



# **J2ME SyncML API Design Document**

November 2006

## Changes History

<i>Date</i>	<i>Author</i>	<i>Changes</i>
2006.11.06	Andrea Gazzaniga	Initial Draft
2006.11.06	Giuseppe Monticelli	Server Alerted Notification
2006.11.09	Ivano Brogonzoli	OTA Configuration
2006.11.10	Andrea Gazzaniga	Version 1.0 (first review completed)
2006.11.15	Giuseppe Monticelli	SAN Java examples, UML class diagrams
2006.11.16	Andrea Gazzaniga	Editorial changes
2006.11.17	Stefano Fornari	Editorial and content review
2006.12.04	Andrea Gazzaniga	Version 1.1 (after public review)
2007.12.07	Edoardo Schepis	SAN and OTAConfig paragraphs moved to common package documentation Added Large Object documentation
2008.09.16	Ivano Brogonzoli	Added The MappingManager part. It explains how lost mappings informations are recovered by the API and how it solves the problem of duplicated items.

# **Table of Contents**

1. <a href="#">Overview</a> .....	4
1.1. <a href="#">Document Conventions</a> .....	4
1.1.1. <a href="#">Sequence Diagrams</a> .....	4
1.1.2. <a href="#">Class Diagrams</a> .....	4
1.2. <a href="#">Functional overview</a> .....	4
2. <a href="#">Data Synchronization Layer</a> .....	6
2.1. <a href="#">The Synchronization Process</a> .....	6
2.1.1. <a href="#">Initialization</a> .....	7
2.1.2. <a href="#">Modifications Exchange</a> .....	9
2.1.3. <a href="#">LUID-GUID Mapping</a> .....	10
2.2. <a href="#">Data Synchronization Layer Architecture</a> .....	10
2.2.1. <a href="#">SyncManager</a> .....	11
2.2.2. <a href="#">The SyncSource Interface</a> .....	12
2.2.3. <a href="#">BaseSyncSource</a> .....	13
2.2.4. <a href="#">SyncConfig</a> .....	14
2.2.5. <a href="#">DeviceConfig</a> .....	14
2.2.6. <a href="#">SourceConfig</a> .....	15
2.3. <a href="#">Synchronization Events Notification</a> .....	16
2.4. <a href="#">Client Capabilities Handling</a> .....	17
2.5. <a href="#">Filtering</a> .....	17
2.6. <a href="#">Large Objects Handling</a> .....	18
3. <a href="#">References</a> .....	20

# 1. Overview

The Funambol SyncML API allows application developers to embed SyncML capabilities in a J2ME application, thus giving it the access to a powerful data synchronization framework to keep the application data in sync with other devices.

This document explains, from a developer point of view, the architecture of the Funambol SyncML API for Java 2 Micro Edition.

## 1.1. Document Conventions

The diagrams used in this document are inspired to the UML sequence and class diagrams, but with some simplification. The conventions used by the diagrams are described in the following sections.

### 1.1.1. Sequence Diagrams

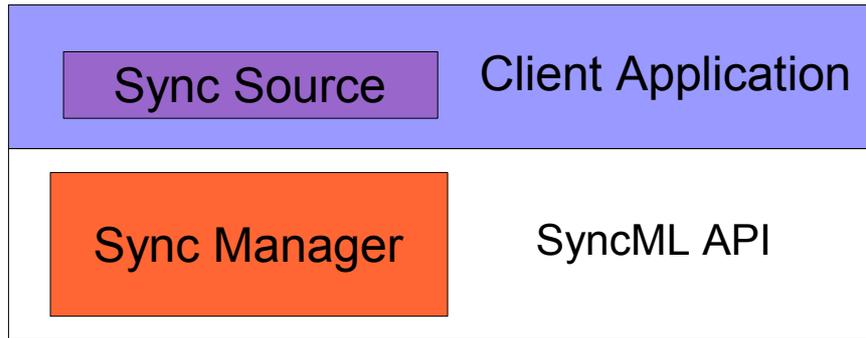
- Each entity is represented as a box;
- a box can represent a class, an instance, an interface or even a conceptual entity; the real meaning depends by the context;
- solid arrows represents method or function calls;
- dashed arrows represent some sort of communication between two entities; it is intended that the communication mechanism is left unspecified or is not important or it is at a different abstraction layer.

### 1.1.2. Class Diagrams

- Each class is represented as a box;
- data members and methods are separated by a horizontal line;
- plain titles represent classes, italicized titles represent interfaces (abstract classes);
- a + next to a method or data member name means "public"
- a - next to a method or data member name means "private"
- a # next to a method or data member name means "protected";
- a > next to a data member name means it is a property with get/set accessors;
- inheritance is represented by an arrow pointing to the base class;
- italicized methods names represent abstract methods.

## 1.2. Functional overview

The Funambol SyncML API for J2ME is currently focused on data synchronization only. The device management protocol (OMA-DM), available on other platforms, is not implemented in this version of the API.



*Figure 1.1 - SyncClient API architecture*

A client application interacts mainly with two entities of the Funambol Client API: the SyncManager and the SyncSource. The SyncManager is the component that handles all the communication and protocol stuff. It hides the complexity of the synchronization process providing a simple interface to the client application. A SyncSource represents the collection of items stored in the local repository. It contains the client logic to discover the items to send to the server and to store the ones obtained from the server. The client feeds a SyncSource with the items changed on the client side, whilst the SyncManager feeds it with the items received by the server.

## 2. Data Synchronization Layer

This chapter describes the design of the Funambol J2ME SyncML API Data Synchronization Layer.

The Data Synchronization process implemented in the Funambol API follows the OMA DS 1.2 specification (see [OMA-DS] and related documents).

### 2.1. The Synchronization Process

In this section, a brief introduction to the synchronization process is provided to better understand the use of the SyncML API. For a more detailed description see the OMA DS specification ([OMA-DS]).

The synchronization process is logically a sequence of three phases:

1. Initialization
2. Modifications exchange
3. IDs Mapping

In the initialization phase the client sends its credentials and which database to synchronize (along with the desired synchronization type) to the server. The server responds with the authentication status and the synchronization type to perform. The client must then perform the synchronization as requested by the server.

After the initialization phase, first the client sends all client-side modifications and receives the status of the execution of the commands on the server; then, the server sends server-side modifications. The client applies the changes and sends the proper status back to the server.

