



# APACHE UNOMI 1.2.X - DOCUMENTATION

Apache Software Foundation

# TABLE OF CONTENTS

1. Concepts	1
1.1. Concepts	1
1.1.1. Items and types	1
1.1.2. Events	2
1.1.3. Profiles	3
1.1.4. Sessions	4
1.2. Extending Unomi via plugins	4
1.2.1. Types vs. instances	4
1.2.2. Plugin structure	4
1.2.3. Extension points	5
1.2.4. Other Unomi entities	6
2. Quick start	7
2.1. Building	7
2.1.1. Initial Setup	7
2.1.2. Building	7
2.1.3. Installing an Elasticsearch server	7
2.1.4. Deploying the generated binary package	8
2.1.5. Deploying into an existing Karaf server	9
2.1.6. JDK Selection on Mac OS X	9
2.1.7. Running the integration tests	10
2.1.8. Running the performance tests	11
2.1.9. Testing with an example page	11
2.1.10. Integrating onto a page	11
2.2. Getting started with Unomi	11
2.2.1. Prerequisites	11
2.2.2. Running Unomi	12
2.3. Configuration	14
2.3.1. Changing the default configuration	14
2.3.2. Secured events configuration	14
2.3.3. Installing the MaxMind GeoIPLite2 IP lookup database	15
2.3.4. Installing Geonames database	15
2.3.5. REST API Security	15
2.3.6. Automatic profile merging	16
2.3.7. Securing a production environment	16
2.3.8. Integrating with an Apache HTTP web server	18
2.3.9. Changing the default tracking location	19
2.3.10. Apache Karaf SSH Console	20
3. Sample	20
3.1. Samples	20
3.2. Login sample	20
3.2.1. Warning !	20
3.2.2. Installing the sample	21
3.3. Twitter sample	21
3.3.1. Overview	21
3.3.2. Interacting with the context server	22
3.3.3. Retrieving context information from Unomi using the context servlet	23
3.4. Example	24
3.4.1. HTML page	24
3.4.2. Javascript	24

3.5. Conclusion	32
3.6. Annex	32
3.7. Weather update sample	32
4. Connectors	32
4.1. Connectors	32
4.1.1. Call for contributors	33
4.2. Apache Unomi Salesforce Connector	33
4.2.1. Getting started	33
4.2.2. Troubleshooting	34
4.2.3. Upgrading the Salesforce connector	35
4.2.4. Using the Salesforce Workbench for testing REST API	36
4.2.5. Setting up Streaming Push queries	36
4.2.6. Executing the unit tests	36
5. Cluster setup	37
5.1. Cluster setup	37
5.1.1. 2 nodes configuration	37
5.1.2. 3 nodes configuration	37
6. Custom extensions	38
6.1. Custom extensions	38
6.1.1. Creating an extension	38
6.1.2. Predefined segments	40
6.1.3. Predefined rules	40
6.1.4. Predefined properties	41
6.1.5. Predefined child conditions	42
6.1.6. Predefined personas	42
6.1.7. Custom actions	43
6.1.8. Custom conditions	45

# 1. CONCEPTS

## 1.1. CONCEPTS

Apache Unomi gathers information about users actions, information that is processed and stored by Unomi services. The collected information can then be used to personalize content, derive insights on user behavior, categorize the user profiles into segments along user-definable dimensions or acted upon by algorithms.

### 1.1.1. ITEMS AND TYPES

Unomi structures the information it collects using the concept of [Item](#) which provides the base information (an identifier and a type) the context server needs to process and store the data. Items are persisted according to their type (structure) and identifier (identity). This base structure can be extended, if needed, using properties in the form of key-value pairs.

These properties are further defined by the [Item](#)'s type definition which explicits the [Item](#)'s structure and semantics. By defining new types, users specify which properties (including the type of values they accept) are available to items of that specific type.

Unomi defines default value types: [date](#), [email](#), [integer](#) and [string](#), all pretty self-explanatory. While you can think of these value types as "primitive" types, it is possible to extend Unomi by providing additional value types.

Additionally, most items are also associated to a scope, which is a concept that Unomi uses to group together related items. A given scope is represented in Unomi by a simple string identifier and usually represents an application or set of applications from which Unomi gathers data, depending on the desired analysis granularity. In the context of web sites, a scope could, for example, represent a site or family of related sites being analyzed. Scopes allow clients accessing the context server to filter data to only see relevant data.

*Base [Item](#) structure:*

```
{
  "itemType": <type of the item>,
  "scope": <scope>,
  "itemId": <item identifier>,
  "properties": <optional properties>
}
```

Some types can be dynamically defined at runtime by calling to the REST API while other extensions are done via Unomi plugins. Part of extending Unomi, therefore, is a matter of defining new types and specifying which kind of Unomi entity (e.g. profiles) they can be affected to. For example, the following

JSON document can be passed to Unomi to declare a new property type identified (and named) `tweetNb`, tagged with the `social` tag, targeting profiles and using the `integer` value type.

*Example JSON type definition:*

```
{
  "itemId": "tweetNb",
  "itemType": "propertyType",
  "metadata": {
    "id": "tweetNb",
    "name": "tweetNb"
  },
  "tags": ["social"],
  "target": "profiles",
  "type": "integer"
}
```

Unomi defines a built-in scope (called `systemscope`) that clients can use to share data across scopes.

### 1.1.2. EVENTS

Users' actions are conveyed from clients to the context server using events. Of course, the required information depends on what is collected and users' interactions with the observed systems but events minimally provide a type, a scope and source and target items. Additionally, events are timestamped. Conceptually, an event can be seen as a sentence, the event's type being the verb, the source the subject and the target the object.

*Event structure:*

```
{
  "eventType": <type of the event>,
  "scope": <scope of the event>,
  "source": <Item>,
  "target": <Item>,
  "properties": <optional properties>
}
```

Source and target can be any Unomi item but are not limited to them. In particular, as long as they can be described using properties and Unomi's type mechanism and can be processed either natively or via extension plugins, source and target can represent just about anything. Events can also be triggered as part of Unomi's internal processes for example when a rule is triggered.

Events are sent to Unomi from client applications using the JSON format and a typical page view event from a web site could look something like the following:

*Example page view event:*

```

{
  "eventType": "view",
  "scope": "ACMESPAC",
  "source": {
    "itemType": "site",
    "scope": "ACMESPAC",
    "itemId": "c4761bbf-d85d-432b-8a94-37e866410375"
  },
  "target": {
    "itemType": "page",
    "scope": "ACMESPAC",
    "itemId": "b6acc7b3-6b9d-4a9f-af98-54800ec13a71",
    "properties": {
      "pageInfo": {
        "pageID": "b6acc7b3-6b9d-4a9f-af98-54800ec13a71",
        "pageName": "Home",
        "pagePath": "/sites/ACMESPAC/home",
        "destinationURL": "http://localhost:8080/sites/ACMESPAC/home.html",
        "referringURL": "http://localhost:8080/",
        "language": "en"
      },
      "category": {},
      "attributes": {}
    }
  }
}

```

### 1.1.3. PROFILES

By processing events, Unomi progressively builds a picture of who the user is and how they behave. This knowledge is embedded in [Profile](#) object. A profile is an [Item](#) with any number of properties and optional segments and scores. Unomi provides default properties to cover common data (name, last name, age, email, etc.) as well as default segments to categorize users. Unomi users are, however, free and even encouraged to create additional properties and segments to better suit their needs.

Contrary to other Unomi items, profiles are not part of a scope since we want to be able to track the associated user across applications. For this reason, data collected for a given profile in a specific scope is still available to any scoped item that accesses the profile information.

It is interesting to note that there is not necessarily a one to one mapping between users and profiles as users can be captured across applications and different observation contexts. As identifying information might not be available in all contexts in which data is collected, resolving profiles to a single physical user can become complex because physical users are not observed directly. Rather, their portrait is progressively patched together and made clearer as Unomi captures more and more traces of their actions. Unomi will merge related profiles as soon as collected data permits positive association between distinct profiles, usually as a result of the user performing some identifying action in a context where the user hadn't already been positively identified.

## 1.1.4. SESSIONS

A session represents a time-bounded interaction between a user (via their associated profile) and a Unomi-enabled application. A session represents the sequence of actions the user performed during its duration. For this reason, events are associated with the session during which they occurred. In the context of web applications, sessions are usually linked to HTTP sessions.

## 1.2. EXTENDING UNOMI VIA PLUGINS

Unomi is architected so that users can provide extensions in the form of plugins.

### 1.2.1. TYPES VS. INSTANCES

Several extension points in Unomi rely on the concept of type: the extension defines a prototype for what the actual items will be once parameterized with values known only at runtime. This is similar to the concept of classes in object-oriented programming: types define classes, providing the expected structure and which fields are expected to be provided at runtime, that are then instantiated when needed with actual values.

### 1.2.2. PLUGIN STRUCTURE

Being built on top of Apache Karaf, Unomi leverages OSGi to support plugins. A Unomi plugin is, thus, an OSGi bundle specifying some specific metadata to tell Unomi the kind of entities it provides. A plugin can provide the following entities to extend Unomi, each with its associated definition (as a JSON file), located in a specific spot within the [META-INF/cxs/](#) directory of the bundle JAR file:

Entity	Location in <a href="#">cxs</a> directory
ActionType	actions
ConditionType	conditions
Persona	personas
PropertyMergeStrategyType	mergers
PropertyType	properties then profiles or sessions subdirectory then <a href="#">&lt;category name&gt;</a> directory
Rule	rules
Scoring	scorings
Segment	segments
Tag	tags then <a href="#">&lt;category name&gt;</a> directory
ValueType	values

[Blueprint](#) is used to declare what the plugin provides and inject any required dependency. The Blueprint file is located, as usual, at [OSGI-INF/blueprint/blueprint.xml](#) in the bundle JAR file.

