

# APACHE UNOMI 1.X - DOCUMENTATION

## TABLE OF CONTENTS

1. Quick start	5
1.1. Five Minutes QuickStart	5
2. First steps with Apache Unomi	6
2.1. Getting started with Unomi	6
2.1.1. Prerequisites	6
2.1.2. Running Unomi	7
2.2. Recipes	7
2.2.1. Introduction	7
2.2.2. How to read a profile	7
2.2.3. How to update a profile from the public internet	8
2.2.4. How to search for profile events	11
2.2.5. How to create a new rule	12
2.2.6. How to search for profiles	12
2.2.7. Getting / updating consents	13
2.2.8. How to send a login event to Unomi	13
2.3. Request examples	14
2.3.1. Retrieving your first context	14
2.3.2. Retrieving a context as a JSON object	15
2.3.3. Accessing profile properties in a context	15
2.3.4. Sending events using the context servlet	15
2.3.5. Sending events using the eventcollector servlet	16
2.3.6. Where to go from here	17
2.4. Web Tracker	17
2.4.1. Getting started	17
2.4.2. How to contribute	18
2.4.3. Tracking page views	18
2.4.4. Tracking form submissions	21
3. Configuration	27
3.1. Centralized configuration	27
3.2. Changing the default configuration using environment variables (i.e. Docker configuration)	27
3.3. Changing the default configuration using property files	28
3.4. Secured events configuration	28
3.5. Installing the MaxMind GeoIPLite2 IP lookup database	30
3.6. Installing Geonames database	30
3.7. REST API Security	30
3.8. Scripting security	31
3.8.1. Multi-layer scripting filtering system	31
3.8.2. Scripts and expressions	32
3.8.3. Scripting expression filtering configuration parameters	34
3.8.4. Groovy Actions	36
3.8.5. Scripting roadmap	39
3.9. Automatic profile merging	39
3.10. Securing a production environment	39
3.11. Integrating with an Apache HTTP web server	41

3.12. Changing the default tracking location	43
3.13. Apache Karaf SSH Console	44
3.14. Elasticsearch authentication and security	44
3.14.1. User authentication !	44
3.14.2. SSL communication	45
4. Migrations	45
4.1. From version 1.4 to 1.5	45
4.1.1. Data model and Elasticsearch 7	45
4.1.2. API changes	46
4.1.3. Migration steps	46
4.2. Important changes in public servlets since version 1.5.5 and 2.0.0	48
5. Queries and aggregations	48
5.1. Query counts	49
5.2. Metrics	49
5.3. Aggregations	50
5.3.1. Aggregation types	51
6. Profile import & export	56
6.1. Importing profiles	56
6.1.1. Import API	57
6.2. Exporting profiles	58
6.2.1. Export API	58
6.3. Configuration in details	60
7. Consent management	61
7.1. Consent API	61
7.1.1. Profiles with consents	62
7.1.2. Consent type definitions	63
7.1.3. Creating / update a visitor consent	63
7.1.4. How it works (internally)	65
8. Privacy management	66
8.1. Setting up access to the privacy endpoint	66
8.2. Anonymizing a profile	67
8.3. Downloading profile data	67
8.4. Deleting a profile	67
8.5. Related	67
9. Cluster setup	68
9.1. Cluster setup	68
10. Reference	68
10.1. Useful Apache Unomi URLs	68
10.2. How profile tracking works	69
10.2.1. Steps	69
10.3. Context Request Flow	70
10.4. Data Model Overview	71
10.5. Scope	72
10.5.1. Example	72
10.6. Item	72
10.6.1. Structure definition	73
10.7. Metadata	74
10.7.1. Structure definition	74
10.7.2. Example	75
10.8. MetadataItem	75
10.8.1. Structure definition	75

10.8.2. Example	76
10.9. Event	76
10.9.1. Fields	76
10.9.2. Event types	77
10.10. Profile	77
10.10.1. Structure definition	78
10.10.2. Example	78
10.11. Persona	81
10.11.1. Structure definition	81
10.11.2. Example	81
10.12. Consent	82
10.12.1. Structure definition	82
10.12.2. Example	82
10.13. Session	82
10.13.1. Structure definition	83
10.13.2. Example	83
10.14. Segment	85
10.14.1. Structure definition	85
10.14.2. Example	86
10.15. Condition	87
10.15.1. Structure definition	88
10.15.2. Example	88
10.16. Rule	89
10.16.1. Structure definition	91
10.16.2. Example	91
10.17. Action	92
10.17.1. Structure definition	92
10.17.2. Example	93
10.18. List	93
10.18.1. Structure definition	93
10.18.2. Example	93
10.19. Goal	94
10.19.1. Structure definition	94
10.19.2. Example	94
10.20. Campaign	95
10.20.1. Structure definition	95
10.20.2. Example	96
10.21. Scoring plan	97
10.21.1. Structure definition	98
10.21.2. Example	98
10.22. Built-in Event types	99
10.22.1. Login event type	99
10.22.2. View event type	101
10.22.3. Form event type	102
10.22.4. Update properties event type	104
10.22.5. Identify event type	106
10.22.6. Session created event type	108
10.22.7. Goal event type	109
10.22.8. Modify consent event type	111
10.23. Built-in condition types	114
10.23.1. Existing condition type descriptors	115