In the case the server issued a new item to the client, the latter creates a new local id for it and therefore needs to communicate such new key to the server (see LUID-GUID mapping in [OMA-DS]). This is usually done in the mapping phase (or at least, the Funambol client API will send the mapping just after the modifications exchange phase).

The synchronization flow is illustrated by the sequence diagram of Figure 2.1Error: Reference source not found, where the entities that participates to the sync have the following meaning:

- Client Application: a final client application using the Funambol SyncML API;
- SyncManager: the SyncML API synchronization engine;
- SyncSource: the interface for a client data source;
- SyncListener: the interface for monitoring a synchronization process;
- Sync Server: the SyncML server

As shown in the figure, the client application just kicks off the synchronization process giving a sync source to synchronize. The SyncManager takes control of the synchronization process. The macros initialization, modifications exchange and client mapping involve the interaction with the server and include the required SyncML messages exchange.

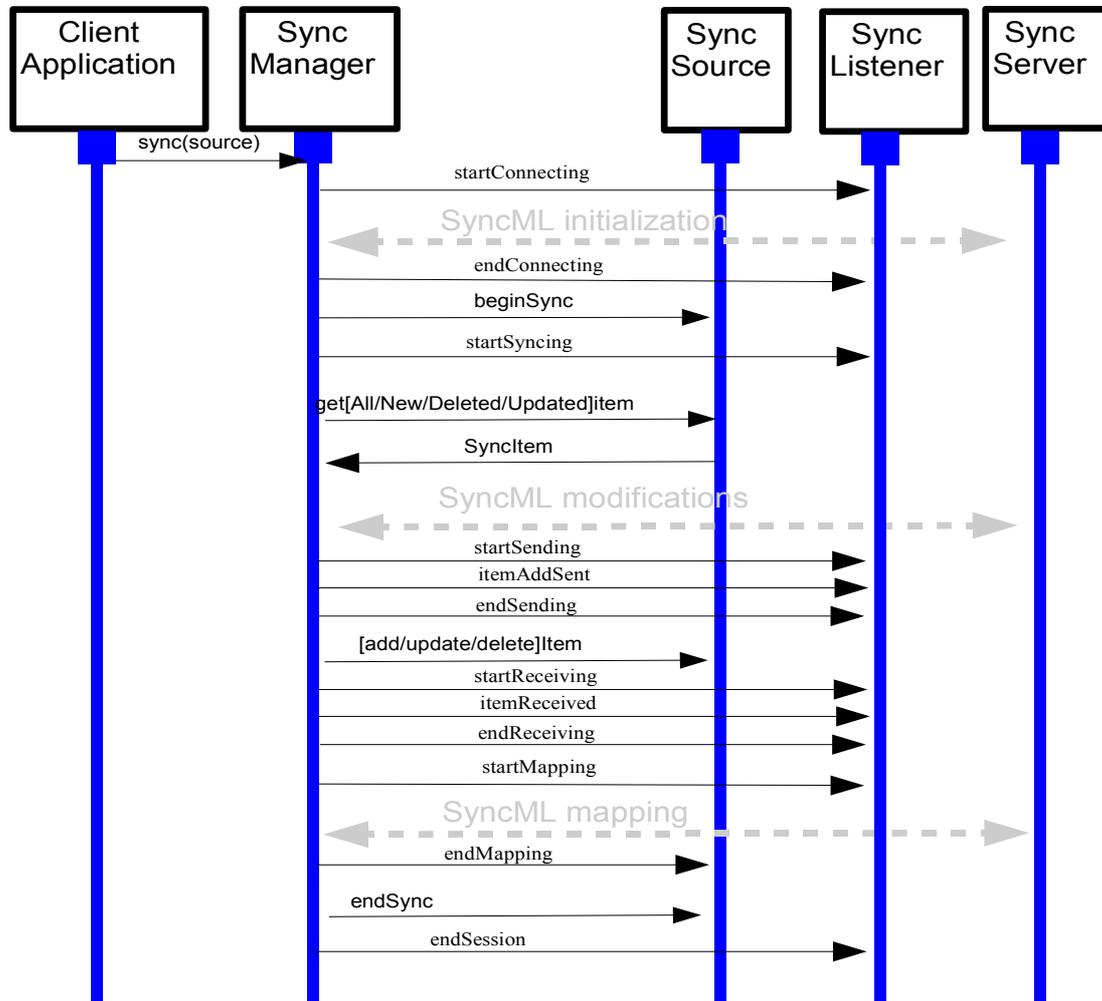


Figure 2.1 - Synchronization process flow

To simplify the figure, some back arrows are not shown as they do not carry information (for example the ones from the listener back to the SyncManager).

### 2.1.1. Initialization

As per the SyncML specification, the initialization phase can be performed in two ways:

1. As a separate package
2. Together with the modifications exchange package

The Funambol Client API implements “separate initialization” only, which is the option that optimizes at best network usage in the most common cases. In fact, with a synchronization without a separate initialization, there is the risk to start a potentially long synchronization, while the server would refuse the sync (for example because it does not authorize the client). Performing separate initialization avoids this issue with a minimal impact on network traffic.

The specification defines three important tasks performed in the initialization phase:

1. client authentication
2. server authentication
3. database alerting

The server authentication is not implemented in the current version of the J2ME SyncML API, so it will not be further described.

### Client Authentication

This section covers how the client sends its authentication credentials to the server and how the authentication process goes.

The SyncML specifications mandate that client implementations must support at least *basic* and *MD5 Digest* authentication.

Client authentication is delivered in the Cred element of the SyncML message header. Plus, even if the client sends in the first message its credentials in one of the supported types, the server can refuse it and challenge the client for a different kind of credentials. For example, if the client starts sending Basic credentials and the server requires MD5 authentication, the server responds with a 401 status to the client credentials and provides a Chal element with the requested authentication.

The Funambol J2ME SyncML API implements only the basic one at the moment. If the server requires an MD5 authentication, the client recognize the request and report an an invalid authentication method exception.

For additional information on Basic and MD5 authentication see section 2.5 of [OMA-DS].

Basic authentication is identified by the URI syncml:auth-basic; it is pretty simple and it is very similar to what happens in Web applications.

