneNumber	Returns or sets the telephone number of the person who is the assistant for the contact
Birthday	Returns or sets the birthday. It is expressed in format YYYY-MM-DD
Body	Returns or sets the clear-text body of the item.
Business2TelephoneNumber	Returns or sets the second business telephone number for the contact.
BusinessAddressCity	Returns or sets the city name portion of the business address for the contact
BusinessAddressCountry	Returns or sets the country code portion of the business address for the contact
BusinessAddressPostalCode	Returns or sets the postal code (zip code) portion of the business address for the contact
BusinessAddressPostOfficeBox	Returns or sets the post office box number portion of the business address for the contact
BusinessAddressState	Returns or sets the state code portion of the business address for the contact
BusinessAddressStreet	Returns or sets the street address portion of the business address for the contact
BusinessFaxNumber	Returns or sets the business fax number for the contact
BusinessLabel	Returns or sets a String with complete business address
BusinessTelephoneNumber	Returns or sets the first business telephone number for the contact
BusinessWebPage	Business web page
CallbackTelephoneNumber	Returns or sets the callback telephone number for the contact
CarTelephoneNumber	Returns or sets the car telephone number for the contact
Categories	Returns or sets the categories assigned to the Outlook item.
Children	Returns or sets the names of the children of the contact
CompanyMainTelephoneNumber	Returns or sets the company main telephone number for the contact
CompanyName	Returns or sets the company name for the contact
Companies	Returns or sets the names of the companies associated with the item.
ComputerNetworkName	Returns or sets the name of the computer network for the contact

Property	Description
Department	Returns or sets the department name for the contact
Email1Address	Returns or sets a String representing the e-mail address of the first e-mail entry for the contact.
Email1AddressType	Returns or sets a String representing the address type (such as EX or SMTP) of the first e-mail entry for the contact. This is a free-form text field, but it must match the actual type of an existing mail transport.
Email2Address	Returns or sets the e-mail address of the second e-mail entry for the contact
Email2AddressType	Returns or sets a String representing the address type (such as EX or SMTP) of the second e-mail entry for the contact. This is a free-form text field, but it must match the actual type of an existing mail transport.
Email3Address	Returns or sets the e-mail address of the third e-mail entry
Email3AddressType	Returns or sets a String representing the address type (such as EX or SMTP) of the third e-mail entry for the contact. This is a free-form text field, but it must match the actual type of an existing mail transport.
FileAs	Returns or sets the default keyword string assigned to the contact when it is filed
FirstName	Returns or sets the first name for the contact.
Folder	Return or sets the Folder the contact has to be written in. The contact in the default contacts folder can be represented with "/" root or without <Folder> tag. Eg. <Folder>/subfolder1/subfolder11/</Folder>. Note the final "/".
Gender	Gender
Hobby	Returns or sets the hobby for the contact
Home2TelephoneNumber	Returns or sets the second home telephone number for the contact
HomeAddressCity	Returns or sets the city portion of the home address for the contact
HomeAddressCountry	Returns or sets the country portion of the home address for the contact
HomeAddressPostalCode	Returns or sets the postal code portion of the home address for the contact
HomeAddressPostOfficeBox	Returns or sets the post office box number portion of the home address
HomeAddressState	Returns or sets the state portion of the home address for the contact
HomeAddressStreet	Returns or sets the street portion of the home address for the contact
HomeFaxNumber	Returns or sets the home fax number for the contact
HomeLabel	Returns or sets a String with complete home address
HomeTelephoneNumber	Returns or sets the first home telephone number for the contact
HomeWebPage	Personal web page
id	Returns or sets the identifier of contact item
Importance	Returns or sets the relative importance level for the Outlook item. Can be one of the following OllImportance constants: ollImportanceHigh(2), ollImportanceLow(0), or ollImportanceNormal(1). This property corresponds to the MAPI property PR_IMPORTANCE.
Initials	Returns or sets the initials for the contact
JobTitle	Returns or sets the job title for the contact
Language	Returns or sets the language for the contact
LastName	Returns or sets the last name for the contact
ManagerName	Returns or sets the manager name for the contact
MiddleName	Returns or sets a String representing the middle name for the contact. This property is parsed from the FullName property, but may be changed or entered independently should it be parsed incorrectly. Note that any such changes or entries to this property will be overwritten by any subsequent changes of entries to FullName.
Mileage	Returns or sets a String representing the mileage for an item. This is a free-form string field and can be used to store mileage information associated with the item (for example, 100 miles documented for an appointment, contact, or task) for purposes of reimbursement.
MobileTelephoneNumber	Returns or sets a String representing the mobile telephone number
NickName	Returns or sets a String representing the nickname for the contact.
OfficeLocation	Returns or sets a String specifying the specific office location (for example, Building 1 Room 1 or Suite 123) for the contact. This property corresponds to the MAPI property PR_OFFICE_LOCATION.