The plugin otherwise follows a regular maven project layout and should depend on the Unomi API maven artifact:

```
<dependency>
  <groupId>org.apache.unomi</groupId>
  <artifactId>unomi-api</artifactId>
  <version>...</version>
</dependency>
```

Some plugins consists only of JSON definitions that are used to instantiate the appropriate structures at runtime while some more involved plugins provide code that extends Unomi in deeper ways.

In both cases, plugins can provide more that one type of extension. For example, a plugin could provide both `ActionType`s and `ConditionType`s.

### 1.2.3. EXTENSION POINTS

#### ACTIONTYPE

`ActionType`s` define new actions that can be used as consequences of Rules being triggered. When a rule triggers, it creates new actions based on the event data and the rule internal processes, providing values for parameters defined in the associated `ActionType`. Example actions include: “Set user property x to value y” or “Send a message to service x”.

#### CONDITIONTYPE

`ConditionType`s define new conditions that can be applied to items (for example to decide whether a rule needs to be triggered or if a profile is considered as taking part in a campaign) or to perform queries against the stored Unomi data. They may be implemented in Java when attempting to define a particularly complex test or one that can better be optimized by coding it. They may also be defined as combination of other conditions. A simple condition could be: “User is male”, while a more generic condition with parameters may test whether a given property has a specific value: “User property x has value y”.

#### PERSONA

A persona is a "virtual" profile used to represent categories of profiles, and may also be used to test how a personalized experience would look like using this virtual profile. A persona can define predefined properties and sessions. Persona definition make it possible to “emulate” a certain type of profile, e.g : US visitor, non-US visitor, etc.

#### PROPERTYMERGESTRATEGYTYPE

A strategy to resolve how to merge properties when merging profile together.

#### PROPERTYTYPE

Definition for a profile or session property, specifying how possible values are constrained, if the value is multi-valued (a vector of values as opposed to a scalar value). `PropertyType`s can also be categorized using tags or file system structure, using sub-directories to organize definition files.



## RULE

`Rule`s are conditional sets of actions to be executed in response to incoming events. Triggering of rules is guarded by a condition: the rule is only triggered if the associated condition is satisfied. That condition can test the event itself, but also the profile or the session. Once a rule triggers, a list of actions can be performed as consequences. Also, when rules trigger, a specific event is raised so that other parts of Unomi can react accordingly.

## SCORING

`Scoring`s are set of conditions associated with a value to assign to profiles when matching so that the associated users can be scored along that dimension. Each scoring element is evaluated and matching profiles' scores are incremented with the associated value.

## SEGMENTS

`Segment`s represent dynamically evaluated groups of similar profiles in order to categorize the associated users. To be considered part of a given segment, users must satisfies the segment's condition. If they match, users are automatically added to the segment. Similarly, if at any given point during, they cease to satisfy the segment's condition, they are automatically removed from it.

## TAG

`Tag`s are simple labels that are used to classify all other objects inside Unomi. Tags can define sub-tags.

## VALUETYPE

Definition for values that can be assigned to properties ("primitive" types).

## 1.2.4. OTHER UNOMI ENTITIES

### USERLIST

User list are simple static lists of users. The associated profile stores the lists it belongs to in a specific property.

### GOAL

Goals represent tracked activities / actions that can be accomplished by site (or more precisely scope) visitors. These are tracked in general because they relate to specific business objectives or are relevant to measure site/scope performance.

Goals can be defined at the scope level or in the context of a particular [Campaign](#). Either types of goals behave exactly the same way with the exception of two notable differences: - duration: scope-level goals are considered until removed while campaign-level goals are only considered for the campaign duration - audience filtering: any visitor is considered for scope-level goals while campaign-level goals only consider visitors who match the campaign's conditions

## CAMPAIGN

A goal-oriented, time-limited marketing operation that needs to be evaluated for return on investment performance by tracking the ratio of visits to conversions.

# 2. QUICK START

## 2.1. BUILDING

### 2.1.1. INITIAL SETUP

- 1) Install J2SE 8.0 SDK (or later), which can be downloaded from <http://www.oracle.com/technetwork/java/javase/downloads/index.html>
- 2) Make sure that your JAVA\_HOME environment variable is set to the newly installed JDK location, and that your PATH includes %JAVA\_HOME%\bin (windows) or \$JAVA\_HOME/bin (unix).
- 3) Install Maven 3.0.3 (or later), which can be downloaded from <http://maven.apache.org/download.html>. Make sure that your PATH includes the MVN\_HOME/bin directory.

### 2.1.2. BUILDING

- 1) Change to the top level directory of Apache Unomi source distribution. 2) Run

```
$> mvn clean install
```

This will compile Apache Unomi and run all of the tests in the Apache Unomi source distribution. Alternatively, you can run

```
$> mvn -P !integration-tests,!performance-tests clean install
```

This will compile Apache Unomi without running the tests and takes less time to build.

- 3) The distributions will be available under "package/target" directory.

### 2.1.3. INSTALLING AN ELASTICSEARCH SERVER

Starting with version 1.2, Apache Unomi no longer embeds an Elasticsearch server as this is no longer supported by the developers of Elasticsearch. Therefore you will need to install a standalone Elasticsearch using the following steps:

.

Download an Elasticsearch version. Here's the version you will need depending on your version of

Apache Unomi.

Apache Unomi <= 1.2 : <https://www.elastic.co/downloads/past-releases/elasticsearch-5-1-2> Apache Unomi  
>= 1.3 : <https://www.elastic.co/downloads/past-releases/elasticsearch-5-6-3>

Uncompress the downloaded package into a directory

In the config/elasticsearch.yml file, uncomment and modify the following line :

```
cluster.name: contextElasticSearch
```

Launch the server using

```
bin/elasticsearch (Mac, Linux)  
bin\elasticsearch.bat (Windows)
```

Check that the ElasticSearch is up and running by accessing the following URL :

<http://localhost:9200>

## 2.1.4. DEPLOYING THE GENERATED BINARY PACKAGE

The "package" sub-project generates a pre-configured Apache Karaf installation that is the simplest way to get started. Simply uncompress the package/target/unomi-VERSION.tar.gz (for Linux or Mac OS X) or package/target/unomi-VERSION.zip (for Windows) archive into the directory of your choice.

You can then start the server simply by using the command on UNIX/Linux/MacOS X :

```
./bin/karaf
```

or on Windows shell :

```
bin\karaf.bat
```

## 2.1.5. DEPLOYING INTO AN EXISTING KARAF SERVER

This is only needed if you didn't use the generated package. Also, this is the preferred way to install a development environment if you intend to re-deploy the context server KAR iteratively.

Additional requirements: \* Apache Karaf 3.x, <http://karaf.apache.org>

Before deploying, make sure that you have Apache Karaf properly installed. You will also have to increase the default maximum memory size and perm gen size by adjusting the following environment values in the bin/setenv(.bat) files (at the end of the file):

```
MY_DIRNAME= `dirname $0`  
MY_KARAF_HOME=`cd "$MY_DIRNAME/.."; pwd`  
export JAVA_MAX_MEM=3G  
export JAVA_MAX_PERM_MEM=384M
```

Install the WAR support, CXF and Karaf Cellar into Karaf by doing the following in the Karaf command line:

```
feature:repo-add cxf 3.0.2  
feature:repo-add cellar 3.0.3  
feature:repo-add mvn:org.apache.unomi/unomi-kar/VERSION/xml/features  
feature:install unomi-kar
```

Create a new \$MY\_KARAF\_HOME/etc/org.apache.cxf.osgi.cfg file and put the following property inside :

```
org.apache.cxf.servlet.context=/cxs
```

If all went smoothly, you should be able to access the context script here : <http://localhost:8181/cxs/cluster> . You should be able to login with karaf / karaf and see basic server information. If not something went wrong during the install.

## 2.1.6. JDK SELECTION ON MAC OS X

You might need to select the JDK to run the tests in the itests subproject. In order to do so you can list the installed JDKs with the following command :

```
/usr/libexec/java_home -V
```

which will output something like this :

```
Matching Java Virtual Machines (7):
 1.7.0_51, x86_64: "Java SE 7" /Library/Java/JavaVirtualMachines/jdk1.7.0_51.jdk/Contents/Home
 1.7.0_45, x86_64: "Java SE 7" /Library/Java/JavaVirtualMachines/jdk1.7.0_45.jdk/Contents/Home
 1.7.0_25, x86_64: "Java SE 7" /Library/Java/JavaVirtualMachines/jdk1.7.0_25.jdk/Contents/Home
 1.6.0_65-b14-462, x86_64: "Java SE 6" /Library/Java/JavaVirtualMachines/1.6.0_65-b14-
462.jdk/Contents/Home
 1.6.0_65-b14-462, i386: "Java SE 6" /Library/Java/JavaVirtualMachines/1.6.0_65-b14-
462.jdk/Contents/Home
 1.6.0_65-b14-462, x86_64: "Java SE 6"
/System/Library/Java/JavaVirtualMachines/1.6.0.jdk/Contents/Home
 1.6.0_65-b14-462, i386: "Java SE 6"
/System/Library/Java/JavaVirtualMachines/1.6.0.jdk/Contents/Home
```

You can then select the one you want using :

```
export JAVA_HOME=`/usr/libexec/java_home -v 1.7.0_51`
```

and then check that it was correctly referenced using:

```
java -version
```

which should give you a result such as this:

```
java version "1.7.0_51"
Java(TM) SE Runtime Environment (build 1.7.0_51-b13)
Java HotSpot(TM) 64-Bit Server VM (build 24.51-b03, mixed mode)
```

## 2.1.7. RUNNING THE INTEGRATION TESTS

The integration tests are not executed by default to make build time minimal, but it is recommended to run the integration tests at least once before using the server to make sure that everything is ok in the build. Another way to use these tests is to run them from a continuous integration server such as Jenkins, Apache Gump, Atlassian Bamboo or others.