10.24. Built-in action types	116
10.24.1. Existing action types descriptors	116
10.25. Updating Events Using the Context Servlet	117
10.25.1. Solution	117
10.25.2. Defining Rules	118
11. Integration samples	118
11.1. Samples	118
11.2. Login sample	118
11.2.1. Warning !	119
11.2.2. Installing the samples	119
11.3. Twitter sample	119
11.3.1. Overview	119
11.3.2. Interacting with the context server	120
11.3.3. Retrieving context information from Unomi using the context servlet	121
11.4. Example	122
11.4.1. HTML page	122
11.4.2. Javascript	122
11.5. Conclusion	134
11.6. Annex	134
11.7. Weather update sample	135
12. Connectors	135
12.1. Connectors	135
12.1.1. Call for contributors	136
12.2. Salesforce Connector	136
12.2.1. Getting started	136
12.2.2. Properties	138
12.2.3. Hot-deploying updates to the Salesforce connector (for developers)	138
12.2.4. Using the Salesforce Workbench for testing REST API	139
12.2.5. Setting up Streaming Push queries	139
12.2.6. Executing the unit tests	139
12.3. MailChimp Connector	140
12.3.1. Getting started	140
13. Developers	142
13.1. Building	142
13.1.1. Initial Setup	142
13.1.2. Building	142
13.1.3. Installing an ElasticSearch server	142
13.1.4. Deploying the generated binary package	143
13.1.5. Deploying into an existing Karaf server	143
13.1.6. JDK Selection on Mac OS X	144
13.1.7. Running the integration tests	145
13.1.8. Testing with an example page	145
13.2. SSH Shell Commands	145
13.2.1. Using the shell	146
13.2.2. Lifecycle commands	146
13.2.3. Runtime commands	147
13.3. Writing Plugins	151
13.4. Types vs. instances	152
13.5. Plugin structure	152
13.6. Extension points	153
13.6.1. ActionType	153

13.6.2. ConditionType .....	153
13.6.3. Persona .....	153
13.6.4. PropertyMergeStrategyType .....	153
13.6.5. PropertyType .....	154
13.6.6. Rule .....	154
13.6.7. Scoring .....	154
13.6.8. Segments .....	154
13.6.9. Tag .....	154
13.6.10. ValueType .....	154
13.7. Custom plugins .....	154
13.7.1. Creating a plugin .....	154
13.7.2. Deployment and custom definition .....	155
13.7.3. Predefined segments .....	156
13.7.4. Predefined rules .....	156
13.7.5. Predefined properties .....	157
13.7.6. Predefined child conditions .....	158
13.7.7. Predefined personas .....	158
13.7.8. Custom action types .....	159
13.7.9. Custom condition types .....	161
13.8. Migration patches .....	162



# 1. QUICK START

## 1.1. FIVE MINUTES QUICKSTART

- 1) Install JDK 8 (<https://www.oracle.com/technetwork/java/javase/downloads/jdk8-downloads-2133151.html>) and make sure you set the JAVA\_HOME variable [https://docs.oracle.com/cd/E19182-01/820-7851/inst\\_cli\\_jdk\\_javahome\\_t/](https://docs.oracle.com/cd/E19182-01/820-7851/inst_cli_jdk_javahome_t/) (see our [Getting Started](#) guide for more information on JDK compatibility)
- 2) Download Elasticsearch here : <https://www.elastic.co/downloads/past-releases/elasticsearch-7-4-2> (please **make sure** you use the proper version : 7.4.2)
- 3) Uncompress it and change the [config/elasticsearch.yml](#) to include the following config :  

```
<code>cluster.name: contextElasticSearch</code>
```
- 4) Launch Elasticsearch using : [bin/elasticsearch](#)
- 5) Download Apache Unomi here : <https://unomi.apache.org/download.html>
- 6) Start it using : `./bin/karaf`
- 7) Start the Apache Unomi packages using `unomi:start` in the Apache Karaf Shell

8) Wait for startup to complete

9) Try accessing <https://localhost:9443/cxs/cluster> with username/password: `karaf/karaf` . You might get a certificate warning in your browser, just accept it despite the warning it is safe.

10) Request your first context by simply accessing : <http://localhost:8181/cxs/context.js?sessionId=1234>

11) If something goes wrong, you should check the logs in `./data/log/karaf.log`. If you get errors on ElasticSearch, make sure you are using the proper version.

Next steps:

- Connect to <http://localhost:8181> to try our some live examples (such as the web tracker)
- Trying our integration [samples page](#)
- Learning more about the [web tracker](#)

## 2. FIRST STEPS WITH APACHE UNOMI

### 2.1. 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.

#### 2.1.1. PREREQUISITES

This document assumes working knowledge of [git](#) to be able to retrieve the code for Unomi and the example. Additionally, you will require a working Java 8 or above install. Refer to <http://www.oracle.com/technetwork/java/javase/> for details on how to download and install Java SE 8 or greater.

#### JDK COMPATIBILITY

Starting with Java 9, Oracle made some big changes to the Java platform releases. This is why Apache Unomi is focused on supporting the Long Term Supported versions of the JDK, currently versions 8 and 11. We do not test with intermediate versions so they may or may not work properly. Currently the most tested version is version 8 and version 11 is also supported.