Property	Description
OrganizationalIDNumber	Returns or sets the organizational ID number for the contact
OtherAddressCity	Returns or sets the city portion of the other address for the contact
OtherAddressCountry	Returns or sets the country portion of the other address for the contact
OtherAddressPostalCode	Returns or sets the postal code portion of the other address for the contact
OtherAddressPostOfficeBox	Returns or sets the post office box portion of the other address for the contact
OtherAddressState	Returns or sets the state portion of the other address for the contact
OtherAddressStreet	Returns or sets the street portion of the other address for the contact
OtherFaxNumber	Returns or sets the other fax number for the contact
OtherLabel	Returns or sets a String with complete other address
OtherTelephoneNumber	Returns or sets the other telephone number for the contact
PagerNumber	Returns or sets the pager number for the contact
PrimaryTelephoneNumber	Returns or sets the primary telephone number for the contact
Profession	Returns or sets the profession for the contact
Revision	Revision number
RadioTelephoneNumber	Returns or sets the radio telephone number for the contact
Sensitivity	Returns or sets the sensitivity for the Outlook item. Can be one of the following OISensitivity constants: olConfidential(3), olNormal(0), olPersonal(1), or olPrivate(2). This property corresponds to the MAPI property PR_SENSITIVITY
Spouse	Returns or sets the spouse name entry for the contact
Subject	Returns or sets the subject for the Outlook item. This property corresponds to the MAPI property PR_SUBJECT. The Subject property is the default property for Outlook items.
Suffix	Returns or sets the name suffix (such as Jr., III, or Ph.D.) for the contact
TelexNumber	Returns or sets the telex number for the contact
Timezone	Returns or set the timezone for the contact
Title	Returns or sets the title for the contact
Uid	Unique ID
WebPage	Returns or sets the URL of the Web page for the contact
YomiCompanyName	Returns or sets a String indicating the Japanese phonetic rendering (yomigana) of the company name for the contact
YomiFirstName	Returns or sets a String indicating the Japanese phonetic rendering (yomigana) of the first name for the contact
YomiLastName	Returns or sets a String indicating the Japanese phonetic rendering (yomigana) of the last name for the contact

You will note that some names and constants are a bit odd. This is for legacy reason since they are used by Microsoft Exchange Server or some Microsoft clients (Outlook, Pocket Outlook).

The following table lists all fields and their use on clients:

Property	Outlook	Pocket PC	Palm	Black Berry	Java GUI	Exchange	VCARD
Anniversary	Y	Y	N	N	N	Y	N
AssistantName	Y	Y	N	v	N	Y	N
AssistantTelephoneNumber	Y	Y	N	N	N	Y	N
Birthday	Y	Y	N	Y	Y	Y	BDAY
Body	Y	Y	Y	N	Y	Y	NOTE
Business2TelephoneNumber	Y	Y	N	N	Y	Y	TEL;VOICE;WORK
BusinessAddressCity	Y	Y	N	Y	Y	Y	ADR;WORK;-;-;-;-;CITY;-;-;-
BusinessAddressCountry	Y	Y	N	Y	Y	Y	ADR;WORK;-;-;-;-;-;COUNT RY
BusinessAddressPostalCode	Y	Y	N	Y	Y	Y	ADR;WORK;-;-;-;-;-;POSTALC ODE;-