Note : the integration tests require a JDK 7 or more recent !

To run the tests simply activate the following profile :

```
mvn -P integration-tests clean install
```

## 2.1.8. RUNNING THE PERFORMANCE TESTS

Performance tests are based on Gatling. You need to have a running context server or cluster of servers before executing the tests.

Test parameters are editable in the `performance-tests/src/test/scala/unomi/Parameters.scala` file. `baseUrl` should contain the URLs of all your cluster nodes

Run the test by using the `gatling.conf` file in `performance-tests/src/test/resources` :

```
export GATLING_CONF=<path>/performance-tests/src/test/resources
gatling.sh
```

Reports are generated in `performance-tests/target/results`.

## 2.1.9. TESTING WITH AN EXAMPLE PAGE

A default test page is provided at the following URL:

```
http://localhost:8181/index.html
```

This test page will trigger the loading of the `/context.js` script, which will try to retrieve the user context or create a new one if it doesn't exist yet. It also contains an experimental integration with Facebook Login, but it doesn't yet save the context back to the context server.

## 2.1.10. INTEGRATING ONTO A PAGE

Simply reference the context script in your HTML as in the following example:

```
<script type="text/javascript">
  (function(){ var u=("https:" == document.location.protocol) ? "https://localhost:8181/" :
"http://localhost:8181/";
  var d=document, g=d.createElement('script'), s=d.getElementsByTagName('script')[0];
g.type='text/javascript'; g.defer=true; g.async=true; g.src=u+'context.js';
  s.parentNode.insertBefore(g,s); })();
</script>
```

## 2.2. GETTING STARTED WITH UNOMI

We will first get you up and running with an example. We will then lift the corner of the cover somewhat and explain in greater details what just happened.

### 2.2.1. PREREQUISITES

This document assumes that you are already familiar with Unomi's [concepts](#). On the technical side, we also assume working knowledge of [git](#) to be able to retrieve the code for Unomi and the example.

Additionally, you will require a working Java 7 or above install. Refer to <http://www.oracle.com/technetwork/java/javase/> for details on how to download and install Java SE 7 or greater.

## 2.2.2. RUNNING UNOMI

### BUILDING UNOMI

1. Get the code: `git clone https://git-wip-us.apache.org/repos/asf/incubator-unomi.git`
2. Build and install according to the [instructions](#) and install Unomi.

### START UNOMI

Start Unomi according to the [instructions](#). Once you have Karaf running, you should wait until you see the following messages on the Karaf console:

```
Initializing user list service endpoint...
Initializing geonames service endpoint...
Initializing segment service endpoint...
Initializing scoring service endpoint...
Initializing campaigns service endpoint...
Initializing rule service endpoint...
Initializing profile service endpoint...
Initializing cluster service endpoint...
```

This indicates that all the Unomi services are started and ready to react to requests. You can then open a browser and go to <http://localhost:8181/cxs> to see the list of available RESTful services or retrieve an initial context at <http://localhost:8181/context.json> (which isn't very useful at this point).

### REQUEST EXAMPLES

#### RETRIEVING YOUR FIRST CONTEXT

You can retrieve a context using curl like this :

```
curl http://localhost:8181/context.js?sessionId=1234
```

This will retrieve a JavaScript script that contains a `cxs` object that contains the context with the current user profile, segments, scores as well as functions that makes it easier to perform further requests (such as collecting events using the `cxs.collectEvents()` function).

#### RETRIEVING A CONTEXT AS A JSON OBJECT.

If you prefer to retrieve a pure JSON object, you can simply use a request formed like this:

```
curl http://localhost:8181/context.json?sessionId=1234
```

## ACCESSING PROFILE PROPERTIES IN A CONTEXT

By default, in order to optimize the amount of data sent over the network, Apache Unomi will not send the content of the profile or session properties. If you need this data, you must send a JSON object to configure the resulting output of the context.js(on) servlet.

Here is an example that will retrieve all the session and profile properties.

```
curl -H "Content-Type: application/json" -X POST -d
'{"source":{"itemId":"homepage","itemType":"page","scope":"example"},"requiredProfileProperties":["*"],"requiredSessionProperties":["*"],"requireSegments":true}'
http://localhost:8181/context.json?sessionId=1234
```

The `requiredProfileProperties` and `requiredSessionProperties` are properties that take an array of property names that should be retrieved. In this case we use the wildcard character `*` to say we want to retrieve all the available properties. The structure of the JSON object that you should send is a JSON-serialized version of the `ContextRequest` Java class.

## SENDING EVENTS USING THE CONTEXT SERVLET

At the same time as you are retrieving the context, you can also directly send events in the `ContextRequest` object as illustrated in the following example:

```
curl -H "Content-Type: application/json" -X POST -d
'{"source":{"itemId":"homepage","itemType":"page","scope":"example"},"events":[{"eventType":"view","scope":"example","source":{"itemType":"site","scope":"example","itemId":"mysite"},"target":{"itemType":"page","scope":"example","itemId":"homepage"},"properties":{"pageInfo":{"referringURL":""}}}]}' http://localhost:8181/context.json?sessionId=1234
```

Upon received events, Apache Unomi will execute all the rules that match the current context, and return an updated context. This way of sending events is usually used upon first loading of a page. If you want to send events after the page has finished loading you could either do a second call and get an updating context, or if you don't need the context and want to send events in a network optimal way you can use the `eventcollector` servlet (see below).

## SENDING EVENTS USING THE EVENTCOLLECTOR SERVLET

If you only need to send events without retrieving a context, you should use the `eventcollector` servlet that is optimized respond quickly and minimize network traffic. Here is an example of using this servlet:

```
curl -H "Content-Type: application/json" -X POST -d '{"events":[{"eventType":"view","scope":"example","source":{"itemType":"site","scope":"example","itemId":"mysite"},"target":{"itemType":"page","scope":"example","itemId":"homepage"},"properties":{"pageInfo":{"referringURL":""}}}]}' http://localhost:8181/eventcollector?sessionId=1234
```

Note that the `eventcollector` executes the rules but does not return a context. It is generally used after a page is loaded to send additional events.



## WHERE TO GO FROM HERE

- Read the [Twitter sample](#) documentation that contains a detailed example of how to integrate with Apache Unomi.

## 2.3. CONFIGURATION

### 2.3.1. CHANGING THE DEFAULT CONFIGURATION

If you want to change the default configuration, you can perform any modification you want in the `$MY_KARAF_HOME/etc` directory.

The context server configuration is kept in the `$MY_KARAF_HOME/etc/org.apache.unomi.cluster.cfg`. It defines the addresses where it can be found :

```
contextserver.publicAddress=https://localhost:9443
contextserver.internalAddress=http://127.0.0.1:8181
```

If you need to specify an Elasticsearch cluster name, or a host and port that are different than the default, it is recommended to do this BEFORE you start the server for the first time, or you will loose all the data you have stored previously.

To change these settings, you will need to modify a file called

```
$MY_KARAF_HOME/etc/org.apache.unomi.persistence.elasticsearch.cfg
```

with the following contents:

```
cluster.name=contextElasticSearch
# The elasticSearchAddresses may be a comma separated list of host names and ports such as
# hostA:9300,hostB:9300
# Note: the port number must be repeated for each host.
elasticSearchAddresses=localhost:9300
index.name=context
```

### 2.3.2. SECURED EVENTS CONFIGURATION

If you need to secure some events, that will be sent only by a trusted third party server, you can update the file :

```
$MY_KARAF_HOME/etc/org.apache.unomi.thirdparty.cfg
```

Usually, login events, which operate on profiles and do merge on protected properties, must be secured. For each trusted third party server, you need to add these 3 lines :

```
thirdparty.provider1.key=secret-key
thirdparty.provider1.ipAddresses=127.0.0.1,::1
thirdparty.provider1.allowedEvents=login,download
```

The events set in `allowedEvents` will be secured and will only be accepted if the call comes from the specified IP address, and if the secret-key is passed in the X-Unomi-Peer header.

### 2.3.3. INSTALLING THE MAXMIND GEOPLITE2 IP LOOKUP DATABASE

The Context Server requires an IP database in order to resolve IP addresses to user location. The GeoLite2 database can be downloaded from MaxMind here : <http://dev.maxmind.com/geoip/geoip2/geolite2/>

Simply download the GeoLite2-City.mmdb file into the "etc" directory.

### 2.3.4. INSTALLING GEONAMES DATABASE

Context server includes a geocoding service based on the geonames database ( <http://www.geonames.org/> ). It can be used to create conditions on countries or cities.

In order to use it, you need to install the Geonames database into . Get the "allCountries.zip" database from here : <http://download.geonames.org/export/dump/>

Download it and put it in the "etc" directory, without unzipping it. Edit `$MY_KARAF_HOME/etc/org.apache.unomi.geonames.cfg` and set `request.geonamesDatabase.forceImport` to true, import should start right away. Otherwise, import should start at the next startup. Import runs in background, but can take about 15 minutes. At the end, you should have about 4 million entries in the geonames index.