In this case, credential data are encoded as follows:

---

```
B64(username': 'password)
```

---

Where B64() is a function that encodes the given string in Base 64. *username* and *password* represent the account's authentication information.

### Database Alerting

The second task performed during initialization is database alerting. Database alerting is the means the client requests to synchronize a particular database; plus, it specifies which type of synchronization should be performed. The SyncML specifications define the following synchronization types:

<b>Sync Type</b>	<b>Description</b>	<b>Alert Code</b>
Two-way	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.	200
Slow	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.	201
One-way 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.	202
Refresh 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.	203
One-way 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.	204
Refresh 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.	205

From the specifications perspective only Slow and Two-way sync types are mandatory. However, the Funambol SyncML API supports issuing all the defined synchronization types.

Note that the server can respond with a different Alert type for a given database. The client must always perform the synchronization type as given by the server.

## 2.1.2. Modifications Exchange

The modification phase is when client modifications are detected on the client and sent to the server who replies with the server side updates. These updates are applied to the local database. Accordingly to the SyncML specs, updates are exchanged with commands like Add, Replace, Delete, Copy, Move, etc.

However only Add, Replace and Delete are mandatory to SyncML implementations; the others are optional (see the specs for details).

Since the initialization phase is done separately, the modification exchange phase is carried on only if initialization completed successfully.

The real synchronization process is driven by the SyncManager. It uses the SyncSource interface in order to interact with a specific data source. The development of the SyncSource is delegated to the application developer. The role of a SyncSource is to

- provide information about the SyncSource to synchronize (preferred synchronization type, name, local and target URIs, MIME type of the data...)
- keep the synchronization state from the data source point of view
- retrieve all items in the data source regardless their state
- retrieve the modified items only from the data source
- apply server side changes
- commit changes

The interaction between the SyncManager and the SyncSource is the core of the synchronization process. Since the number of items in the local data source may be relevant, and thinking of the limited resources of a J2ME device, the SyncSource works like a cursor: it provides getItem()/getNewItem()/getUpdatedItem()/getDeletedItem() to retrieve the first and next items in the data store. The SyncSource implementation must reset the cursor after a call to the beginSync() method, and return the next item of the requested type at each call.

In this way the SyncManager has the opportunity to choose to add the item in the current message or in the next one, to implement the multi-message mechanism, described later.

When the server receives a modification from the client, it applies the change in the server database and returns the status code for the operation. The SyncManager communicates to the SyncSource the status of the performed command by calling setItemStatus() so that the SyncSource can take appropriate actions (for example, in the case of a success code, it can reset the item dirty flag or permanently delete a soft deleted item).

On the opposite direction, when modifications are received from the server, the SyncManager decodes each command and call add/update/deleteItem() on the SyncSource implementation. These calls apply the change and return a status code (the same value that will be sent in the SyncML message).

Note that addItem() returns also the new local ID generated by the data store (LUID); this is stored in the SyncSource along with the GUID received by the server so that the mapping will be communicated back to the server in a Map command (see the next phase).

### Multiple messages in one package

OMA-DS specification defines a way to split the modification package over multiple messages, which means the client can send its modified items in more messages and get the server items in the same way. This process is commonly called "multi-messaging" and it is used to meet the limited resources requirements of mobile devices. The Funambol J2ME SyncML API supports multi-messaging, as summarized below:

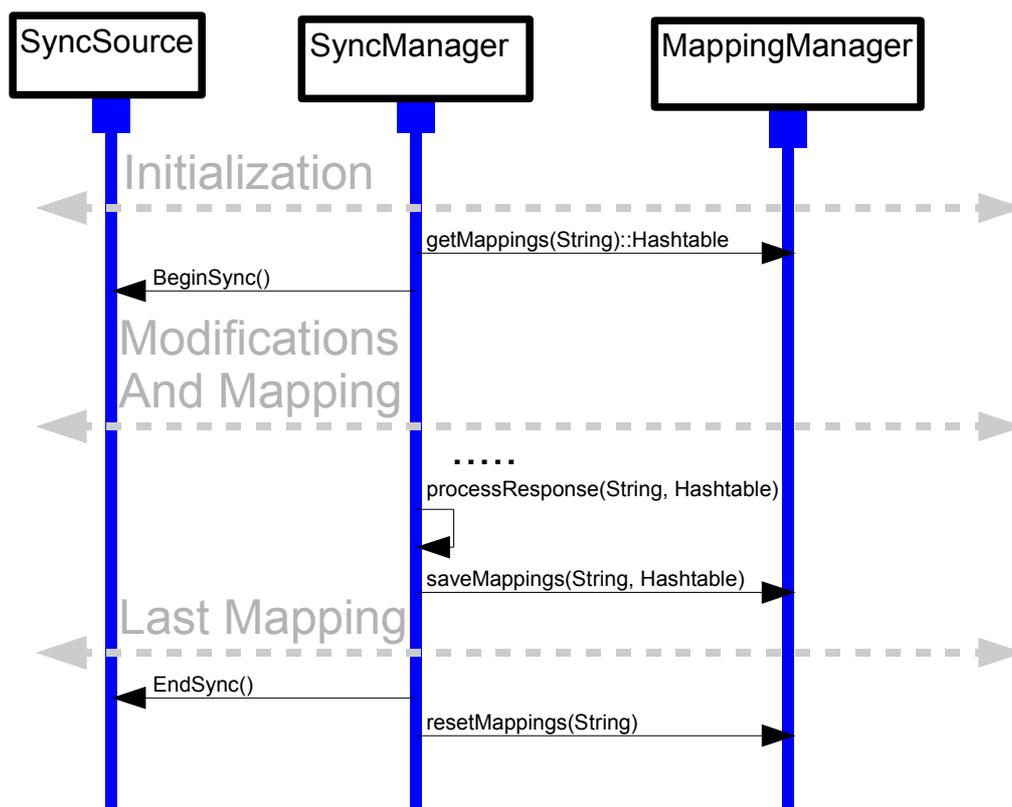
1. Client sends the first n items in the first message of PKG #3, without the <Final/> element;
2. Server applies the changes and replies with a message containing only status commands; the server is not allowed to send its own changes until client finished to send client side changes (which means unless the client ends the package with the <Final/> element);

3. Client sends the next message with the next items; if the message contains the last items, the package is terminated with the <Final/> tag;
4. Server applies the client changes and returns a message with only status; if the client sent its last message, Server goes to the next step; otherwise, the process goes back to step 2;
5. Client replies to Server simply with a status command (optionally with a 222 Alert code);
6. Server starts sending server side updates; if the message is not the last one, the <Final/> element is not added to the message; otherwise, <Final/> indicates the end of the package;
7. Client applies the received changes and replies with status command only; when the message obtained by Server contains the <Final/> element, the process ends.