<i>Property</i>	<i>Outlook</i>	<i>Pocket PC</i>	<i>Palm</i>	<i>Black Berry</i>	<i>Java GUI</i>	<i>Exchange</i>	<i>VCARD</i>
BusinessAddressPostOfficeBox	Y	N	N	N	N	Y	ADR;WORK:POSTOFFICE;:-;:-;:-
BusinessAddressState	Y	Y	N	Y	Y	Y	ADR;WORK:-;:-;:-;STATE;:-
BusinessAddressStreet	Y	Y	N	Y	Y	Y	ADR;WORK:-;:-;STREET;:-;:-
BusinessFaxNumber	Y	Y	Y	Y	Y	Y	TEL;FAX;WORK
BusinessLabel	N	N	N	N	Y	N	LABEL;WORK
BusinessTelephoneNumber	Y	Y	Y	Y	Y	Y	TEL;VOICE;WORK
BusinessWebPage	Y	N	N	N	Y	Y	URL;WORK
CallbackTelephoneNumber	Y	N	N	N	N	Y	N
CarTelephoneNumber	Y	Y	N	N	N	Y	TEL;CAR;VOICE
Categories	Y	Y	N	N	N	Y	N
Children	Y	Y	N	N	N	Y	N
CompanyMainTelephoneNumber	Y	N	N	N	N	Y	TEL;WORK;PREF
CompanyName	Y	Y	Y	Y	Y	Y	ORG:COMPANYNAME;-
Companies	Y	N	N	N	N	Y	N
ComputerNetworkName	Y	N	N	N	N	Y	N
Department	Y	Y	N	Y	Y	Y	ORG:-;DEPARTMENT
Email1Address	Y	Y	Y	N	Y	Y	EMAIL;INTERNET
Email1AddressType	Y	N	N	N	N	Y	N
Email2Address	Y	Y	N	N	Y	Y	EMAIL;INTERNET;HOME
Email2AddressType	Y	N	N	N	N	Y	N
Email3Address	Y	Y	N	N	Y	Y	EMAIL;INTERNET;WORK
Email3AddressType	Y	N	N	N	N	Y	N
FileAs	Y	Y	N	N	Y	Y	FN
FirstName	Y	Y	Y	Y	Y	Y	N:-;FIRSTNAME;:-;:-
Folder	Y	N	N	N	N	N	N
Gender	N	N	N	N	N	N	N
Hobby	Y	N	N	N	N	Y	N
Home2TelephoneNumber	Y	Y	N	N	Y	Y	TEL;VOICE;HOME
HomeAddressCity	Y	Y	Y	N	Y	Y	ADR;HOME:-;:-;:-;CITY;:-;:-
HomeAddressCountry	Y	Y	Y	N	Y	Y	ADR;HOME:-;:-;:-;:-;COUNTRY
HomeAddressPostalCode	Y	Y	Y	N	Y	Y	ADR;HOME:-;:-;:-;:-;POSTALCODE;-
HomeAddressPostOfficeBox	Y	N	N	N	N	Y	ADR;HOME:POSTOFFICE;:-;:-;:-
HomeAddressState	Y	Y	Y	N	Y	Y	ADR;HOME:-;:-;:-;STATE;:-
HomeAddressStreet	Y	Y	Y	N	Y	Y	ADR;HOME:-;:-;STREET;:-;:-;:-
HomeFaxNumber	Y	Y	N	N	Y	Y	TEL;FAX;HOME
HomeLabel	N	N	N	N	Y	N	LABEL;HOME
HomeTelephoneNumber	Y	Y	Y	Y	Y	Y	TEL;VOICE;HOME
HomeWebPage	Y	N	N	N	N	Y	URL;HOME
id	N	N	N	N	Y	N	N
Importance	Y	N	N	N	Y	Y	N
Initials	Y	N	N	N	N	Y	N
InstantMessenger	N	N	N	N	N	N	N
JobTitle	Y	Y	Y	Y	Y	Y	TITLE
Language	Y	N	N	N	N	Y	N

<i>Property</i>	<i>Outlook</i>	<i>Pocket PC</i>	<i>Palm</i>	<i>Black Berry</i>	<i>Java GUI</i>	<i>Exchange</i>	<i>VCARD</i>
LastName	Y	Y	Y	Y	Y	Y	N:LASTNAME;,-;-;
ManagerName	Y	N	N	N	N	Y	N
MiddleName	Y	Y	N	N	Y	Y	N:,-;-;MIDDLENAME;,-;-;
Mileage	Y	N	N	N	Y	Y	N
MobileTelephoneNumber	Y	Y	Y	Y	Y	Y	TEL;CELL
NickName	Y	N	N	N	Y	Y	NICKNAME
OfficeLocation	Y	Y	N	N	N	Y	N
OrganizationalIDNumber	Y	N	N	N	N	N	N
OtherAddressCity	Y	Y	N	N	N	Y	ADR:,-;-;-;CITY;,-;-;
OtherAddressCountry	Y	Y	N	N	N	Y	ADR:,-;-;-;-;-;COUNTRY
OtherAddressPostalCode	Y	Y	N	N	N	Y	ADR:,-;-;-;-;-;POSTALCODE;,-;-;
OtherAddressPostOfficeBox	Y	N	N	N	N	N	ADR:POSTOFFICE;,-;-;-;-;-;
OtherAddressState	Y	N	N	N	N	Y	ADR:,-;-;-;-;STATE;,-;-;
OtherAddressStreet	Y	Y	N	N	N	Y	ADR:,-;-;-;STREET;,-;-;-;
OtherFaxNumber	Y	N	N	N	N	Y	TEL;FAX
OtherLabel	N	N	N	N	N	N	N
OtherTelephoneNumber	Y	N	Y	N	Y	Y	TEL;VOICE
PagerNumber	Y	Y	Y	N	Y	Y	TEL;PAGER
PrimaryTelephoneNumber	Y	N	Y	N	N	N	TEL;PREF;VOICE
Profession	Y	N	N	N	Y	Y	ROLE
Revision	N	N	N	N	Y	N	REV
RadioTelephoneNumber	Y	Y	N	N	N	N	N
Sensitivity	Y	N	N	N	Y	Y	N
Spouse	Y	Y	N	N	N	Y	N
Subject	Y	N	N	N	N	Y	N
Suffix	Y	Y	N	N	Y	Y	N:,-;-;-;-;SUFFIX
TelexNumber	Y	N	N	N	N	Y	N
Timezone	N	N	N	N	N	N	TZ:
Title	Y	Y	N	N	Y	Y	N:,-;-;-;-;SALUTATION;,-;-;
Uid	N	N	N	N	N	N	UID:
WebPage	Y	Y	N	Y	N	Y	URL:
YomiCompanyName	Y	Y	N	N	N	N	N
YomiFirstName	Y	Y	N	N	N	N	N
YomiLastName	Y	Y	N	N	N	N	N