### 2.3.5. REST API SECURITY

The Context Server REST API is protected using JAAS authentication and using Basic or Digest HTTP auth. By default, the login/password for the REST API full administrative access is "karaf/karaf".

The generated package is also configured with a default SSL certificate. You can change it by following these steps :

Replace the existing keystore in `$MY_KARAF_HOME/etc/keystore` by your own certificate :

[http://wiki.eclipse.org/Jetty/Howto/Configure\\_SSL](http://wiki.eclipse.org/Jetty/Howto/Configure_SSL)

Update the keystore and certificate password in `$MY_KARAF_HOME/etc/custom.properties` file :

```
org.osgi.service.http.secure.enabled = true
org.ops4j.pax.web.ssl.keystore=${karaf.etc}/keystore
org.ops4j.pax.web.ssl.password=changeme
org.ops4j.pax.web.ssl.keypassword=changeme
org.osgi.service.http.port.secure=9443
```

You should now have SSL setup on Karaf with your certificate, and you can test it by trying to access it on port 9443.

.

Changing the default Karaf password can be done by modifying the etc/users.properties file

.

You will also need to change the user/password information in the org.apache.unomi.cluster.cfg file :

```
cluster.group=default
cluster.jmxUsername=karaf
cluster.jmxPassword=karaf
cluster.jmxPort=1099
```

### 2.3.6. AUTOMATIC PROFILE MERGING

The context server is capable of merging profiles based on a common property value. In order to use this, you must add the MergeProfileOnPropertyAction to a rule (such as a login rule for example), and configure it with the name of the property that will be used to identify the profiles to be merged. An example could be the "email" property, meaning that if two (or more) profiles are found to have the same value for the "email" property they will be merged by this action.

Upon merge, the old profiles are marked with a "mergedWith" property that will be used on next profile access to delete the original profile and replace it with the merged profile (aka "master" profile). Once this is done, all cookie tracking will use the merged profile.

To test, simply configure the action in the "login" or "facebookLogin" rules and set it up on the "email" property. Upon sending one of the events, all matching profiles will be merged.

### 2.3.7. SECURING A PRODUCTION ENVIRONMENT

Before going live with a project, you should *absolutely* read the following section that will help you setup a proper secure environment for running your context server.

Step 1: Install and configure a firewall

You should setup a firewall around your cluster of context servers and/or Elasticsearch nodes. If you have an application-level firewall you should only allow the following connections open to the whole world :

- <http://localhost:8181/context.js>
- <http://localhost:8181/eventcollector>

All other ports should not be accessible to the world.

For your Context Server client applications (such as the Jahia CMS), you will need to make the following ports accessible :

```
8181 (Context Server HTTP port)
9443 (Context Server HTTPS port)
```

The context server actually requires HTTP Basic Auth for access to the Context Server administration REST API, so it is highly recommended that you design your client applications to use the HTTPS port for accessing the REST API.

The user accounts to access the REST API are actually routed through Karaf's JAAS support, which you may find the documentation for here :

- <http://karaf.apache.org/manual/latest/users-guide/security.html>

The default username/password is

```
karaf/karaf
```

You should really change this default username/password as soon as possible. To do so, simply modify the following file :

```
$MY_KARAF_HOME/etc/users.properties
```

For your context servers, and for any standalone Elasticsearch nodes you will need to open the following ports for proper node-to-node communication : 9200 (Elasticsearch REST API), 9300 (Elasticsearch TCP transport)

Of course any ports listed here are the default ports configured in each server, you may adjust them if needed.

Step 2 : Follow industry recommended best practices for securing Elasticsearch

You may find more valuable recommendations here :

- <https://www.elastic.co/blog/found-elasticsearch-security>
- <https://www.elastic.co/blog/scripting-security>

Step 4 : Setup a proxy in front of the context server

As an alternative to an application-level firewall, you could also route all traffic to the context server through a proxy, and use it to filter any communication.

### 2.3.8. INTEGRATING WITH AN APACHE HTTP WEB SERVER

If you want to setup an Apache HTTP web server in from of Apache Unomi, here is an example configuration using `mod_proxy`.

In your Unomi package directory, in `/etc/org.apache.unomi.cluster.cfg` for `unomi.apache.org`

```
contextserver.publicAddress=https://unomi.apache.org/  
contextserver.internalAddress=http://192.168.1.1:8181
```

and you will also need to change the `contextserver.domain` in the `/etc/org.apache.unomi.web.cfg` file

```
contextserver.domain=apache.org
```

Main virtual host config:

```
<VirtualHost *:80>  
    Include /var/www/vhosts/unomi.apache.org/conf/common.conf  
</VirtualHost>  
  
<IfModule mod_ssl.c>  
    <VirtualHost *:443>  
        Include /var/www/vhosts/unomi.apache.org/conf/common.conf  
  
        SSLEngine on  
  
        SSLCertificateFile /var/www/vhosts/unomi.apache.org/conf/ssl/24d5b9691e96eafa.crt  
        SSLCertificateKeyFile /var/www/vhosts/unomi.apache.org/conf/ssl/apache.org.key  
        SSLCertificateChainFile /var/www/vhosts/unomi.apache.org/conf/ssl/gd_bundle-g2-g1.crt  
  
        <FilesMatch "\.(cgi|shtml|phtml|php)$">  
            SSLOptions +StdEnvVars  
        </FilesMatch>  
        <Directory /usr/lib/cgi-bin>  
            SSLOptions +StdEnvVars  
        </Directory>  
        BrowserMatch "MSIE [2-6]" \  
            nokeepalive ssl-unclean-shutdown \  
            downgrade-1.0 force-response-1.0  
        BrowserMatch "MSIE [17-9]" ssl-unclean-shutdown  
  
    </VirtualHost>  
</IfModule>
```

common.conf:

```

ServerName unomi.apache.org
ServerAdmin webmaster@apache.org

DocumentRoot /var/www/vhosts/unomi.apache.org/html
CustomLog /var/log/apache2/access-unomi.apache.org.log combined
<Directory />
    Options FollowSymLinks
    AllowOverride None
</Directory>
<Directory /var/www/vhosts/unomi.apache.org/html>
    Options FollowSymLinks MultiViews
    AllowOverride None
    Order allow,deny
    allow from all
</Directory>
<Location /cxs>
    Order deny,allow
    deny from all
    allow from 88.198.26.2
    allow from www.apache.org
</Location>

RewriteEngine On
RewriteCond %{REQUEST_METHOD} ^(TRACE|TRACK)
RewriteRule .* - [F]
ProxyPreserveHost On
ProxyPass /server-status !
ProxyPass /robots.txt !

RewriteCond %{HTTP_USER_AGENT} Googlebot [OR]
RewriteCond %{HTTP_USER_AGENT} msnbot [OR]
RewriteCond %{HTTP_USER_AGENT} Slurp
RewriteRule ^.* - [F,L]

ProxyPass / http://localhost:8181/ connectiontimeout=20 timeout=300 ttl=120
ProxyPassReverse / http://localhost:8181/

```

### 2.3.9. CHANGING THE DEFAULT TRACKING LOCATION

When performing localhost requests to Apache Unomi, a default location will be used to insert values into the session to make the location-based personalization still work. You can find the default location settings in the file :

```
org.apache.unomi.plugins.request.cfg
```

that contains the following default settings:

```
# The following settings represent the default position that is used for localhost requests
defaultSessionCountryCode=CH
defaultSessionCountryName=Switzerland
defaultSessionCity=Geneva
defaultSessionAdminSubDiv1=2660645
defaultSessionAdminSubDiv2=6458783
defaultSessionIsp=Cablecom
defaultLatitude=46.1884341
defaultLongitude=6.1282508
```

You might want to change these for testing or for demonstration purposes.

### 2.3.10. APACHE KARAF SSH CONSOLE

The Apache Karaf SSH console is available inside Apache Unomi, but the port has been changed from the default value of 8101 to 8102 to avoid conflicts with other Karaf-based products. So to connect to the SSH console you should use:

```
ssh -p 8102 karaf@localhost
```

or the user/password you have setup to protect the system if you have changed it.

## 3. SAMPLE

### 3.1. SAMPLES

Apache Unomi provides the following samples:

- [Twitter integration](#)
- [Login integration](#)

### 3.2. LOGIN SAMPLE

This sample is an example of what is involved in integrated a login with Apache Unomi.

#### 3.2.1. WARNING !

The example code uses client-side Javascript code to send the login event. This is only done this way for the sake of sample simplicity but it should NEVER BE DONE THIS WAY in real cases.

The login event should always be sent from the server performing the actual login since it must only be sent if the user has authenticated properly, and only the authentication server can validate this.

## 3.2.2. INSTALLING THE SAMPLE

Login into the Unomi Karaf SSH shell using something like this :

```
ssh -p 8102 karaf@localhost (default password is karaf)
```

Install the login sample using the following command:

```
bundle:install mvn:org.apache.unomi/login-integration-sample/1.2.0-incubating-SNAPSHOT
```

when the bundle is successfully install you will get an bundle ID back we will call it BUNDLE\_ID.

You can then do:

```
bundle:start BUNDLE_ID
```

If all went well you can access the login sample HTML page here :

```
http://localhost:8181/login/index.html
```

You can fill in the form to test it. Note that the hardcoded password is:

```
test1234
```

## 3.3. TWITTER SAMPLE

### 3.3.1. OVERVIEW

We will examine how a simple HTML page can interact with Unomi to enrich a user's profile. The use case we will follow is a rather simple one: we use a Twitter button to record the number of times the visitor tweeted (as a `tweetNb` profile integer property) as well as the URLs they tweeted from (as a `tweetedFrom` multi-valued string profile property). A javascript script will use the Twitter API to react to clicks on this button and update the user profile using a `ContextServlet` request triggering a custom



event. This event will, in turn, trigger a Unomi action on the server implemented using a Unomi plugin, a standard extension point for the server.