### 2.1.3. LUID-GUID Mapping

After the server has done with its modifications, along with the latest status command, the client is requested to send also the LUID-GUID mapping, if there is any. As already mentioned, such mapping is stored temporarily in the SyncSource implementation's instance, therefore the SyncManager can easily extract them and embed them into the message. However, it may happen that the communication falls in the middle of sending the message or before receiving the response from the server. For this reason, mappings must be permanently stored before sending them and can be removed only after having received a successful status for the sent Map command.

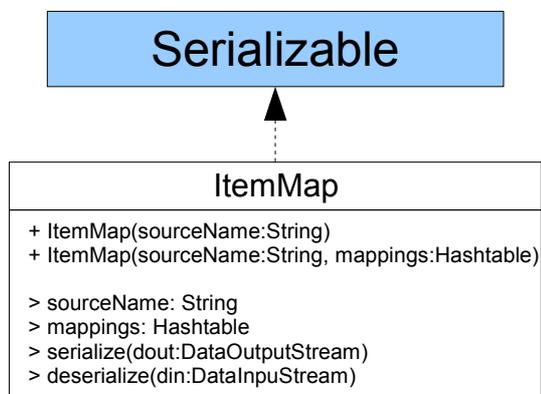
If, when a new synchronization starts, the persistent mapping is still present, it must be sent again. The Class that manages this task is called MappingManager and it is instantiated by the sync manager, that invokes it during the sync process as per the diagram below:



The MappingManager class refers to a particular Serializable object called the ItemMap, that contains the name of the source mappings to be persisted and the hashtable with the mappings informations.

Before the sync begins the SyncManager checks if there are mappings to be sent for the given source. Due to the fact that the sync implements the multimessage technique every time a mapping information must be sent, the SyncManager saves the mappings before sending the message: when a network problem occurs the persisted mappings will be sent on the next sync process only if the next sync type is two way or one way from server (in case of slow syncs or refresh syncs no mapping message recovery is needed). The call to the reset method is done when the sync ended successfully.

The ItemMap class is owned by the MappingManager and its diagram is shown below:



The ItemMap class implements the Serializable interface. The two methods serialize(..) and deserialize(...) are implicitly used by the store when the source mappings are persisted into the RMS. Due to performance issues the mappings table is persisted every time a mapping list is filled after the response by the server is processed. This is a bit different from the C++ API implementation that saves every mappings anytime a new item is added on the client. Doing the same on the RMS would create loss of performance into the sync process, so it has been decided to implement the mapping caching once per message.

## 2.2. Data Synchronization Layer Architecture

This section defines the design of the data synchronization layer. The fundamental components are: SyncManager, SyncSource and SyncItem. Roles and responsibilities of these components are summarized in the table below.

<b>Component</b>	<b>Implementing class / Interface</b>	<b>Roles and responsibilities</b>
SyncManager	com.funambol.syncml.spds.SyncManager	This is responsible for the implementation of each single phase of the synchronization. It is also responsible of driving the process of building the proper SyncML messages to send to the server through the TransportAgent and to interpret the server response and take the consequent actions.
SyncSource	com.funambol.syncml.spds.SyncSource	An implementation of this interface represents a client database. It is used by the SyncManager when it needs to read or store data from and to the local database. SyncSource is purely an interface, which client developers must implement in order to access the most disparate data sources. This componentn
SyncItem	com.funambol.syncml.spds.SyncItem	This class is a container for the items exchanged between the SyncManager and the SyncSource. It is designed to let the SyncManager handle any data in the same way.
SyncListener	com.funambol.util.SyncListener	This class is a generic listener to monitor a synchronization session.

Figure 2.2 represents the class diagram for the data synchronization layer (only the main classes with the relevant details are shown, for a complete description of the classes please refer to the Javadoc).