Also, as there are new licensing restrictions on JDKs provided by Oracle for production usages, Apache Unomi has also added support for OpenJDK builds. Other JDK distributions might also work but are not regularly tested so you should use them at your own risks.

#### ELASTICSEARCH COMPATIBILITY

Starting with version 1.5.0 Apache Unomi adds compatibility with ElasticSearch 7.4 . It is highly recommended to use the ElasticSearch version provided by the documentation when possible. However minor versions (7.4.x) should also work, and one version higher (7.5) will usually work. Going higher

than that is risky given the way that Elasticsearch is developed and breaking changes are introduced quite often. If in doubt, don't hesitate to check with the Apache Unomi community to get the latest information about Elasticsearch version compatibility.

## 2.1.2. RUNNING UNOMI

### START UNOMI

Start Unomi according to the [five minutes quick start](#) or by compiling using the [building 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/cxs/context.json> (which isn't very useful at this point).

You can now find an introduction page at the following location: <http://localhost:8181>

Also now that your service is up and running you can go look at the [request examples](#) to learn basic requests you can do once your server is up and running.

## 2.2. RECIPES

### 2.2.1. INTRODUCTION

In this section of the documentation we provide quick recipes focused on helping you achieve a specific result with Apache Unomi.

### 2.2.2. HOW TO READ A PROFILE

The simplest way to retrieve profile data for the current profile is to simply send a request to the `/cxs/context.json` endpoint. However you will need to send a body along with that request. Here's an example:

Here is an example that will retrieve all the session and profile properties, as well as the profile's segments and scores

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

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.

Note that it is also possible to access a profile's data through the `/cxs/profiles/` endpoint but that really should be reserved to administrative purposes. All public accesses should always use the `/cxs/context.json` endpoint for consistency and security.

### 2.2.3. HOW TO UPDATE A PROFILE FROM THE PUBLIC INTERNET

Before we get into how to update a profile directly from a request coming from the public internet, we'll quickly talk first about how NOT to do it, because we often see users using the following anti-patterns.

#### HOW NOT TO UPDATE A PROFILE FROM THE PUBLIC INTERNET

Please avoid using the `/cxs/profile` endpoint. This endpoint was initially the only way to update a profile but it has multiple issues:

- it requires authenticated access. The temptation can be great to use this endpoint because it is simple to access but the risk is that developers might include the credentials to access it in non-secure parts of code such as client-side code. Since there is no difference between this endpoint and any other administration-focused endpoints, attackers could easily re-use stolen credentials to wreak havoc on the whole platform.
- No history of profile modifications is kept: this can be a problem for multiple reasons: you might want to keep a trail of profile modifications, or even a history of profile values in case you want to understand how a profile property was modified.
- Even when protected using some kind of proxy, potentially the whole profile properties might be modified, including ones that you might not want to be overridden.



## RECOMMENDED WAYS TO UPDATE A PROFILE

Instead you can use the following solutions to update profiles:

- (Preferred) Use your own custom event(s) to send data you want to be inserted in a profile, and use rules to map the event data to the profile. This is simpler than it sounds, as usually all it requires is setting up a simple rule and you're ready to update profiles using events. This is also the safest way to update a profile because if you design your events to be as specific as possible to your needs, only the data that you specified will be copied to the profile, making sure that even in the case an attacker tries to send more data using your custom event it will simply be ignored.
- Use the protected built-in "updateProperties" event. This event is designed to be used for administrative purposes only. Again, prefer the custom events solution because as this is a protected event it will require sending the Unomi key as a request header, and as Unomi only supports a single key for the moment it could be problematic if the key is intercepted. But at least by using an event you will get the benefits of auditing and historical property modification tracing.

Let's go into more detail about the preferred way to update a profile. Let's consider the following example of a rule:

```

curl -X POST http://localhost:8181/cxs/rules \
--user karaf:karaf \
-H "Content-Type: application/json" \
-d @- <<'EOF'
{
  "metadata": {
    "id": "setContactInfo",
    "name": "Copy the received contact info to the current profile",
    "description": "Copies the contact info received in a custom event called 'contactInfoSubmitted' to
the current profile"
  },
  "raiseEventOnlyOnceForSession": false,
  "condition": {
    "type": "eventTypeCondition",
    "parameterValues": {
      "eventType": "contactInfoSubmitted"
    }
  },
  "actions": [
    {
      "type": "setPropertyAction",
      "parameterValues": {
        "propertyName": "properties(firstName)",
        "propertyValue": "eventProperty::properties(firstName)",
        "setPropertyStrategy": "alwaysSet"
      }
    },
    {
      "type": "setPropertyAction",
      "parameterValues": {
        "propertyName": "properties.lastName)",
        "propertyValue": "eventProperty::properties.lastName)",
        "setPropertyStrategy": "alwaysSet"
      }
    },
    {
      "type": "setPropertyAction",
      "parameterValues": {
        "propertyName": "properties(email)",
        "propertyValue": "eventProperty::properties(email)",
        "setPropertyStrategy": "alwaysSet"
      }
    }
  ]
}
EOF

```

What this rule does is that it listen for a custom event (events don't need any registration, you can simply start sending them to Apache Unomi whenever you like) of type 'contactInfoSubmitted' and it will search for properties called 'firstName', 'lastName' and 'email' and copy them over to the profile with corresponding property names. You could of course change any of the property names to find your needs. For example you might want to prefix the profile properties with the source of the event, such as 'mobileApp:firstName'.