## BUILDING THE TWEET BUTTON SAMPLE

In your local copy of the Unomi repository and run:

```
cd samples/tweet-button-plugin
mvn clean install
```

This will compile and create the OSGi bundle that can be deployed on Unomi to extend it.

## DEPLOYING THE TWEET BUTTON SAMPLE

In standard Karaf fashion, you will need to copy the sample bundle to your Karaf [deploy](#) directory.

If you are using the packaged version of Unomi (as opposed to deploying it to your own Karaf version), you can simply run, assuming your current directory is [samples/tweet-button-plugin](#) and that you uncompressed the archive in the directory it was created:

```
cp target/tweet-button-plugin-1.0.0-incubating-SNAPSHOT.jar ../../package/target/unomi-1.0.0-incubating-SNAPSHOT/deploy
```

## TESTING THE SAMPLE

You can now go to <http://localhost:8181/index.html> to test the sample code. The page is very simple, you will see a Twitter button, which, once clicked, will open a new window to tweet about the current page. The original page should be updated with the new values of the properties coming from Unomi. Additionally, the raw JSON response is displayed.

We will now explain in greater details some concepts and see how the example works.

### 3.3.2. INTERACTING WITH THE CONTEXT SERVER

There are essentially two modalities to interact with the context server, reflecting different types of Unomi users: context server clients and context server integrators.

**Context server clients** are usually web applications or content management systems. They interact with Unomi by providing raw, uninterpreted contextual data in the form of events and associated metadata. That contextual data is then processed by the context server to be fed to clients once actionable. In that sense context server clients are both consumers and producers of contextual data. Context server clients will mostly interact with Unomi using a single entry point called the [ContextServlet](#), requesting context for the current user and providing any triggered events along the way.

On the other hand, **context server integrators** provide ways to feed more structured data to the context server either to integrate with third party services or to provide analysis of the uninterpreted data

provided by context server clients. Such integration will mostly be done using Unomi's API either directly using Unomi plugins or via the provided REST APIs. However, access to REST APIs is restricted due for security reasons, requiring privileged access to the Unomi server, making things a little more complex to set up.

For simplicity's sake, this document will focus solely on the first use case and will interact only with the context servlet.

### 3.3.3. RETRIEVING CONTEXT INFORMATION FROM UNOMI USING THE CONTEXT SERVLET

Unomi provides two ways to retrieve context: either as a pure JSON object containing strictly context information or as a couple of JSON objects augmented with javascript functions that can be used to interact with the Unomi server using the `<context server base URL>/context.json` or `<context server base URL>/context.js` URLs, respectively.

Below is an example of asynchronously loading the initial context using the javascript version, assuming a default Unomi install running on <http://localhost:8181>:

```
// Load context from Unomi asynchronously
(function (document, elementToCreate, id) {
  var js, fjs = document.getElementsByTagName(elementToCreate)[0];
  if (document.getElementById(id)) return;
  js = document.createElement(elementToCreate);
  js.id = id;
  js.src = 'http://localhost:8181/context.js';
  fjs.parentNode.insertBefore(js, fjs);
})(document, 'script', 'context');
```

This initial context results in a javascript file providing some functions to interact with the context server from javascript along with two objects: a `cxs` object containing information about the context for the current user and a `digitalData` object that is injected into the browser's `window` object (leveraging the [Customer Experience Digital Data Layer](#) standard). Note that this last object is not under control of the context server and clients are free to use it or not. Our example will not make use of it.

On the other hand, the `cxs` top level object contains interesting contextual information about the current user:

```
{
  "profileId":<identifier of the profile associated with the current user>,
  "sessionId":<identifier of the current user session>,
  "profileProperties":<requested profile properties, if any>,
  "sessionProperties":<requested session properties, if any>,
  "profileSegments":<segments the profile is part of if requested>,
  "filteringResults":<result of the evaluation of personalization filters>,
  "trackedConditions":<tracked conditions in the source page, if any>
}
```

We will look at the details of the context request and response later.

## 3.4. EXAMPLE

### 3.4.1. HTML PAGE

The code for the HTML page with our Tweet button can be found at <https://github.com/apache/incubator-unomi/blob/master/wab/src/main/webapp/index.html>.

This HTML page is fairly straightforward: we create a tweet button using the Twitter API while a Javascript script performs the actual logic.

### 3.4.2. JAVASCRIPT

Globally, the script loads both the twitter widget and the initial context asynchronously (as shown previously). This is accomplished using fairly standard javascript code and we won't look at it here. Using the Twitter API, we react to the `tweet` event and call the Unomi server to update the user's profile with the required information, triggering a custom `tweetEvent` event. This is accomplished using a `contextRequest` function which is an extended version of a classic `AJAX` request:

```

function contextRequest(successCallback, errorCallback, payload) {
  var data = JSON.stringify(payload);
  // if we don't already have a session id, generate one
  var sessionId = cxs.sessionId || generateUUID();
  var url = 'http://localhost:8181/context.json?sessionId=' + sessionId;
  var xhr = new XMLHttpRequest();
  var isGet = data.length < 100;
  if (isGet) {
    xhr.withCredentials = true;
    xhr.open("GET", url + "&payload=" + encodeURIComponent(data), true);
  } else if ("withCredentials" in xhr) {
    xhr.open("POST", url, true);
    xhr.withCredentials = true;
  } else if (typeof XMLHttpRequest != "undefined") {
    xhr = new XMLHttpRequest();
    xhr.open("POST", url);
  }
  xhr.onreadystatechange = function () {
    if (xhr.readyState != 4) {
      return;
    }
    if (xhr.status == 200) {
      var response = xhr.responseText ? JSON.parse(xhr.responseText) : undefined;
      if (response) {
        cxs.sessionId = response.sessionId;
        successCallback(response);
      }
    } else {
      console.log("contextserver: " + xhr.status + " ERROR: " + xhr.statusText);
      if (errorCallback) {
        errorCallback(xhr);
      }
    }
  };
  xhr.setRequestHeader("Content-Type", "text/plain;charset=UTF-8"); // Use text/plain to avoid CORS
  preflight
  if (isGet) {
    xhr.send();
  } else {
    xhr.send(data);
  }
}

```

There are a couple of things to note here:

- If we specify a payload, it is expected to use the JSON format so we [stringify](#) it and encode it if passed as a URL parameter in a [GET](#) request.
- We need to make a [CORS](#) request since the Unomi server is most likely not running on the same host than the one from which the request originates. The specific details are fairly standard and we will not explain them here.
- We need to either retrieve (from the initial context we retrieved previously using [cxs.sessionId](#)) or generate a session identifier for our request since Unomi currently requires one.

- We’re calling the `ContextServlet` using the default install URI, specifying the session identifier: `<a href="http://localhost:8181/context.json?sessionId=&#39" class="bare">http://localhost:8181/context.json?sessionId=&#39</a> + sessionId</code>. This URI requests context from Unomi, resulting in an updated cxs</code> object in the javascript global scope. The context server can reply to this request either by returning a JSON-only object containing solely the context information as is the case when the requested URI is context.json</code>. However, if the client requests context.js</code> then useful functions to interact with Unomi are added to the cxs</code> object in addition to the context information as depicted above.`
- We don’t need to provide any authentication at all to interact with this part of Unomi since we only have access to read-only data (as well as providing events as we shall see later on). If we had been using the REST API, we would have needed to provide authentication information as well.

## CONTEXT REQUEST AND RESPONSE STRUCTURE

The interesting part, though, is the payload. This is where we provide Unomi with contextual information as well as ask for data in return. This allows clients to specify which type of information they are interested in getting from the context server as well as specify incoming events or content filtering or property/segment overrides for personalization or impersonation. This conditions what the context server will return with its response.

Let’s look at the context request structure:

```
{
  source: <Item source of the context request>,
  events: <optional array of triggered events>,
  requiredProfileProperties: <optional array of property identifiers>,
  requiredSessionProperties: <optional array of property identifiers>,
  filters: <optional array of filters to evaluate>,
  profileOverrides: <optional profile containing segments,scores or profile properties to override>,
    - segments: <optional array of segment identifiers>,
    - profileProperties: <optional map of property name / value pairs>,
    - scores: <optional map of score id / value pairs>
  sessionPropertiesOverrides: <optional map of property name / value pairs>,
  requireSegments: <boolean, whether to return the associated segments>
}
```

We will now look at each part in greater details.

### SOURCE

A context request payload needs to at least specify some information about the source of the request in the form of an [Item](#) (meaning identifier, type and scope plus any additional properties we might have to provide), via the `source` property of the payload. Of course the more information can be provided about the source, the better.

## FILTERS

A client wishing to perform content personalization might also specify filtering conditions to be evaluated by the context server so that it can tell the client whether the content associated with the filter should be activated for this profile/session. This is accomplished by providing a list of filter definitions to be evaluated by the context server via the `filters` field of the payload. If provided, the evaluation results will be provided in the `filteringResults` field of the resulting `cxs` object the context server will send.

## OVERRIDES

It is also possible for clients wishing to perform user impersonation to specify properties or segments to override the proper ones so as to emulate a specific profile, in which case the overridden value will temporarily replace the proper values so that all rules will be evaluated with these values instead of the proper ones. The `segments` (array of segment identifiers), `profileProperties` (maps of property name and associated object value) and `scores` (maps of score id and value) all wrapped in a `profileOverrides` object and the `sessionPropertiesOverrides` (maps of property name and associated object value) fields allow to provide such information. Providing such overrides will, of course, impact content filtering results and segments matching for this specific request.