15.2. SIF-E

A calendar event represents an event scheduled for a day at a particular time or a series of days at a particular time. An example of a SIF event is:

```

<?xml version="1.0" encoding="UTF-8"?>
<appointment>
  <Start>20040930T133000Z</Start>
  <End>20040930T140000Z</End>
  <AllDayEvent>0</AllDayEvent>
  <Body/>
  <BusyStatus>2</BusyStatus>
  <Categories/>
  <Companies/>
  <Importance>1</Importance>
  <IsRecurring>0</IsRecurring>
  <Location>Milan</Location>

```

```

<MeetingStatus>0</MeetingStatus>
<Mileage/>
<ReminderMinutesBeforeStart>15</ReminderMinutesBeforeStart>
<ReminderSet>1</ReminderSet>
<ReminderSoundFile>CalenAlarmSound</ReminderSoundFile>
<ReminderOptions>8</ReminderOptions>
<ReminderInterval>5</ReminderInterval>
<ReminderRepeatCount>2</ReminderRepeatCount>
<ReplyTime/>
<Sensitivity/>
<Subject>Meeting with Maga developers</Subject>
<RecurrenceType>1</RecurrenceType>
<Interval>1</Interval>
<MonthOfYear>0</MonthOfYear>
<DayOfMonth>0</DayOfMonth>
<DayOfWeekMask>16</DayOfWeekMask>
<Instance>0</Instance>
<PatternStartDate>20040930T230000Z</PatternStartDate>
<NoEndDate>1</NoEndDate>
<Occurrences>0</Occurrences>
<PatternEndDate></PatternEndDate>
</appointment >

```

The fields defined and used by Sync4j are listed in the table below.

Property	Description
AllDayEvent	True if the appointment is an all-day event (as opposed to a specified time). Corresponds to the All day event check box on the Appointment page of an AppointmentItem. True is 1.
Body	Returns or sets the clear-text body of the item.
BusyStatus	Returns or sets the busy status of the user for the appointment. Can be one of the following OIBusyStatus constants: olBusy(2), olFree(0), olOutOfOffice(3), or olTentative(1).
Calscale	Defines the calendar scale used for the appointment
Categories	Returns or sets the categories assigned to the Outlook item.
Companies	Returns or sets the names of the companies associated with the Outlook item. This is a free-form text field
Contact	Represent the contact information or a reference to contact information associated with the appointment.
CREATED	Specifies the date and time that the calendar information was created by the calendar user agent in the calendar store.
DALARM	Display alarms that usually trigger a dialog box to be displayed by the client program
DCREATED	Specifies the date and time that the calendar information was created by the calendar user agent in the calendar store.
Dtstamp	Indicates the date/time that the instance of the appointment was created.
End	Returns or sets the end date and time of an appointment or journal entry. Expressed in UTC YYYYMMDDTHHMMSSZ
Geo	Specifies information related to the global position for the activity specified by an appointment.
id	Returns or sets the identifier of the calendar item
Importance	Returns or sets the relative importance level for the Outlook item. Can be one of the following OIImportance constants: olImportanceHigh(2), olImportanceLow(0), or olImportanceNormal(1). This property corresponds to the MAPI property PR_IMPORTANCE.
IsRecurring	True if the appointment is recurring. True is 1.
Last-modified	Specifies the date and time that the information associated with the appointment was last revised in the calendar store
Location	Returns or sets the specific location
MeetingStatus	Returns or sets an OIMeetingStatus constant specifying the meeting status of the appointment. The constants are: olNonMeeting (0), olMeeting (1), olMeetingReceived (3), olMeetingCanceled (5)
Mileage	Returns or sets a String representing the mileage for an item. This is a free-form string field and can be used to store mileage information associated with the item (for example, 100 miles documented for an appointment, contact, or task) for purposes of reimbursement.
PALARM	Is a procedure reminder, or application executable that will be run as an alarm for a calendar component.
ProdId	Specifies the identifier for the product that created the appointment.