You could then simply send the `contactInfoSubmitted` event using a request similar to this one:

```
curl -X POST http://localhost:8181/cxs/eventcollector \
-H "Content-Type: application/json" \
-d @- <<'EOF'
{
  "sessionId" : "1234",
  "events":[
    {
      "eventType":"contactInfoSubmitted",
      "scope": "example",
      "source":{
        "itemType": "site",
        "scope":"example",
        "itemId": "mysite"
      },
      "target":{
        "itemType":"form",
        "scope":"example",
        "itemId":"contactForm"
      },
      "properties" : {
        "firstName" : "John",
        "lastName" : "Doe",
        "email" : "john.doe@acme.com"
      }
    }
  ]
}
EOF
```

## 2.2.4. HOW TO SEARCH FOR PROFILE EVENTS

Sometimes you want to retrieve events for a known profile. You will need to provide a query in the body of the request that looks something like this (and [documentation is available in the REST API](#)) :

```
curl -X POST http://localhost:8181/cxs/events/search \
--user karaf:karaf \
-H "Content-Type: application/json" \
-d @- <<'EOF'
{ "offset" : 0,
  "limit" : 20,
  "condition" : {
    "type": "eventPropertyCondition",
    "parameterValues" : {
      "propertyName" : "profileId",
      "comparisonOperator" : "equals",
      "propertyValue" : "PROFILE_ID"
    }
  }
}
EOF
```

where PROFILE\_ID is a profile identifier. This will indeed retrieve all the events for a given profile.

## 2.2.5. HOW TO CREATE A NEW RULE

There are basically two ways to create a new rule :

- Using the REST API
- Packaging it as a predefined rule in a plugin

In both cases the JSON structure for the rule will be exactly the same, and in most scenarios it will be more interesting to use the REST API to create and manipulate rules, as they don't require any development or deployments on the Apache Unomi server.

```
curl -X POST http://localhost:8181/cxs/rules \  
--user karaf:karaf \  
-H "Content-Type: application/json" \  
-d @- <<'EOF'  
{  
  "metadata": {  
    "id": "exampleEventCopy",  
    "name": "Example Copy Event to Profile",  
    "description": "Copy event properties to profile properties"  
  },  
  "condition": {  
    "type": "eventTypeCondition",  
    "parameterValues": {  
      "eventType": "myEvent"  
    }  
  },  
  "actions": [  
    {  
      "parameterValues": {  
      },  
      "type": "allEventToProfilePropertiesAction"  
    }  
  ]  
}
```

The above rule will be executed if the incoming event is of type `myEvent` and will simply copy all the properties contained in the event to the current profile.

## 2.2.6. HOW TO SEARCH FOR PROFILES

In order to search for profiles you will have to use the `/cxs/profiles/search` endpoint that requires a Query JSON structure. Here's an example of a profile search with a Query object:

```

curl -X POST http://localhost:8181/cxs/profiles/search \
--user karaf:karaf \
-H "Content-Type: application/json" \
-d @- <<'EOF'
{
  "text" : "unomi",
  "offset" : 0,
  "limit" : 10,
  "sortBy" : "properties.lastName:asc,properties.firstName:desc",
  "condition" : {
    "type" : "booleanCondition",
    "parameterValues" : {
      "operator" : "and",
      "subConditions" : [
        {
          "type": "profilePropertyCondition",
          "parameterValues": {
            "propertyName": "properties.leadAssignedTo",
            "comparisonOperator": "exists"
          }
        },
        {
          "type": "profilePropertyCondition",
          "parameterValues": {
            "propertyName": "properties.lastName",
            "comparisonOperator": "exists"
          }
        }
      ]
    }
  }
}
EOF

```

In the above example, you search for all the profiles that have the `leadAssignedTo` and `lastName` properties and that have the `unomi` value anywhere in their profile property values. You are also specifying that you only want 10 results beginning at offset 0. The results will be also sorted in alphabetical order for the `lastName` property value, and then by reverse alphabetical order for the `firstName` property value.

As you can see, queries can be quite complex. Please remember that the more complex the more resources it will consume on the server and potentially this could affect performance.

## 2.2.7. GETTING / UPDATING CONSENTS

You can find information on how to retrieve or create/update consents in the [Consent API](#) section.

## 2.2.8. HOW TO SEND A LOGIN EVENT TO UNOMI

Tracking logins must be done carefully with Unomi. A login event is considered a "privileged" event and therefore for not be initiated from the public internet. Ideally user authentication should always be validated by a trusted third- party even if it is a well-known social platform such as Facebook or Twitter.

Basically what should NEVER be done:

1. Login to a social platform
2. Call back to the originating page
3. Send a login event to Unomi from the page originating the login in step 1

The problem with this, is that any attacker could simply directly call step 3 without any kind of security. Instead the flow should look something like this:

1. Login to a social platform
2. Call back to a special secured system that performs an server-to-server call to send the login event to Apache Unomi using the Unomi key.