## CONTROLLING THE CONTENT OF THE RESPONSE

The clients can also specify which information to include in the response by setting the `requireSegments` property to true if segments the current profile matches should be returned or provide an array of property identifiers for `requiredProfileProperties` or `requiredSessionProperties` fields to ask the context server to return the values for the specified profile or session properties, respectively. This information is provided by the `profileProperties`, `sessionProperties` and `profileSegments` fields of the context server response.

Additionally, the context server will also return any tracked conditions associated with the source of the context request. Upon evaluating the incoming request, the context server will determine if there are any rules marked with the `trackedCondition` tag and which source condition matches the source of the incoming request and return these tracked conditions to the client. The client can use these tracked conditions to learn that the context server can react to events matching the tracked condition and coming from that source. This is, in particular, used to implement form mapping (a solution that allows clients to update user profiles based on values provided when a form is submitted).

## EVENTS

Finally, the client can specify any events triggered by the user actions, so that the context server can process them, via the `events` field of the context request.

## DEFAULT RESPONSE

If no payload is specified, the context server will simply return the minimal information deemed necessary for client applications to properly function: profile identifier, session identifier and any tracked conditions that might exist for the source of the request.

## CONTEXT REQUEST FOR OUR EXAMPLE

Now that we've seen the structure of the request and what we can expect from the context response, let's examine the request our component is doing.

In our case, our `source` item looks as follows: we specify a scope for our application (`unomi-tweet-button-sample`), specify that the item type (i.e. the kind of element that is the source of our event) is a `page` (which corresponds, as would be expected, to a web page), provide an identifier (in our case, a Base-64 encoded version of the page's URL) and finally, specify extra properties (here, simply a `url` property corresponding to the page's URL that will be used when we process our event in our Unomi extension).

```
var scope = 'unomi-tweet-button-sample';
var itemId = btoa(window.location.href);
var source = {
  itemType: 'page',
  scope: scope,
  itemId: itemId,
  properties: {
    url: window.location.href
  }
};
```

We also specify that we want the context server to return the values of the `tweetNb` and `tweetedFrom` profile properties in its response. Finally, we provide a custom event of type `tweetEvent` with associated scope and source information, which matches the source of our context request in this case.

```
var contextPayload = {
  source: source,
  events: [
    {
      eventType: 'tweetEvent',
      scope: scope,
      source: source
    }
  ],
  requiredProfileProperties: [
    'tweetNb',
    'tweetedFrom'
  ]
};
```

The `tweetEvent` event type is not defined by default in Unomi. This is where our Unomi plugin comes into play since we need to tell Unomi how to react when it encounters such events.

## UNOMI PLUGIN OVERVIEW

In order to react to `tweetEvent` events, we will define a new Unomi rule since this is exactly what Unomi rules are supposed to do. Rules are guarded by conditions and if these conditions match, the associated set of actions will be executed. In our case, we want our new `incrementTweetNumber` rule to only react

to [tweetEvent](#) events and we want it to perform the profile update accordingly: create the property types for our custom properties if they don't exist and update them. To do so, we will create a custom [incrementTweetNumberAction](#) action that will be triggered any time our rule matches. An action is some custom code that is deployed in the context server and can access the Unomi API to perform what it is that it needs to do.

## RULE DEFINITION

Let's look at how our custom [incrementTweetNumber](#) rule is defined:

```
{
  "metadata": {
    "id": "smp:incrementTweetNumber",
    "name": "Increment tweet number",
    "description": "Increments the number of times a user has tweeted after they click on a tweet
button"
  },
  "raiseEventOnlyOnceForSession": false,
  "condition": {
    "type": "eventTypeCondition",
    "parameterValues": {
      "eventId": "tweetEvent"
    }
  },
  "actions": [
    {
      "type": "incrementTweetNumberAction",
      "parameterValues": {}
    }
  ]
}
```

Rules define a metadata section where we specify the rule name, identifier and description.

When rules trigger, a specific event is raised so that other parts of Unomi can react to it accordingly. We can control how that event should be raised. Here we specify that the event should be raised each time the rule triggers and not only once per session by setting [raiseEventOnlyOnceForSession](#) to `false`, which is not strictly required since that is the default. A similar setting ([raiseEventOnlyOnceForProfile](#)) can be used to specify that the event should only be raised once per profile if needed.

We could also specify a priority for our rule in case it needs to be executed before other ones when similar conditions match. This is accomplished using the [priority](#) property. We're using the default priority here since we don't have other rules triggering on `tweetEvent`'s and don't need any special ordering.

We then tell Unomi which condition should trigger the rule via the [condition](#) property. Here, we specify that we want our rule to trigger on an [eventTypeCondition](#) condition. Unomi can be extended by adding new condition types that can enrich how matching or querying is performed. The condition type definition file specifies which parameters are expected for our condition to be complete. In our case, we use the built-in event type condition that will match if Unomi receives an event of the type specified in the condition's [eventId](#) parameter value: `tweetEvent` here.



Finally, we specify a list of actions that should be performed as consequences of the rule matching. We only need one action of type `incrementTweetNumberAction` that doesn't require any parameters.

## ACTION DEFINITION

Let's now look at our custom `incrementTweetNumberAction` action type definition:

```
{
  "id": "incrementTweetNumberAction",
  "actionExecutor": "incrementTweetNumber",
  "tags": [
    "event"
  ],
  "parameters": []
}
```

We specify the identifier for the action type, a list of tags if needed: here we say that our action is a consequence of events using the `event` tag. Our actions does not require any parameters so we don't define any.

Finally, we provide a mysterious `actionExecutor` identifier: `incrementTweetNumber`.

## ACTION EXECUTOR DEFINITION

The action executor references the actual implementation of the action as defined in our [blueprint definition](#):

```
<blueprint xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xsi:schemaLocation="http://www.osgi.org/xmlns/blueprint/v1.0.0
http://www.osgi.org/xmlns/blueprint/v1.0.0/blueprint.xsd">

  <reference id="profileService" interface="org.apache.unomi.api.services.ProfileService"/>

  <!-- Action executor -->
  <service id="incrementTweetNumberAction" auto-export="interfaces">
    <service-properties>
      <entry key="actionExecutorId" value="incrementTweetNumber"/>
    </service-properties>
    <bean
class="org.apache.unomi.examples.unomi_tweet_button_plugin.actions.IncrementTweetNumberAct
ion">
      <property name="profileService" ref="profileService"/>
    </bean>
  </service>
</blueprint>
```

In standard Blueprint fashion, we specify that we will need the `profileService` defined by Unomi and then define a service of our own to be exported for Unomi to use. Our service specifies one property: `actionExecutorId` which matches the identifier we specified in our action definition. We then inject the

profile service in our executor and we're done for the configuration side of things!

## ACTION EXECUTOR IMPLEMENTATION

Our action executor definition specifies that the bean providing the service is implemented in the `org.apache.unomi.samples.tweet_button_plugin.actions.IncrementTweetNumberAction` class. This class implements the Unomi `ActionExecutor` interface which provides a single `int execute(Action action, Event event)` method: the executor gets the action instance to execute along with the event that triggered it, performs its work and returns an integer status corresponding to what happened as defined by public constants of the `EventService` interface of Unomi: `NO_CHANGE`, `SESSION_UPDATED` or `PROFILE_UPDATED`.

Let's now look at the implementation of the method:

```
final Profile profile = event.getProfile();
Integer tweetNb = (Integer) profile.getProperty(TWEET_NB_PROPERTY);
List<String> tweetedFrom = (List<String>) profile.getProperty(TWEETED_FROM_PROPERTY);

if (tweetNb == null || tweetedFrom == null) {
    // create tweet number property type
    PropertyType propertyType = new PropertyType(new Metadata(event.getScope(),
TWEET_NB_PROPERTY, TWEET_NB_PROPERTY, "Number of times a user tweeted"));
    propertyType.setValueTypeId("integer");
    service.createPropertyType(propertyType);

    // create tweeted from property type
    propertyType = new PropertyType(new Metadata(event.getScope(), TWEETED_FROM_PROPERTY,
TWEETED_FROM_PROPERTY, "The list of pages a user tweeted from"));
    propertyType.setValueTypeId("string");
    propertyType.setMultivalued(true);
    service.createPropertyType(propertyType);

    tweetNb = 0;
    tweetedFrom = new ArrayList<>();
}

profile.setProperty(TWEET_NB_PROPERTY, tweetNb + 1);
final String sourceURL = extractSourceURL(event);
if (sourceURL != null) {
    tweetedFrom.add(sourceURL);
}
profile.setProperty(TWEETED_FROM_PROPERTY, tweetedFrom);

return EventService.PROFILE_UPDATED;
```

It is fairly straightforward: we retrieve the profile associated with the event that triggered the rule and check whether it already has the properties we are interested in. If not, we create the associated property types and initialize the property values.

Note that it is not an issue to attempt to create the same property type multiple times as Unomi will not add a new property type if an identical type already exists.

Once this is done, we update our profile with the new property values based on the previous values and the metadata extracted from the event using the `extractSourceURL` method which uses our `url` property that we've specified for our event source. We then return that the profile was updated as a result of our action and Unomi will properly save it for us when appropriate. That's it!