Property	Description
ReminderInterval	Returns or sets the interval in which the reminder has to be repeated.
ReminderMinutesBeforeStart	Returns or sets the number of minutes the reminder should occur prior to the start of the appointment
ReminderOptions	Returns or sets the type of a reminder. Sum of any of the following constants. <i>oLED</i> activates the LED (light emitting diode) on a device. <i>oVibrate</i> activates any vibration indicator on a device. <i>oDialog</i> displays a dialog. <i>oSound</i> plays the file specified by <i>ReminderSoundFile</i> . <i>oRepeat</i> repeats the reminder.
ReminderRepeatCount	Returns or sets the number of times that the reminder has to be repeated.
ReminderSet	True if a reminder has been set for this appointment, item or task.
ReminderSoundFile	Returns or sets the path and file name of the sound file to play when the reminder occurs for the Appointment. This property is valid only if the <i>ReminderSet</i> property is TRUE and the <i>ReminderOptions</i> property includes <i>oSound</i> . The default for this is the current setting for the Calendar application or Alarm1.wav if none.
ReplyTime	Returns or sets a Date indicating the reply time for the appointment. Read/write
Revision	Represents the date and time that this event's information was last modified
Rrule	Defines a rule or repeating pattern for recurring events, to-dos, or time zone definitions.
Sensitivity	Returns or sets the sensitivity for the Outlook item. Can be one of the following <i>OSensitivity</i> constants: <i>oConfidential</i> (3), <i>oNormal</i> (0), <i>oPersonal</i> (1), or <i>oPrivate</i> (2). This property corresponds to the MAPI property PR_SENSITIVITY
Sequence	The sequence field specifies the sequence number of a version of an appointment.
Start	Returns or sets the starting date and time for the appointment or journal entry. Expressed in UTC YYYYMMDDTHMMSSZ
Status	Is the overall status or confirmation of the appointment. Can be one of the following values: Tentative, Confirmed, Cancelled.
Subject	The subject for the appointment
Timezoneid	The <i>timezoneid</i> field specifies the time zone identifier of an appointment or meeting
Transp	Defines whether an event is transparent or not to busy time searches.
Uid	Defines the persistent, globally unique identifier for the appointment.
Url	Defines a Uniform Resource Locator (URL) associated with the appointment.
Version	Specifies the identifier corresponding to the highest version number or the minimum and maximum range of the iCalendar specification that is required in order to interpret the iCalendar object
Recurrence properties	
DayOfMonth	The single day of the month from 1 to 31.
DayOfWeekMask	The combination days of the week constants (i. e. event recurring on Monday and Wednesday. The <i>DayOfWeekMask</i> should be <i>oMonday</i> + <i>oWednesday</i>)
Interval	Is the interval of the recurrence. If <i>RecurrenceType</i> is <i>oRecurDaily</i> , event occurs every <interval> day. If <i>oRecurWeekly</i> , event occurs every <interval> week..
Instance	The ordinal number of the day, week, month
MonthOfYear	Month of year
NoEndDate	True if there is no end date. True is 1.
Occurrences	Number of occurrences for the recurrence.
PatternEndDate	The end date of the recurrence
PatternStartDate	The start day of the recurrence
RecurrenceType	Returns or set a <i>RecurrenceType</i> . See below for the possible values.

The following table lists all fields and their use on clients:

Property	Outlook	Pocket PC	BlackBerry	Java GUI	Exchange	iCAL
AllDayEvent	Y	Y	N	N	Y	N
Body	Y	Y	Y	Y	Y	DESCRIPTION
BusyStatus	Y	Y	N	N	Y	N