For simplicity reasons, in our login example, the first method is used, but it really should never be done like this in production because of the aforementioned security issues. The second method, although a little more involved, is much preferred.

When sending a login event, you can setup a rule that can check a profile property to see if profiles can be merged on an universal identifier such as an email address.

In our login sample we provide an example of such a rule. You can find it here:

<https://github.com/apache/unomi/blob/master/samples/login-integration/src/main/resources/META-INF/cxs/rules/exampleLogin.json>

As you can see in this rule, we call an action called :

```
mergeProfilesOnPropertyAction
```

with as a parameter value the name of the property on which to perform the merge (the email). What this means is that upon successful login using an email, Unomi will look for other profiles that have the same email and merge them into a single profile. Because of the merge, this should only be done for authenticated profiles, otherwise this could be a security issue since it could be a way to load data from other profiles by merging their data !

## 2.3. REQUEST EXAMPLES

### 2.3.1. RETRIEVING YOUR FIRST CONTEXT

You can retrieve a context using curl like this :

```
curl http://localhost:8181/cxs/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).

### 2.3.2. 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/cxs/context.json?sessionId=1234
```

### 2.3.3. 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, as well as the profile's segments and scores

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

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.

### 2.3.4. 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 -X POST http://localhost:8181/cxs/context.json?sessionId=1234 \
-H "Content-Type: application/json" \
-d @- <<'EOF'
{
  "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":""
          }
        }
      }
    }
  ]
}
EOF

```

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).

### 2.3.5. 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 -X POST http://localhost:8181/cxs/eventcollector \
-H "Content-Type: application/json" \
-d @- <<'EOF'
{
  "sessionId" : "1234",
  "events":[
    {
      "eventType":"view",
      "scope": "example",
      "source":{
        "itemType": "site",
        "scope":"example",
        "itemId": "mysite"
      },
      "target":{
        "itemType":"page",
        "scope":"example",
        "itemId":"homepage",
        "properties":{
          "pageInfo":{
            "referringURL":""
          }
        }
      }
    }
  ]
}
EOF
```

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.

### 2.3.6. WHERE TO GO FROM HERE

- You can find more [useful Apache Unomi URLs](#) that can be used in the same way as the above examples.
- You may want to know integrate the provided [web tracker](#) into your web site.
- Read the [Twitter sample](#) documentation that contains a detailed example of how to integrate with Apache Unomi.

## 2.4. WEB TRACKER

This extension is providing the web tracker to start collecting visitors data on your website. The tracker is implemented as an integration of [analytics.js](#) for Unomi.

### 2.4.1. GETTING STARTED

Extension can be tested at : <http://localhost:8181/tracker/index.html>

In your page include unomiOptions and include code snippet from [snippet.min.js](#) :

```

<script type="text/javascript">
  var unomiOption = {
    scope: 'realEstateManager',
    url: 'http://localhost:8181'
  };
  window.unomiTracker || (window.unomiTracker={},function(){function
e(e){for(unomiTracker.initialize({"Apache Unomi":unomiOption});n.length>0){var
r=n.shift(),t=r.shift();unomiTracker[t]&&unomiTracker[t].apply(unomiTracker,r)}}for(var
n=[],r=["trackSubmit","trackClick","trackLink","trackForm","initialize","pageview","identify","reset","
group","track","ready","alias","debug","page","once","off","on","personalize"],t=0;t<r.length;t++){var
i=r[t];window.unomiTracker[i]=function(e){return function(){var
r=Array.prototype.slice.call(arguments);return
r.unshift(e),n.push(r),window.unomiTracker}}(i)}unomiTracker.load=function(){var
n=document.createElement("script");n.type="text/javascript",n.async=!0,n.src=unomiOption.url+"/tr
acker/unomi-
tracker.min.js",n.addEventListener?n.addEventListener("load",function(n){"function"==typeof
e&&e(n),!1):n.onreadystatechange=function(){"complete"!==this.readyState&&"loaded"!==this.ready
State || e(window.event)};var
r=document.getElementsByTagName("script")[0];r.parentNode.insertBefore(n,r)},document.addEve
ntListener("DOMContentLoaded",unomiTracker.load),unomiTracker.page()});
</script>

```

`window.unomiTracker` can be used to send additional events when needed.

Check analytics.js API [here](#). All methods can be used on `unomiTracker` object, although not all event types are supported by Unomi intergation.

## 2.4.2. HOW TO CONTRIBUTE

The source code is in the folder javascript with a package.json, the file to update is `analytics.js-integration-apache-unomi.js` apply your modification in this file then use the command `yarn build` to compile a new JS file. Then you can use the test page to try your changes <http://localhost:8181/tracker/index.html>.