For reference, here's the `extractSourceURL` method implementation:

```
private String extractSourceURL(Event event) {
    final Item sourceAsItem = event.getSource();
    if (sourceAsItem instanceof CustomItem) {
        CustomItem source = (CustomItem) sourceAsItem;
        final String url = (String) source.getProperties().get("url");
        if (url != null) {
            return url;
        }
    }

    return null;
}
```

## 3.5. CONCLUSION

We have seen a simple example how to interact with Unomi using a combination of client-side code and Unomi plugin. Hopefully, this provided an introduction to the power of what Unomi can do and how it can be extended to suit your needs.

## 3.6. ANNEX

Here is an overview of how Unomi processes incoming requests to the `ContextServlet`. [Unomi request overview]

## 3.7. WEATHER UPDATE SAMPLE

# 4. CONNECTORS

## 4.1. CONNECTORS

Apache Unomi provides the following connector:

- [Salesforce CRM connector](#)

### 4.1.1. CALL FOR CONTRIBUTORS

We are looking for help with the development of additional connectors. Any contribution (large or small) is more than welcome. Feel free to discuss this in our [mailing list](#).

## 4.2. APACHE UNOMI SALESFORCE CONNECTOR

This connector makes it possible to push and pull data to/from the Salesforce CRM. It can copy information between Apache Unomi profiles and Salesforce Leads.

### 4.2.1. GETTING STARTED

.

Create a new developer account here:

<https://developer.salesforce.com/signup>

.

Create a new Connected App, by going into Setup -> App Manager and click "Create Connected App"

.

In the settings, make sure you do the following:

Enable OAuth settings -> Activated  
Enable for device flow -> Activated (no need for a callback URL)  
Add all the selected OAuth scopes you want (or put all of them)  
Make sure Require Secret for Web Server flow is activated

.

Make sure you retrieve the following information once you have created the app in the API (Enable OAuth Settings):

Consumer key  
Consumer secret (click to see it)

.

You must also retrieve your user's security token, or create it if you don't have one already. To do this simply click on your user at the top right, select "Settings", then click on "Reset my security token". You will receive an email with the security token.

.

You are now ready to configure the Apache Unomi Salesforce Connector. In the `etc/org.apache.unomi.sfdc.cfg` file change the following settings:

```
sfdc.user.username=YOUR_USER_NAME
sfdc.user.password=YOUR_PASSWORD
sfdc.user.securityToken=YOUR_USER_SECURITY_TOKEN
sfdc.consumer.key=CONNECTED_APP_CONSUMER_KEY
sfdc.consumer.secret=CONNECTED_APP_SECRET
```

Connected to the Apache Unomi Karaf Shell using :

```
ssh -p 8102 karaf@localhost (default password is karaf)
```

Deploy into Apache Unomi using the following commands from the Apache Karaf shell:

```
feature:repo-add mvn:org.apache.unomi/unomi-salesforce-connector-karaf-kar/1.2.0-incubating-SNAPSHOT/xml/features
feature:install unomi-salesforce-connector-karaf-kar
```

You can then test the connection to Salesforce by accessing the following URLs:

```
https://localhost:9443/cxs/sfdc/version
https://localhost:9443/cxs/sfdc/limits
```

The first URL will give you information about the version of the connector, so this makes it easy to check that the plugin is properly deployed, started and the correct version. The second URL will actually make a request to the Salesforce REST API to retrieve the limits of the Salesforce API.

Both URLs are password protected by the Apache Unomi (Karaf) password. You can find this user and password information in the `etc/users.properties` file.

You can now use the connector's defined actions in rules to push or pull data to/from the Salesforce CRM. You can find more information about rules in the [Concepts](#) and the [Getting Started](#) pages.

## 4.2.2. TROUBLESHOOTING

In case you're having problem retrieving or sending data to Salesforce, make sure you check the logs

(log:display from the Karaf shell) or in the data/log/karaf.log file. It should usually tell you what went wrong.

Common problems include trying to send profiles that don't include mandatory fields in Salesforce. By default the mandatory fields are : Last name, Company and Email. If you are missing some mandatory fields, the connector will not pushed any new leads to Salesforce.

### 4.2.3. UPGRADING THE SALESFORCE CONNECTOR

If you followed all the steps in the Getting Started section, you can upgrade the Salesforce connector by using the following steps:

.

Compile the connector using:

```
cd extensions/salesforce-connector  
mvn clean install
```

.

Login to the Unomi Karaf Shell using:

```
ssh -p 8102 karaf@localhost (password by default is karaf)
```

.

Execute the following commands in the Karaf shell

```
feature:repo-refresh  
feature:uninstall unomi-salesforce-connector-karaf-feature  
feature:install unomi-salesforce-connector-karaf-feature
```

.

You can then check that the new version is properly deployed by accessing the following URL and checking the build date:

```
https://localhost:9443/cxs/sfdc/version
```

(if asked for a password it's the same karaf/karaf default)

## 4.2.4. USING THE SALESFORCE WORKBENCH FOR TESTING REST API

The Salesforce Workbench contains a REST API Explorer that is very useful to test requests. You may find it here :

```
https://workbench.developerforce.com/restExplorer.php
```

## 4.2.5. SETTING UP STREAMING PUSH QUERIES

Using the Salesforce Workbench, you can setting streaming push queries (Queries->Streaming push topics) such as the following example:

```
Name: LeadUpdates  
Query : SELECT Id,FirstName,LastName,Email,Company FROM Lead
```

## 4.2.6. EXECUTING THE UNIT TESTS

Before running the tests, make sure you have completed all the steps above, including the streaming push queries setup.

By default the unit tests will not run as they need proper Salesforce credentials to run. To set this up create a properties file like the following one:

test.properties

```
#  
# Licensed to the Apache Software Foundation (ASF) under one or more  
# contributor license agreements. See the NOTICE file distributed with  
# this work for additional information regarding copyright ownership.  
# The ASF licenses this file to You under the Apache License, Version 2.0  
# (the "License"); you may not use this file except in compliance with  
# the License. You may obtain a copy of the License at  
#  
# http://www.apache.org/licenses/LICENSE-2.0  
#  
# Unless required by applicable law or agreed to in writing, software  
# distributed under the License is distributed on an "AS IS" BASIS,  
# WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.  
# See the License for the specific language governing permissions and  
# limitations under the License.  
#  
sfdc.user.username=YOUR_USER_NAME  
sfdc.user.password=YOUR_PASSWORD  
sfdc.user.securityToken=YOUR_USER_SECURITY_TOKEN  
sfdc.consumer.key=CONNECTED_APP_CONSUMER_KEY  
sfdc.consumer.secret=CONNECTED_APP_SECRET
```

and then use the following command line to reference the file:

```
cd extensions/salesforce-connector
mvn clean install -DsfcdProperties=./test.properties
```

(in case you're wondering the ../ is because the test is located in the services sub-directory)

## 5. CLUSTER SETUP

### 5.1. CLUSTER SETUP

Apache Karaf relies on Apache Karaf Cellar, which in turn uses Hazelcast to discover and configure its cluster. You just need to install multiple context servers on the same network, and then (optionally) change the Hazelcast configuration in the following file :

```
etc/hazelcast.xml
```

All nodes on the same network, sharing the same cluster name will be part of the same cluster.

For the actual ElasticSearch configuration however, this must be done using the following file:

```
etc/org.apache.unomi.persistence.elasticsearch.cfg
```

Depending on the cluster size, you will want to adjust the following parameters to make sure your setup is optimal in terms of performance and safety.

#### 5.1.1. 2 NODES CONFIGURATION

One node dedicated to context server, 1 node for elasticsearch storage.

Node A :

```
numberOfReplicas=0
monthlyIndex.numberOfReplicas=0
```

Node B :

```
numberOfReplicas=0
monthlyIndex.numberOfReplicas=0
```

#### 5.1.2. 3 NODES CONFIGURATION

One node dedicated to context server, 2 nodes for elasticsearch storage with fault-tolerance



Node A :

```
numberOfReplicas=1  
monthlyIndex.numberOfReplicas=1
```

Node B :

```
numberOfReplicas=1  
monthlyIndex.numberOfReplicas=1
```

Node C :

```
numberOfReplicas=1  
monthlyIndex.numberOfReplicas=1
```

## 6. CUSTOM EXTENSIONS

### 6.1. CUSTOM EXTENSIONS

Apache Unomi is a pluggable server that may be extended in many ways. This document assumes you are familiar with the [Apache Unomi concepts](#) . This document is mostly a reference document on the different things that may be used inside an extension. If you are looking for complete samples, please see the [samples page](#).

#### 6.1.1. CREATING AN EXTENSION

An extension is simply a Maven project, with a Maven pom that looks like this:

```

<project xmlns="http://maven.apache.org/POM/4.0.0"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-
4.0.0.xsd">
  <parent>
    <groupId>org.apache.unomi</groupId>
    <artifactId>unomi-extensions</artifactId>
    <version>1.2.0-incubating-SNAPSHOT</version>
  </parent>

  <modelVersion>4.0.0</modelVersion>

  <artifactId>unomi-extension-example</artifactId>
  <name>Apache Unomi :: Extensions :: Example</name>
  <description>Service implementation for the Apache Unomi Context Server extension that
integrates with the Geonames database</description>
  <version>1.2.0-incubating-SNAPSHOT</version>
  <packaging>bundle</packaging>

  <dependencies>
    <!-- This dependency is not required but generally used in extensions -->
    <dependency>
      <groupId>org.apache.unomi</groupId>
      <artifactId>unomi-api</artifactId>
      <version>1.2.0-incubating-SNAPSHOT</version>
      <scope>provided</scope>
    </dependency>
  </dependencies>

  <build>
    <plugins>
      <plugin>
        <groupId>org.apache.felix</groupId>
        <artifactId>maven-bundle-plugin</artifactId>
        <extensions>>true</extensions>
        <configuration>
          <instructions>
            <Embed-Dependency>*;scope=compile|runtime</Embed-Dependency>
            <Import-Package>
              sun.misc;resolution:=optional,
              *
            </Import-Package>
          </instructions>
        </configuration>
      </plugin>
    </plugins>
  </build>
</project>

```

An extension may contain many different kinds of Apache Unomi objects, as well as custom OSGi services or anything that is needed to build your application.