<i>Property</i>	<i>Outlook</i>	<i>Pocket PC</i>	<i>BlackBerry</i>	<i>Java GUI</i>	<i>Exchange</i>	<i>iCAL</i>
Calscale	N	N	N	N	N	CALSCALE
Categories	Y	Y	N	N	Y	CATEGORIES
Companies	Y	N	N	N	N	ORGANIZER
Contact	N	N	N	N	N	CONTACT
CREATED	N	N	N	N	N	CREATED
DALARM	N	N	N	N	N	DALARM
DCREATED	N	N	N	N	N	DCREATED
Dtstamp	N	N	N	N	N	DTSTAMP
End	Y	Y	Y	Y	Y	DTEND
Geo	N	N	N	N	N	GEO:-;-
id	N	N	N	Y	N	N
Importance	Y	N	N	N	Y	PRIORITY
IsRecurring	Y	Y	N	N	N	N
Last-Modified	N	N	N	N	N	LAST-MODIFIED
Location	Y	Y	Y	Y	Y	LOCATION
MeetingStatus	Y	N	N	N	Y	N
Method	N	N	N	N	N	METHOD
Mileage	Y	N	N	N	Y	N
PALARM	N	N	N	N	N	PALARM
ProdId	N	N	N	N	N	PRODID
ReminderInterval	N	N	N	N	N	AALARM:-;INTERVAL:-;;
ReminderMinutesBeforeStart	Y	Y	Y	N	Y	Is used with DTSTART to calculate the reminder start date.
ReminderOptions	N	Y	N	N	N	N
ReminderRepeatCount	N	N	N	N	N	AALARM:-;;COUNT:-;
ReminderSet	Y	Y	N	N	Y	AALARM
ReminderSoundFile	N	Y	N	N	N	AALARM:-;-;SOUNDFILE
ReplyTime	Y	N	N	N	Y	N
Revision	N	N	Y	N	N	N
Rrule	N	N	N	N	N	RRULE
Sensitivity	Y	Y	N	N	Y	CLASS
Sequence	N	N	N	N	N	SEQUENCE
Start	Y	Y	Y	Y	Y	DTSTART
Status	N	N	N	N	N	STATUS
Subject	Y	Y	Y	Y	Y	SUMMARY
Timezoneld	N	N	N	N	N	N
Transp	N	N	N	N	N	TRANSP
Uid	N	N	N	N	N	UID
Url	N	N	N	N	N	URL
Version	N	N	N	N	N	VERSION
Recurrence properties						
DayOfMonth	Y	Y	N	N	N	N
DayOfWeekMask	Y	Y	N	N	N	N
Interval	Y	Y	N	N	N	N
Instance	Y	Y	N	N	N	N
MonthOfYear	Y	Y	N	N	N	N

<i>Property</i>	<i>Outlook</i>	<i>Pocket PC</i>	<i>BlackBerry</i>	<i>Java GUI</i>	<i>Exchange</i>	<i>iCAL</i>
NoEndDate	Y	Y	N	N	N	N
Occurrences	Y	Y	N	N	Y	N
PatternEndDate	Y	Y	N	N	N	N
PatternStartDate	Y	Y	N	N	N	N
RecurrenceType	Y	Y	N	N	N	N

15.2.1. Constants

The following constants are defined:

OlDaysOfWeek

```

olSunday           = 1;
olMonday           = 2;
olTuesday          = 4;
olWednesday        = 8;
olThursday         = 16;
olFriday           = 32;
olSaturday         = 64;

```

OlRecurrenceType

```

olRecurDaily       = 0;
olRecurWeekly      = 1;
olRecurMonthly     = 2;
olRecurMonthNth    = 3;
olRecurYearly      = 5;
olRecurYearNth     = 6;

```

OlSensitivity

```

olNormal           = 0;
olPersonal          = 1;
olPrivate           = 2;
olConfidential     = 3;

```

OlBusyStatus

```

olFree             = 0;
olTentative        = 1;
olBusy             = 2;
olOutOfOffice      = 3;

```

OlImportance

```

olLow              = 0;
olNormal           = 1;
olHigh             = 2;

```

15.2.2. Recurrent event examples

An event happening every 2 weeks on Sunday and Monday:

```

RecurrenceType     = olRecurWeekly
Instance           = 2
DayOfWeekMask      = olSunday + olMonday

```

An event scheduled the 2th Wednesday of April of every year:

```

RecurrenceType = olRecurYearNth
Interval       = 12
MonthofYear   = 4
DayOfWeekMask = olWednesday
Instance      = 2
NoEndDate     = True

```

To learn more about recurrence and how setting all the properties have a look at: <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/vbaol10/html/olobjRecurrencePattern.asp>

Note for Outlook and PPC Client: the recurrent properties must be set in a proper order and only some properties must be set related to the particular RecurrenceType property. If not, no event can be added successfully.

RecurrenceType must be set as the first property and outlook client do it getting the value from the SIF item.

Keeping “**olRecurWeekly**” as example only the properties “*Interval*” and “*DayOfWeekMask*” have to be set meanwhile all the others must be 0. The Outlook Client skips the 0 value properties.

The same thing is for “**olRecurMonthNth**”: only the “*Interval*”, “*Instance*” and “*DayOfWeekMask*” must be set to create a right event.