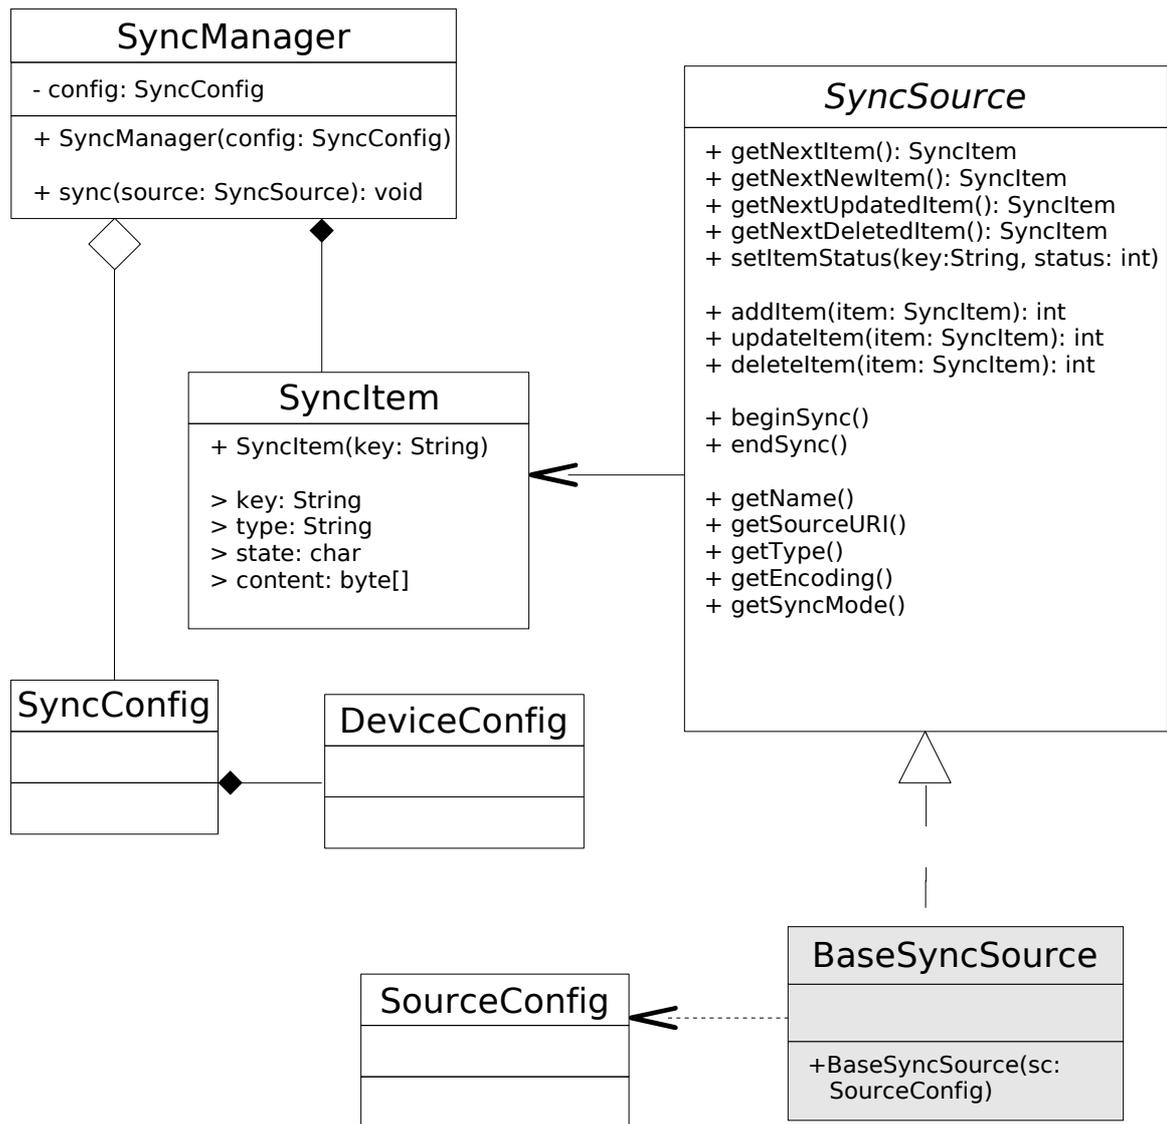


Figure 2.2: SyncManager class diagram

### 2.2.1. SyncManager

SyncManager represents the core of the data synchronization API, and is the contact point between the application and the synchronization engine.

The SyncManager needs some configuration parameters to work, and these parameters are passed during the creation of a new instance of the class using the SyncConfig class. It is client's responsibility to fill or retrieve the config before creating the SyncManager, but the SyncConfig already implements the Serializable interface (see [COMMON]) to make it easier for the client to store and retrieve it to and from the device persistent storage. See below for a description of the SyncConfig class.

A simplified example of SyncManager creation can be:

```

SyncConfig conf = new SyncConfig();
// set something in the conf...
conf.serverUrl = "http://www.somewhere.com/funambol.ds"
SyncManager sm = new SyncManager(conf);
  
```

### Trigger synchronization

To actually start a synchronization, the client application must call the `sync()` method, passing the `SyncSource` to synchronize as a parameter (see the `SyncSource` section for more details on how to create a `SyncSource`):

```
sm.sync(source)
```

When `sync()` is invoked the `SyncManager` start the sync process as described in section 2.1.

This version of the API does not implement large object handling yet (which means the ability to split a big item amongst multiple messages). This feature will be implemented in a future version of the API.

### Detecting, Sending and Committing Client Modifications

Note that the way changes are detected is left to the `SyncSource` implementor. It may be based on flag modification, change log handling, versioning... this is hidden to the `SyncManager` who just asks for modifications calling `getNext***Item()` on a `SyncSource` implementation's instance.

`getNextItem()` is used in the case of a slow sync and retrieve the entire data source content.

`getNext[New/Updated/Deleted]Item()` are used in the case of fast sync and retrieve the items in the data store that are in the corresponding state.

From the `SyncManager` perspective, a `SyncSource` is like a database cursor: the cursor is initialized when `beginSync()` is called and each `getNextItem()` call returns an item or `null` if the cursor reached the end of the result set.

The so retrieved client modifications are embedded in the `SyncML` message and sent to the server. Note that not necessarily all modifications go in one single message. `SyncManager` determines if an item must go in the current message or in the next message based on the number of items in the current message and their size (see the multi message section earlier in this document).

When the `SyncManager` receives a status for a modification command previously sent, it calls the method `setItemStatus()` of the corresponding `SyncSource`. This gives the `SyncSource` the opportunity to commit the change, for example resetting the item modification flag.

### Applying Server Modifications

After the client has done with its modifications the server starts sending its own. These are interpreted by the `SyncManager` which calls the `SyncSource`'s `add/update/deleteItem()` methods. They return a status code as defined by the `SyncML` specifications.

Note that the item parameter passed to `addItem()` on return may have a changed key. This represents the local id (LUID) that will be returned to the server in a `Map` command. See the table in the `SyncSource` Interface paragraph for details on the `addItem()` method.

### Device capabilities handling

`SyncManager` also provides device capabilities. This information combines device specific information with data source related information, therefore properties from many configuration objects will be used to build the `DevInf` command to send to the server. `DevInf` properties are retrieved from the configuration (`SyncConfig` object) passed to the `SyncManager` constructor (see `SyncManager` and `SyncSource` configuration).

If it is the first time `SyncManager` connects to a server, client capabilities are sent in a `Put` command (see the client capabilities handling session).

In the same way, if during initialization the server requests client capabilities sending, `SyncManager` will send them in a `Results` command.

## **2.2.2. The SyncSource Interface**

`SyncSource` is an interface that implementors must implement in order to access specific data sources. The following table gives a list of the main methods with a brief description. Please refer to the Javadoc for a complete description of the interface.