## 2.4.3. TRACKING PAGE VIEWS

In the initialize call, the tracker will generate an implicit page view event, which by default will be populated with the following information:

```

window.digitalData.page = window.digitalData.page || {
  path: location.pathname + location.hash,
  pageInfo: {
    pageName: document.title,
    pageID : location.pathname + location.hash,
    pagePath : location.pathname + location.hash,
    destinationURL: location.href
  }
}

```

Now if you want to provide your own custom page information for the initial page view, you can simply

do it like this:

```
unomiTracker.initialize({
  scope: 'myScope',
  url: 'http://unomi:8181', // we use an empty URL to make it relative to this page.
  initialPageProperties: {
    path: path,
    pageInfo: {
      destinationURL: location.href,
      tags: ["tag1", "tag2", "tag3"],
      categories: ["category1", "category2", "category3"]
    },
    interests: {
      "interest1": 1,
      "interest2": 2,
      "interest3": 3
    }
  }
});
```

Also note that the FIRST call to `unomiTracker.page()` will be IGNORED because of this initial page view. This is the way that the Analytics.js library handles it. So make sure you are aware of this when calling it. This is to avoid having two page views on a single call and to be compatible with old versions that did use the explicit call.

By default the script will track page views, but maybe you want to take control over this mechanism of add page views to a single page application. In order to generate a page view programmatically from Javascript you can use code similar to this :

```

<script type="text/javascript">
  // This is an example of how to provide more details page properties to the view event. This can
  be useful
  // in the case of an SPA that wants to provide information about a view that has metadata such
  as categories,
  // tags or interests.
  path = location.pathname + location.hash;
  properties = {
    path: path,
    pageInfo: {
      destinationURL: location.href,
      tags : [ "tag1", "tag2", "tag3"],
      categories : ["category1", "category2", "category3"],
    },
    interests : {
      "interest1" : 1,
      "interest2" : 2,
      "interest3" : 3
    }
  };
  console.log(properties);
  // this will trigger a second page view for the same page (the first page view is in the tracker
  snippet).
  window.unomiTracker.page(properties);
</script>

```

Here is a more detail view of what you may include in the pageInfo object :

*Table 1. PageInfo Properties*

Name	Description
pageID	A unique identifier in string format for the page. Default value : page path
pageName	A user-displayed name for the page. Default value : page title
pagePath	The path of the page, stored by Unomi. This value should be the same as the one passed in the <a href="#">page</a> property of the object passed to the unomiTracker call. Default value : page path
destinationURL	The full URL for the page view. This doesn't have to be a real existing URL it could be an internal SPA route. Default value : page URL
referringURL	The referringURL also known as the previous URL of the page/screen viewed. Default value : page referrer URL
tags	A String array of tag identifiers. For example ['tag1', 'tag2', 'tag3']

Name	Description
categories	A String array of category identifiers. For example ['category1', 'category2', 'category3']

The `interests` object is basically list of interests with "weights" attached to them. These interests will be accumulated in Apache Unomi on profiles to indicate growing interest over time for specific topics. These are freely defined and will be accepted by Apache Unomi without needing to declare them previously anywhere (the same is true for tags and categories).

#### 2.4.4. TRACKING FORM SUBMISSIONS

Using the web tracker you can also track form submissions. In order to do this a few steps are required to get a form's submission to be tracked and then its form values to be sent as events to Apache Unomi. Finally setting up a rule to react to the incoming event will help use the form values to perform any action that is desired.

Let's look at a concrete example. Before we get started you should know that this example is already available to directly test in Apache Unomi at the following URL :

<http://localhost:8181/tracker>

Simply modify the form values and click submit and it will perform all the steps we are describing below.

So here is the form we want to track :

```
<form id="testFormTracking" action="#" name="testFormTracking">
  <label for="firstName">First name</label>
  <input type="text" id="firstName" name="firstName" value="John"/>

  <label for="lastName">Last name</label>
  <input type="text" id="lastName" name="lastName" value="Doe"/>

  <label for="email">Email</label>
  <input type="email" id="email" name="email" value="johndoe@acme.com"/>

  <input type="submit" name="submitButton" value="Submit"/>
</form>
```

As you can see it's composed of three fields - `firstName`, `lastName` and `email` - as well as a submit button. In order to track it we can add directly under the following snippet :

```
<script type="text/javascript">
  window.addEventListener("load", function () {
    var form = document.getElementById('testFormTracking');
    unomiTracker.trackForm(form, 'formSubmitted', {formName: form.name});
  });
</script>
```

What this snippet does is retrieve the form using its element ID and then uses the `unomiTracker` to track form submissions. Be careful to always use in the form event name a string that starts with `form` in order for the event to be sent back to Unomi. Also the form name is also a mandatory parameter that will be passed to Unomi inside a event of type `form` under the `target.itemId` property name.