15.3. SIF-T

A Task object is very similar to an event object. An example of SIF task is the following:

```

<?xml version="1.0" encoding="UTF-8"?>
<task>
  <Body/>
  <Categories/>
  <Companies>Maga S.p.A.</Companies>
  <Complete>0</Complete>
  <DueDate>20041008T230000Z</DueDate>
  ..<DueDate>20041008T230000Z </DueDate>
  ..<BillingInformation>information</BillingInformation>
  ..<ActualWork>10</ActualWork >
  <Importance>1</Importance>
  <IsRecurring>0</IsRecurring>
  <Mileage/>
  <PercentComplete>0</PercentComplete>
  <ReminderSet>1</ReminderSet>
  <ReminderTime/>
  <Sensitivity/>
  <StartDate>20041110T230000Z</StartDate>
  <Status>0</Status>
  <Subject>new task</Subject>
  <TeamTask>0</TeamTask>
  <TotalWork>0</TotalWork>
  <RecurrenceType>1</RecurrenceType>
  <Interval>1</Interval>
  <MonthOfYear>0</MonthOfYear>
  <DayOfMonth>0</DayOfMonth>
  <DayOfWeekMask>4</DayOfWeekMask>
  <Instance>0</Instance>
  <PatternStartDate>20040930T230000Z</PatternStartDate>
  <NoEndDate>1</NoEndDate>
  <Occurrences>0</Occurrences>
  <PatternEndDate></PatternEndDate>
</task>

```

The fields defined and used by Sync4j are listed in the following table:

<i>Property</i>	<i>Description</i>
ActualWork	Returns or sets the actual effort (in minutes) spent on the task.
BillingInformation	Returns or sets the billing information associated with the Outlook item. This is a free-form text field
Body	Returns or sets the clear-text body of the Outlook item.
Categories	Returns or sets the categories assigned to the Outlook item.
Companies	Returns or sets the names of the companies associated with the Outlook item. This is

Property	Description
	a free-form text field
Complete	True if the task is completed. True is 1
DueDate	Returns or sets a Date indicating the due date for the task.
DateCompleted	Returns when the task is completed. It gets set to the current date on the device when you set the Complete property
Importance	Returns or sets the relative importance level for the Outlook item. Can be one of the following OIImportance constants: oIImportanceHigh(2), oIImportanceLow(0), or oIImportanceNormal(1). This property corresponds to the MAPI property PR_IMPORTANCE.
Mileage	Returns or sets a String representing the mileage for an item. This is a free-form string field and can be used to store mileage information associated with the item (for example, 100 miles documented for an appointment, contact, or task) for purposes of reimbursement.
PercentComplete	Returns or sets the percentage of the task completed at the current date and time
ReminderSet	True if a reminder has been set for this appointment, mail item or task. True is 1
ReminderTime	Returns or sets the date and time at which the reminder should occur for this item.
ReminderSoundFile	Returns or sets the path and file name of the sound file to play when the reminder occurs for the Appointment. This property is valid only if the ReminderSet property is TRUE and the ReminderOptions property includes oISound. The default for this is the current setting for the Calendar application or Alarm1.wav if none.
ReminderOptions	Returns or sets the type of a reminder. Sum of any of the following constants. oILED activates the LED (light emitting diode) on a device. oIVibrate activates any vibration indicator on a device. oIDialog displays a dialog. oISound plays the file specified by ReminderSoundFile. oIRepeat repeats the reminder.
IsRecurring	True if the appointment is recurring. True is 1
Sensitivity	Returns or sets the sensitivity for the Outlook item. Can be one of the following OISensitivity constants: oIConfidential(3), oINormal(0), oIPersonal(1), or oIPrivate(2). This property corresponds to the MAPI property PR_SENSITIVITY
StartDate	Returns or sets the starting date and time for the task
Status	Returns or sets the status for the task. Can be one of the following OITaskStatus constants: oITaskComplete(2), oITaskDeferred(4), oITaskInProgress(1), oITaskNotStarted(0), or oITaskWaiting(3).
Subject	Returns or sets the subject for the Outlook item. This property corresponds to the MAPI property PR_SUBJECT. The Subject property is the default property for Outlook items. IT IS READ ONLY FOR NOTES
TeamTask	True if the task is a team task. True is 1
TotalWork	Returns or sets the total work for the task
Recurrence properties	
DayOfMonth	For recurrence (see calendar property)
DayOfWeekMask	For recurrence (see calendar property)
Interval	Return the interval
Instance	For recurrence (see calendar property)
MonthOfYear	For recurrence (see calendar property)
NoEndDate	For recurrence (see calendar property)
Occurrences	For recurrence (see calendar property)
PatternStartDate	For recurrence (see calendar property)
PatternEndDate	For recurrence (see calendar property)
RecurrenceType	Returns or set a RecurrenceType. values are orRecursDaily...