<i>Method</i>	<i>Description</i>
void beginSync(int syncmode)	Called by the SyncManager after the initialization phase, when the sync mode sent by the server is known. The syncmode is normally used by the SyncSource implementation to prepare the lists of items.
void endSync()	Called by the SyncManager before committing the source. The SyncSource implementation can stop the commit phase throwing an exception here.
SyncItem getNextItem()	Returns the next item regardless its modification state, or NULL if the source contains no items.
SyncItem getNextNewItem()	Returns the next new item or NULL if there are no more new items.
SyncItem getNextUpdatedItem()	Returns the next updated item or NULL if there are no more updated items.
SyncItem getNextDeletedItem()	Returns the next deleted item or NULL if there are no more deleted items.
void setItemStatus(String key, int status)	Called when a Status for a modification command previously sent is returned by the server. This gives the opportunity to the SyncSource to know if the modification has succeeded and to reset the related modification flag.
int addItem(SyncItem item)	Adds a new item to the underlying data source. The call returns the status code that SyncManager shall return to the server in the Status command.  If the SyncSource implementation generates a new local ID for the added item, it must set the new key in the Item object, so that the SyncManager will be able to send it back to server in a Map command.
int updateItem(SyncItem item)	Updates an item into the underlying data source. The call returns the status code that SyncManager shall return to the server in the Status command.
int deleteItem(SyncItem item)	Deletes an item from the underlying data source. The call returns the status code that SyncManager shall return to the server in the Status command.
void dataReceived(String date, int size)	Called from the engine when new data arrives from the server. The date is in the format of the HTTP header, if available, or null otherwise. The size is the total size of the message received, including the protocol overhead.
SyncFilter getFilter()	Returns an implementation of the SyncFilter interface, with the filter to be used for this SyncSource. See section 2.5. for more details.
void setFilter(SyncFilter filter)	Sets the filter that should be used in preparation for the synchronization of this SyncSource. See section 2.5 for more details.

### 2.2.3. BaseSyncSource

BaseSyncSource is an abstract class implementing the SyncSource interface. It contains a SourceConfig (see 2.2.6) instance with the source configuration data, and implements the most common methods. An implementor can extend this class implementing only the getXXXItem() and the add/update/deleteItem() methods.

The constructor of BaseSyncSource, and generally of all the SyncSource implementations, requires a SourceConfig parameter (a SyncSource implementation can also use a class derived from SourceConfig). Given a TestSyncSource that extends BaseSyncSource, and a NamedObjectStore instance used to retrieve the config (see [COMMON]), an example of use can be:

```
SourceConfig sc = new SourceConfig();
namedObjectStore.retrieve("source.test", sc);
TestSyncSource tss = new TestSyncSource(sc);
// Perform a sync: the process can change the SourceConfig (the anchors,
// for instance)
sm.sync(tss);
// Save the modified config
namedObjectStore.store("source.test", sc);
```

## 2.2.4. SyncConfig

SyncConfig is a container class that groups the configuration information needed by the SyncManager. The properties of SyncConfig are described below, for a complete description of the fields, please refer to the Javadoc.

<i>Property</i>	<i>Description</i>
syncURL	The syncURL value. If the URL does not start with http:// (or HTTP://) or https:// (or HTTPS://), http:// is prepended to the given string.
username	The username to use for client authentication
password	The password to use for client authentication
userAgent	The user agent string, will be sent in the HTTP header to identify the client on server side. It should be a short description with the client name plus its version.
deviceConfig	The device settings, used to build the DevInf tag.
devInfHash	This is a hash value generated from all properties that are used for the <DevInf> element, plus the syncURL property from SyncConfig. Initial value = "0". TODO: check this on J2ME

Notes:

'userAgent' value should be specified by the client. If the property is left empty, the user agent will have the form: "Profile/<MID profile> Configuration/[CLDC 1.x]".

'devInfHash' property is used to verify if any DevInf element (or syncURL) has changed since last sync; in that case the devInf is sent to the server (see next paragraph and Client Capabilities Handling chapter).

## 2.2.5. DeviceConfig

The DeviceConfig class is a container for properties related to the device and to the client application. Most of DeviceConfig properties are used to generate the <DevInf> element for client capabilities (see Client Capabilities Handling chapter). DeviceConfig properties are described below.

<i>Property</i>	<i>Description</i>
verDTD	Specifies the major and minor version identifier of the Device Information DTD used in the representation of the Device Information. The value MUST be "1.2". This property is mandatory.
man	Specifies the name of the manufacturer of the device. This property is optional.
mod	Specifies the model name or model number of the device. This property is optional.
oem	Specifies the OEM (Original Equipment Manufacturer) of the device. This property is optional.
fwv	Specifies the firmware version of the device. This property is optional.
swv	Specifies the software version of the device. This property is optional.
hvv	Specifies the hardware version of the device. This property is optional.
devID	Specifies the identifier of the source synchronization device. The content information MUST specify a theoretically, globally unique identifier. This property is mandatory.

<i>Property</i>	<i>Description</i>
devType	Specifies the type of the source synchronization device. Type values for this element type can be e.g. "pager", "handheld", "pda", "phone", "smartphone", "server", "workstation". Other values can also be specified. This property is mandatory.
dsV	Specifies the implemented DS version. This property is optional.
utc	Boolean. Specifies that the device supports UTC based time. If utc = TRUE, the server SHOULD send time in UTC format, else MUST send in local time. Client MAY send time in local or UTC format. Default value is TRUE.
loSupport	Boolean. Specifies that the device supports handling of large objects. Default value is FALSE.
nocSupport	Boolean. Specifies that the device supports number of changes. Default value is FALSE.
maxMsgSize	The maximum message size (byte) accepted for XML messages received from server. The same size is used for messages sent to the server.
maxObjSize	Specifies the maximum object size allowed by the device. Default value is 0 (no maxObjSize set).

Notes:

If there is no firmware/software/hardware version of the device available (fwv, swv, hwv), then their content information can also be a date, for example, 19980114 or 19990714T133000Z. Only hours, minutes and second MUST be specified in the time component.

## 2.2.6. SourceConfig

The class SourceConfig is a container for the configuration parameter of a SyncSource. It is used by BaseSyncSource. The most important SourceConfig properties are described in the table below.