Here is an example of the event that gets sent back to Apache Unomi:

```

{
  "itemId" : "cd627012-963e-4bb5-97f0-480990b41254",
  "itemType" : "event",
  "scope" : "realEstateManager",
  "version" : 1,
  "eventType" : "form",
  "sessionId" : "aaad09aa-88c2-67bd-b106-5a47ded43ead",
  "profileId" : "48563fd0-6319-4260-8dba-ae421beba26f",
  "timeStamp" : "2018-11-23T16:32:26Z",
  "properties" : {
    "firstName" : "John",
    "lastName" : "Doe",
    "email" : "johndoe@acme.com",
    "submitButton" : "Submit"
  },
  "source" : {
    "itemId" : "/tracker/",
    "itemType" : "page",
    "scope" : "realEstateManager",
    "version" : null,
    "properties" : {
      "pageInfo" : {
        "destinationURL" :
"http://localhost:8181/tracker/?firstName=Bill&lastName=Gates&email=bgates%40microsoft.com",
        "pageID" : "/tracker/",
        "pagePath" : "/tracker/",
        "pageName" : "Apache Unomi Web Tracker Test Page",
        "referringURL" :
"http://localhost:8181/tracker/?firstName=John&lastName=Doe&email=johndoe%40acme.com"
      },
      "attributes" : [ ],
      "consentTypes" : [ ],
      "interests" : { }
    }
  },
  "target" : {
    "itemId" : "testFormTracking",
    "itemType" : "form",
    "scope" : "realEstateManager",
    "version" : null,
    "properties" : { }
  },
  "persistent" : true
}

```

You can see in this event that the form values are sent as properties of the event itself, while the form name is sent as the `target.itemId`

While setting up form tracking, it can be very useful to use the Apache Unomi Karaf SSH shell commands : `event-tail` and `event-view` to check if you are properly receiving the form submission events and that they contain the expected data. If not, check your tracking code for any errors.

Now that the data is properly sent using an event to Apache Unomi, we must still use it to perform some kind of actions. Using rules, we could do anything from updating the profile to sending the data to a

third-party server (using a custom- developed action of course). In this example we will illustrate how to update the profile.

In order to do so we will deploy a rule that will copy data coming from the event into a profile. But we will need to map the form field names to profile names, and this can be done using the [setPropertyAction](#) that's available out of the box in the Apache Unomi server.

There are two ways to register rules : either by building a custom OSGi bundle plugin or using the REST API to directly send a JSON representation of the rule to be saved. We will in this example use the CURL shell command to make a call to the REST API.



```

curl -X POST -k -u karaf:karaf https://localhost:9443/cxs/rules \
  --header "Content-Type: application/json" \
-d @- << EOF
{
  "itemId": "form-mapping-example",
  "itemType": "rule",
  "linkedItems": null,
  "raiseEventOnlyOnceForProfile": false,
  "raiseEventOnlyOnceForSession": false,
  "priority": -1,
  "metadata": {
    "id": "form-mapping-example",
    "name": "Example Form Mapping",
    "description": "An example of how to map event properties to profile properties",
    "scope": "realEstateManager",
    "tags": [],
    "enabled": true,
    "missingPlugins": false,
    "hidden": false,
    "readOnly": false
  },
  "condition": {
    "type": "formEventCondition",
    "parameterValues": {
      "formId": "testFormTracking",
      "pagePath" : "/tracker/"
    }
  },
  "actions": [
    {
      "type": "setPropertyAction",
      "parameterValues": {
        "setPropertyName": "properties(firstName)",
        "setPropertyValue": "eventProperty::properties(firstName)",
        "setPropertyStrategy": "alwaysSet"
      }
    },
    {
      "type": "setPropertyAction",
      "parameterValues": {
        "setPropertyName": "properties.lastName)",
        "setPropertyValue": "eventProperty::properties.lastName)",
        "setPropertyStrategy": "alwaysSet"
      }
    },
    {
      "type": "setPropertyAction",
      "parameterValues": {
        "setPropertyName": "properties(email)",
        "setPropertyValue": "eventProperty::properties(email)",
        "setPropertyStrategy": "alwaysSet"
      }
    }
  ]
}
EOF

```

As you can see in this request, we have a few parameters that need explaining:

- `-k` is used to accept any certificate as we are in this example using a default Apache Unomi server configuration that comes with its predefined HTTPS certificates
- `-u karaf:karaf` is the default username/password for authenticating to the REST API. To change this value you should edit the ``etc/users.properties`` file and it is required to modify this login before going to production.

Finally the rule itself should be pretty self-explanatory but there are a few important things to note :

- the `itemId` and `metadata.id` values should be the same
- the `scope` should be the same as the scope that was setup in the tracker initialization
- the `formId` parameter must have the form name value
- the `pagePath` should be the pagePath passed through the event (if you're not sure of its value, you could either using network debugging in the browser or use the `event-tail` and `event-view` commands in the Apache Unomi Karaf SSH shell).
- the `setPropertyAction` may be repeated as many times as desired to copy the values from the event to the profile. Note that the `setPropertyName` will define the property to set on the profile and the `setPropertyValue` will define where the value is coming from. In this example the name and the value are the same but that is no way a requirement. It could even be possible to using multiple `setPropertyAction` instances to copy the same event property into different profile properties.

To check if your rule is properly deployed you can use the following SSH shell command :

```
unomi:rule-view form-mapping-example
```

The parameter is the `itemId` of the rule. If you want to see all the rules deployed in the system you can use the command :

```
unomi:rule-list 1000
```

The `1000` parameter is the limit of number of objects to retrieve. As the number of rules can grow quickly in an Apache Unomi instance, it is recommended to put this value a bit high to make sure you get the full list of rules.