## 6.1.2. PREDEFINED SEGMENTS

You may provide pre-defined segments by simply adding a JSON file in the `src/main/resources/META-INF/cxs/segments` directory of your Maven project. Here is an example of a pre-defined segment:

```
{
  "metadata": {
    "id": "leads",
    "name": "Leads",
    "scope": "systemscope",
    "description": "You can customize the list below by editing the leads segment.",
    "readOnly": true
  },
  "condition": {
    "parameterValues": {
      "subConditions": [
        {
          "parameterValues": {
            "propertyName": "properties.leadAssignedTo",
            "comparisonOperator": "exists"
          },
          "type": "profilePropertyCondition"
        }
      ],
      "operator": "and"
    },
    "type": "booleanCondition"
  }
}
```

Basically this segment uses a condition to test if the profile has a property `leadAssignedTo` that exists. All profiles that match this condition will be part of the pre-defined segment.

## 6.1.3. PREDEFINED RULES

You may provide pre-defined rules by simply adding a JSON file in the `src/main/resources/META-INF/cxs/rules` directory of your Maven project. Here is an example of a pre-defined rule:

```

{
  "metadata" : {
    "id": "evaluateProfileSegments",
    "name": "Evaluate segments",
    "description" : "Evaluate segments when a profile is modified",
    "readOnly":true
  },

  "condition" : {
    "type": "profileUpdatedEventCondition",
    "parameterValues": {
    }
  },

  "actions" : [
    {
      "type": "evaluateProfileSegmentsAction",
      "parameterValues": {
      }
    }
  ]
}

```

In this example we provide a rule that will execute when a predefined composed condition of type "profileUpdatedEventCondition" is received. See below to see how predefined composed conditions are declared. Once the condition is matched, the actions will be executed in sequence. In this example there is only a single action of type "evaluateProfileSegmentsAction" that is defined so it will be executed by Apache Unomi's rule engine. You can also see below how custom actions may be defined.

#### 6.1.4. PREDEFINED PROPERTIES

By default Apache Unomi comes with a set of pre-defined properties, but in many cases it is useful to add additional predefined property definitions. You can create property definitions for session or profile properties by creating them in different directories.

For session properties you must create a JSON file in the following directory in your Maven project:

```
src/main/resources/META-INF/cxs/properties/sessions
```

For profile properties you must create the JSON file inside the directory in your Maven project:

```
src/main/resources/META-INF/cxs/properties/profiles
```

Here is an example of a property definition JSON file

```
{
  "metadata": { "id": "city", "name": "City" },
  "type": "string",
  "tags": ["contactProfileProperties"],
  "defaultValue": "",
  "automaticMappingsFrom": [],
  "rank": "304.0"
}
```

### 6.1.5. PREDEFINED CHILD CONDITIONS

You can define new predefined conditions that are actually conditions inheriting from a parent condition and setting pre-defined parameter values. You can do this by creating a JSON file in:

```
src/main/resources/META-INF/cxs/conditions
```

Here is an example of a JSON file that defines a `profileUpdateEventCondition` that inherits from a parent condition of type `eventTypeCondition`.

```
{
  "metadata": {
    "id": "profileUpdatedEventCondition",
    "name": "profileUpdatedEventCondition",
    "description": "",
    "tags": [
      "event",
      "eventCondition"
    ],
    "readOnly": true
  },
  "parentCondition": {
    "type": "eventTypeCondition",
    "parameterValues": {
      "eventType": "profileUpdated"
    }
  },
  "parameters": [
  ]
}
```

### 6.1.6. PREDEFINED PERSONAS

Personas may also be pre-defined by creating JSON files in the following directory:

```
src/main/resources/META-INF/cxs/personas
```

Here is an example of a persona definition JSON file:

```

{
  "persona": {
    "itemId": "usVisitor",
    "properties": {
      "description": "Represents a visitor browsing from inside the continental US",
      "firstName": "U.S.",
      "lastName": "Visitor"
    },
    "segments": []
  },
  "sessions": [
    {
      "itemId": "aa3b04bd-8f4d-4a07-8e96-d33ffa04d3d9",
      "profileId": "usVisitor",
      "properties": {
        "operatingSystemName": "OS X 10.9 Mavericks",
        "sessionCountryName": "United States",
        "location": {
          "lat": 37.422,
          "lon": -122.084058
        },
        "userAgentVersion": "37.0.2062.120",
        "sessionCountryCode": "US",
        "deviceCategory": "Personal computer",
        "operatingSystemFamily": "OS X",
        "userAgentName": "Chrome",
        "sessionCity": "Mountain View",
        "remoteHost": "www.google.com",
        "remoteAddr": "66.249.66.1"
      },
      "timeStamp": "2014-09-18T11:40:54Z",
      "lastEventDate": "2014-09-18T11:40:59Z",
      "duration": 4790
    }
  ]
}

```

You can see that it's also possible to define sessions for personas.

### 6.1.7. CUSTOM ACTIONS

Custom actions are a powerful way to integrate with external systems by being able to define custom logic that will be executed by an Apache Unomi rule. An action is defined by a JSON file created in the following directory:

```
src/main/resources/META-INF/cxs/actions
```

Here is an example of a JSON action definition:

```

{
  "metadata": {
    "id": "addToListsAction",
    "name": "addToListsAction",
    "description": "",
    "tags": [
      "demographic",
      "hidden.availableToEndUser"
    ],
    "readOnly": true
  },
  "actionExecutor": "addToLists",
  "parameters": [
    {
      "id": "listIdentifiers",
      "type": "string",
      "multivalued": true
    }
  ]
}

```

The `actionExecutor` identifier refers to a service property that is defined in the OSGi Blueprint service registration. Note that any OSGi service registration may be used, but in these examples we use OSGi Blueprint. The definition for the above JSON file will be found in a file called `src/main/resources/OSGI-INF/blueprint/blueprint.xml` with the following content:

```

<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xsi:schemaLocation="http://www.osgi.org/xmlns/blueprint/v1.0.0
  http://www.osgi.org/xmlns/blueprint/v1.0.0/blueprint.xsd">

  <reference id="profileService" interface="org.apache.unomi.api.services.ProfileService"/>
  <reference id="eventService" interface="org.apache.unomi.api.services.EventService"/>

  <!-- Action executors -->

  <service auto-export="interfaces">
    <service-properties>
      <entry key="actionExecutorId" value="addToLists"/>
    </service-properties>
    <bean class="org.apache.unomi.lists.actions.AddToListsAction">
      <property name="profileService" ref="profileService"/>
      <property name="eventService" ref="eventService"/>
    </bean>
  </service>

</blueprint>

```

You can note here the `actionExecutorId` that corresponds to the `actionExecutor` in the JSON file.

The implementation of the action is available here : `org.apache.unomi.lists.actions.AddToListsAction`

## 6.1.8. CUSTOM CONDITIONS

Custom conditions are different from predefined child conditions because they implement their logic using Java classes. They are also declared by adding a JSON file into the conditions directory:

```
src/main/resources/META-INF/cxs/conditions
```

Here is an example of JSON custom condition definition:

```
{
  "metadata": {
    "id": "matchAllCondition",
    "name": "matchAllCondition",
    "description": "",
    "tags": [
      "logical",
      "profileCondition",
      "eventCondition",
      "sessionCondition",
      "sourceEventCondition"
    ],
    "readOnly": true
  },
  "conditionEvaluator": "matchAllConditionEvaluator",
  "queryBuilder": "matchAllConditionESQueryBuilder",

  "parameters": [
  ]
}
```

Note the `conditionEvaluator` and the `queryBuilder` values. These reference OSGi service properties that are declared in an OSGi Blueprint configuration file (other service definitions may also be used such as Declarative Services or even Java registered services). Here is an example of an OSGi Blueprint definition corresponding to the above JSON condition definition file.



src/main/resources/OSGI-INF/blueprint/blueprint.xml

```
<blueprint xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
xsi:schemaLocation="http://www.osgi.org/xmlns/blueprint/v1.0.0
http://www.osgi.org/xmlns/blueprint/v1.0.0/blueprint.xsd">

  <service

interface="org.apache.unomi.persistence.elasticsearch.conditions.ConditionESQueryBuilder">
  <service-properties>
    <entry key="queryBuilderId" value="matchAllConditionESQueryBuilder"/>
  </service-properties>
  <bean
class="org.apache.unomi.plugins.baseplugin.conditions.MatchAllConditionESQueryBuilder"/>
</service>

  <service interface="org.apache.unomi.persistence.elasticsearch.conditions.ConditionEvaluator">
  <service-properties>
    <entry key="conditionEvaluatorId" value="matchAllConditionEvaluator"/>
  </service-properties>
  <bean class="org.apache.unomi.plugins.baseplugin.conditions.MatchAllConditionEvaluator"/>
</service>

</blueprint>
```

You can find the implementation of the two classes here :

- [org.apache.unomi.plugins.baseplugin.conditions.MatchAllConditionESQueryBuilder](#)
- [org.apache.unomi.plugins.baseplugin.conditions.MatchAllConditionEvaluator](#)