<i>Property</i>	<i>Description</i>
version	A version number to handle software upgrades
encoding	Specifies how the content of an item should be encoded. The form of this parameter is a semi-colon separated list of formats that must be applied in sequence from the leftmost to the rightmost. For example, if format is "des;b64", when the item will be output in the message, the content must be first transformed with the "des" encoder and then with the Base64 encoder.
syncMode	This is the current sync mode used by the client. The parameter MUST be one of the sync modes specified in property "syncModes".
type	The mime type of the items sent by this source.
remoteUri	The remote URI of this source on the SyncML server.
lastAnchor	Long value that specifies the last timestamp for this source.
nextAnchor	Long value that specifies the last timestamp for this source.
ctCap	Specifies the content type capabilities for this SyncSource. This information describes which properties and which values the client can store. <b>Not implemented yet.</b>

## 2.3. Synchronization Events Notification

The synchronization process can be monitored by installing a proper listener that can be associated to a SyncSource. The model for monitoring the synchronization is different in J2ME (compared to J2SE [API-J2SE] and C++ API [API-CPP]).

A general SyncListener interface can be implemented or a utility BasicSyncListener can be extended to implement custom listener.

From a synchronization listener point of view, a synchronization is a sequence of events that happen in a certain order.

There are three different phases (as described in the SyncML synchronization process).

1. Connection set up
2. Modification exchange
3. Mapping

The phases are encapsulated into a sync session which starts with a startSession and is terminated by an endSession. Each phase is mandatory and the order cannot change.

The Modification exchange phase is the most interesting from the point of view of the SyncListener, because it can be further split into more sub events. The modifications are exchanged with a sequence of send/receive sessions. Each time modifications are about to be sent to the server a startSending event is generated. Then we can have an arbitrary sequence of itemAddSent, itemReplaceSent and itemDeleteSent. At the end of the sending step a endSending event is generated. Similarly for the receiving phase.

The sequence of events can be described by the following grammar:

S -> startSession C U M endSession

C -> startConnecting endConnecting

U -> startSyncing (S R)+ endSyncing

M -> startMapping endMapping

S -> startSending (itemAddSent | itemUpdateSent | itemDeleteSent) endSending

R -> startReceiving (itemReceived | itemUpdated | itemDeleted | dataReceived)\*  
endReceiving

Unlike in other APIs, an event is not described by an EventInfo. This is just a pragmatic choice as today such an object would not contain useful information but it would require extra memory and CPU cycles. In the future we may need and introduce such an object to describe events.

The SyncManager notifies both the SyncSource and the SyncListener for some events (such as the server provides new items), but in some case it only notifies the SyncListener (for example on the session start and end). The two objects have a different vision of the synchronization process. The listener is notified about everything that happens (with adequate information that describe the events), while the SyncSource is only notified on events that require an action on its side.

Providing a SyncListener is a possibility, but not a requirement. If a simple client is not interested at observing the synchronization, it can avoid registering any listener.

The class diagram of Figure 2.3 shows the classes used for event notification.

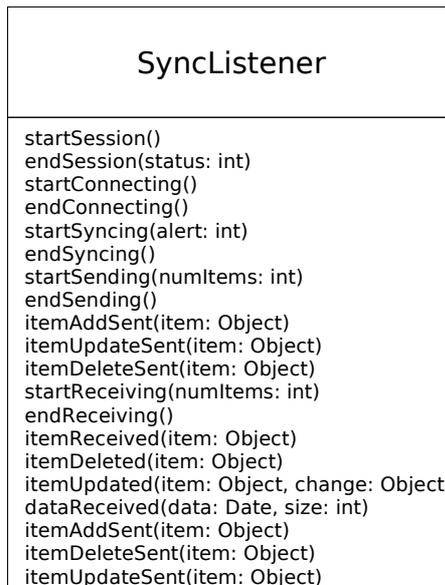


Figure 2.3: Listeners class diagram

## 2.4. Client Capabilities Handling

The SyncML specification ([OMA-DS], § 6.7) state:

*"The sync client MUST send its device information to the server when the first synchronization is done with a server or when the static device information has been updated in the client. The client MUST also be able to transmit its device information if it is asked by the server. The client SHOULD also support the receiving of the server device information".*

Therefore, we have two scenarios:

1. The client connects to a new server;
2. The server issues a Get command requesting for client capabilities.

In both cases the client shall send its capabilities. Plus, if the client has already sent its capabilities to a server in a previous synchronization, the DevInf element is not sent anymore.

For case 1, the SyncManager determines if it has to include client capabilities calculating the value of the devInfHash: if its value is changed it means that some DevInf parameter is changed, or the syncUrl property has been modified.

Note: the syncUrl parameter has been added to the generation of devInfHash, to ensure the client sends device informations when connecting to a different server.

However, a client application can decide to reset devInfHash when any other relevant configuration settings is changed.

In case 2, the SyncManager determines if it has to send the client capabilities inspecting the messages received by the server during initialization. If any contains a Get command, in the reply the SyncManager will send a Results command containing the DevInf build from the various configuration objects.

## 2.5. Filtering

Filtering is a new feature, introduced in the SyncML DS 1.2 specification, that allows a SyncML client to synchronize a database restricting the set of data returned by the server to the items that fall into a given filter. See [OMA-DS], paragraph 5.13, for details.

For the J2ME platform, a simplified interface to add filtering capabilities is defined, so that a client can send a filter to the server, if needed, without overhead for a client that does not need it.

The sync engine uses the SyncFilter interface, that provide a method to format an XML string representing the SyncML filter. The SyncSource interface provide a method getFilter() to obtain the filter defined for that SyncSource, if defined.

If a client need to filter data for a certain SyncSource, it has to implement the getFilter() method of the SyncSource, returning the needed filer. The class BaseSyncSource already implement the interface: a client extending the class has to set the desired filter calling setFilter() before starting the sync.

### The SyncFilter interface

<i>Method</i>	<i>Description</i>
toSyncML()	Return a string containing the SyncML filter to apply for the session.

## 2.6. Large Objects Handling

The SyncML specification ([OMA-DS], § 6.7) state:

*"While synchronizing, object reception can be limited by two factors: the maximum message size the target device can receive (declared in <MaxMsgSize> tag), and the maximum object size the target device can receive (declared in <MaxObjSize> tag).*

*This feature provides a means to synchronize an object whose size exceeds that which can be transmitted within one message (e.g. the maximum message size – declared in <MaxMsgSize> element – that the target device can receive). This is achieved by splitting the object into chunks that will each fit within one message and by sending them contiguously. The first chunk of data is sent with the overall size of the object and a <MoreData/> signaling that more chunks will be sent. Every subsequent chunk is sent with a <MoreData/> tag, except from the last one: the final chunk is sent with no <MoreData/> tag. The target device, having received the final chunk, has to re-construct the object and consequently acts as it had received it in one piece (e.g. apply the requested command). The appropriate status MUST then be sent to the originator. A command on a chunked object MUST implicitly be treated as atomic, i.e. the recipient can only commit the object once all chunks have been successfully received and reassembled."*

In order to implement the support of Large Objects the SyncML API provide a simple mechanism as reported in the following picture.