The following table lists all fields and their use on clients:

Property	Outlook	Pocket PC	Exchange
ActualWork	Y	Y	Y
BillingInformation	Y	Y	Y

<i>Property</i>	<i>Outlook</i>	<i>Pocket PC</i>	<i>Exchange</i>
Body	Y	Y	Y
Categories	Y	Y	Y
Companies	Y	N	N
Complete	Y	Y	Y
DueDate	Y	Y	Y
DateCompleted	Y	Y	Y
Importance	Y	Y	Y
Mileage	Y	N	Y
PercentComplete	Y	N	Y
ReminderSet	Y	Y	Y
ReminderTime	Y	Y	Y
ReminderSoundFile	N	Y	N
ReminderOptions	N	Y	N
IsRecurring	Y	Y	N
Sensitivity	Y	Y	Y
StartDate	Y	Y	Y
Status	Y	N	Y
Subject	Y	Y	Y
TeamTask	Y	Y	Y
TotalWork	Y	N	N
Recurrence properties			
DayOfMonth	Y	Y	N
DayOfWeekMask	Y	Y	N
Interval	Y	Y	N
Instance	Y	Y	N
MonthOfYear	Y	Y	N
NoEndDate	Y	Y	N
Occurrences	Y	Y	Y
PatternStartDate	Y	Y	N
PatternEndDate	Y	Y	N
RecurrenceType	Y	Y	N

15.4. SIF-N

SIF notes are also represented in a SIF format:

```
<?xml version="1.0" encoding="UTF-8"?>
<note>
  <Body>New Note the first note</Body>
  <Categories/>
  <Subject>New Note</Subject>
  <Color>3</Color>
  <Height>166</Height>
  <Width>200</Width>
  <Left>80</Left>
  <Top>80</Top>
</note>
```

The fields defined and used by Sync4j are listed in the following table.

<i>Property</i>	<i>Description</i>
Body	Returns or sets the clear-text body of the note item.

<i>Property</i>	<i>Description</i>
Categories	Returns or sets the categories assigned to the note item.
Color	Color of note
Date	Date of received note
Height	Height of the box note
Left	Left position of the box of the note
Subject	Returns or sets the subject for the note item. This property corresponds to the MAPI property PR_SUBJECT. The Subject property is the default property for Outlook items. IT IS READ ONLY FOR NOTES. Its value is retrieved by the first line of the body.
Top	Top position of the box of the note
Width	Width of the box note

The following table lists all fields and their use on clients:

<i>Property</i>	<i>Outlook</i>	<i>Pocket PC</i>	<i>Exchange</i>
Body	Y	Y	Y
Categories	Y	N	N
Color	Y	N	N
Date	N	N	Y
Height	Y	N	N
Left	Y	N	N
Subject	Y	Y	Y
Top	Y	N	N
Width	Y	N	N

15.4.1. Constants

The following constants are defined:

OINoteColor

```

olBlue           = 0;
olGreen          = 1;
olPink           = 2;
olYellow         = 3;
olWhite          = 4;

```

16. Appendix B – List of acronyms

API	Application Programming Interface
CTP	Client TCP Push
CVS	Concurrent Versions System
DB	DataBase
DS	Data Synchronization
GUI	Graphical User Interface
GUID	Globally Unique Identifier
HTTP	HyperText Transfer Protocol
IDE	Integrated Development Environment
IMAP	Internet Message Access Protocol
IP	Internet Protocol
JDK	Java Development Kit
JVM	Java Virtual Machine
LUID	Local Unique Identifier
OMA	Open Mobile Alliance
OTA	Over The Air
PDA	Personal Digital Assistant
PIM	Personal Information Management
POM	Project Object Model
POP	Post Office Protocol
REST	REpresentational State Transfer
SDK	Software Development Kit
SIF	Sync4j Interchange Format
SMS	Short Message Service
SMTP	Simple Mail Transfer Protocol
TCP	Transmission Control Protocol
UI	User Interface
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
WS	Web Service
XML	eXtensible Markup Language

17. Resources

This section lists resources you may find useful.

- [1] *SyncML Representation Protocol, version 1.1*,
http://www.syncml.org/docs/syncml_represent_v11_20020215.pdf
- [2] *SyncML Sync Protocol, version 1.1*,
http://www.syncml.org/docs/syncml_sync_protocol_v11_20020215.pdf
- [3] Funambol Administration Guide
- [4] Apache maven, <http://maven.apache.org/index.html>
- [5] Community projects, <http://www.funambol.com/opensource/projects.php>
- [6] Funambol public CVS, http://forge.objectweb.org/scm/?group_id=96
- [7] Funambol Software Development Kit, <http://download.forge.objectweb.org/sync4j/funambol-sdk-7.0.0-SNAPSHOT.tar.gz>
- [8] OpenXchange connector, <https://funamboloxconnector.forge.funambol.org>
- [9] Exchange connector, <https://exchange-connector.forge.funambol.org>
- [10] AGPL, <http://www.fsf.org/licensing/licenses/agpl-3.0.html>