As stated above, the Large Objects are splitted in chunks during a synchronization.

The SyncManager is responsible to get the chunks and for each chunk creates a SyncItem. The single chunk is set in the SyncItem internal data structure.

Each SyncItem created during the sync of a large object will be discarded as soon as the SyncSource consume its data.

The SyncSource is responsible to manage the chunks, reassembling the data of the different sync items and is in charge also to maintain the entire large object in memory or serialize it depending on the needs of the running platform.

For java platforms with restrictions in terms of runtime memory, the SyncSource should store the large object in local storage as soon as the chunk is set in the item. In this way the whole large object won't be loaded in the heap space, but only the single chunk. This allows to implement the support of Large Objects also on java platforms with small heap space.

The SyncSource is able to check if a new chunk is coming or not checking the *moreData* property of the SyncItem. This property is updated by the SyncManager checking the <MoreData/> SyncML tag.

After the SyncManager has retrieved all the chunks, the SyncItem moreData property is set to false and the SyncSource can finalize the Large Object.

Only after the end of the large object synchronization the SyncSource is able to create the large object client representation (i.e. File, Contact Photo, Message Attachment) converting the reassembled data chunks in the specific format.

The following picture shows the sequence diagram for the large objects client handling in case of new Item from Server.

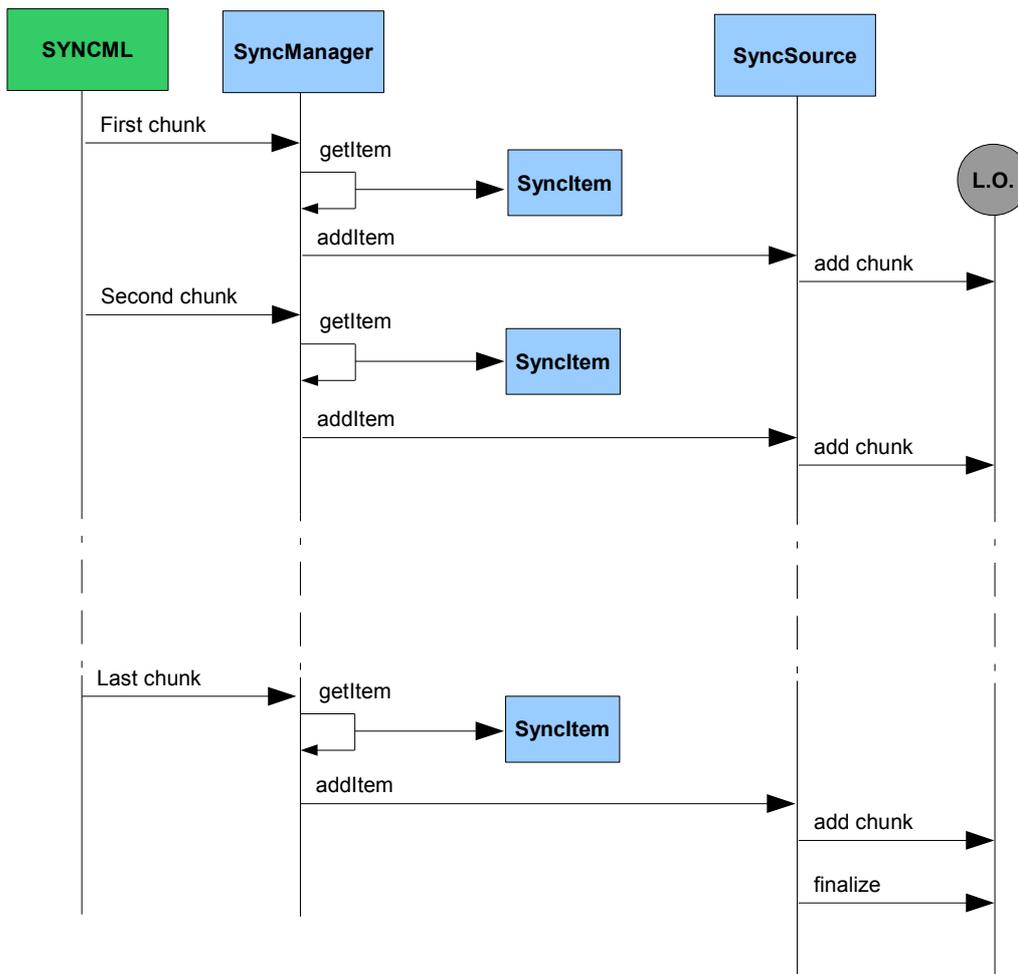


Figure 2.4: Large Objects Client Handling - New Item from Server

A similar approach is used in case of item sent from the client to the server: the SyncSource splits the large object in chunk, getting them from the storage and the SyncManager sends a SyncItem for each chunk.

There are cases where delegating the chunks management to the engine can be a better choice. In fact the large objects support will be hid to top layers (client) and the SyncSource implementation can just use a SyncItem with no knowledge of large objects.

In this case the SyncManager handles all the chunks, assembling them in memory and only at the end creates the SyncItem containing the whole data.

In order to handle this case the SyncSource provides a method that lets the SyncManager to know who is in charge of re-assembling the chunks.

The method `SyncSource.isReassembling()` will be the “switch” between the two different approach: in the first one (true) the SyncSource manages the chunks, in the second one (false) the SyncManager handles the chunks.

The disadvantage of this second approach is that the large object is alive in memory and it cannot be applied to platforms with low runtime memory like low-end MIDP handset.

## 3. References

- [OMA-DS] SyncML Data Synchronization Protocol, version 1.2, Open Mobile Alliance
- [COMMON] Funambol J2ME Common API Design Document, version 1.0, Funambol Inc.
- [API-J2SE] Funambol Java API J2SE Developer Guide, version 1.0, Funambol Inc.
- [API-CPP] Funambol 3.0 Client API C++ Design Document, version 1.0, Funambol Inc.