Once the rule is in place, try submitting the form with some values and check that the profile is properly updated. One recommend way of doing this is to use the `event-tail` command that will output something like this :

```
-----
ID                |Type      |Session                |Profile                |Timestamp
|Scope            |Persi |
-----
cef09b89-6b99-4e4f-a99c-a4159a66b42b|form      |aaad09aa-88c2-67bd-b106-
5a47ded43ead|48563fd0-6319-4260-8dba-ae421beba26f|Fri Nov 23 17:52:33 CET 2018
|realEstateManag|true |
```

You can directly see the profile that is being used, so you can then simply use the

```
unomi:profile-view 48563fd0-6319-4260-8dba-ae421beba26f
```

command to see a JSON dump of the profile and check that the form values have been properly positioned.

## 3. CONFIGURATION

### 3.1. CENTRALIZED CONFIGURATION

Apache Unomi uses a centralized configuration file that contains both system properties and configuration properties. These settings are then fed to the OSGi and other configuration files using placeholder that look something like this:

```
contextserver.publicAddress=${org.apache.unomi.cluster.public.address:-http://localhost:8181}
contextserver.internalAddress=${org.apache.unomi.cluster.internal.address:-https://localhost:9443}
```

Default values are stored in a file called `$MY_KARAF_HOME/etc/custom.system.properties` but you should never modify this file directly, as an override mechanism is available. Simply create a file called:

```
unomi.custom.system.properties
```

and put your own property values in their to override the defaults OR you can use environment variables to also override the values in the `$MY_KARAF_HOME/etc/custom.system.properties`. See the next section for more information about that.

### 3.2. CHANGING THE DEFAULT CONFIGURATION USING ENVIRONMENT VARIABLES (I.E. DOCKER CONFIGURATION)

You might want to use environment variables to change the default system configuration, especially if you intend to run Apache Unomi inside a Docker container. You can find the list of all the environment variable names in the following file:

```
https://github.com/apache/unomi/blob/master/package/src/main/resources/etc/custom.system.properties
```

If you are using Docker Container, simply pass the environment variables on the docker command line or if you are using Docker Compose you can put the environment variables in the `docker-compose.yml` file.

If you want to "save" the environment values in a file, you can use the `bin/setenv(.bat)` to setup the environment variables you want to use.

## 3.3. CHANGING THE DEFAULT CONFIGURATION USING PROPERTY FILES

If you want to change the default configuration using property files instead of environment variables, you can perform any modification you want in the `$MY_KARAF_HOME/etc/unomi.custom.system.properties` file.

By default this file does not exist and is designed to be a file that will contain only your custom modifications to the default configuration.

For example, if you want to change the HTTP ports that the server is listening on, you will need to create the following lines in the `$MY_KARAF_HOME/etc/unomi.custom.system.properties` (and create it if you haven't yet) file:

```
org.osgi.service.http.port.secure=9443
org.osgi.service.http.port=8181
```

If you change these ports, also make sure you adjust the following settings in the same file :

```
org.apache.unomi.cluster.public.address=http://localhost:8181
org.apache.unomi.cluster.internal.address=https://localhost:9443
```

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 lose all the data you have stored previously.

You can use the following properties for the Elasticsearch configuration

```
org.apache.unomi.elasticsearch.cluster.name=contextElasticSearch
# The elasticsearch.addresses may be a comma separated list of host names and ports such as
# hostA:9200,hostB:9200
# Note: the port number must be repeated for each host.
org.apache.unomi.elasticsearch.addresses=localhost:9200
```

## 3.4. SECURED EVENTS CONFIGURATION

Apache Unomi secures some events by default. It comes out of the box with a default configuration that you can adjust by using the centralized configuration file override in `$MY_KARAF_HOME/etc/unomi.custom.system.properties`

You can find the default configuration in the following file:

```
$MY_KARAF_HOME/etc/custom.system.properties
```

The properties start with the prefix : `org.apache.unomi.thirdparty.*` and here are the default values :

```
org.apache.unomi.thirdparty.provider1.key=${env:UNOMI_THIRDPARTY_PROVIDER1_KEY:-670c26d1cc413346c3b2fd9ce65dab41}
org.apache.unomi.thirdparty.provider1.ipAddresses=${env:UNOMI_THIRDPARTY_PROVIDER1_IPADDRESSES:-127.0.0.1,::1}
org.apache.unomi.thirdparty.provider1.allowedEvents=${env:UNOMI_THIRDPARTY_PROVIDER1_ALLOWED_EVENTS:-login,updateProperties}
```

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 HTTP request header. The "env:" part means that it will attempt to read an environment variable by that name, and if it's not found it will default to the value after the ":-" marker.

It is now also possible to use IP address ranges instead of having to list all valid IP addresses for event sources. This is very useful when working in cluster deployments where servers may be added or removed dynamically. In order to support this Apache Unomi uses a library called [IPAddress](#) that supports IP ranges and subnets. Here is an example of how to setup a range:

```
org.apache.unomi.thirdparty.provider1.ipAddresses=${env:UNOMI_THIRDPARTY_PROVIDER1_IPADDRESSES:-192.168.1.1-100,::1}
```

The above configuration will allow a range of IP addresses between 192.168.1.1 and 192.168.1.100 as well as the IPv6 loopback.

Here's another example using the subnet format:

```
org.apache.unomi.thirdparty.provider1.ipAddresses=${env:UNOMI_THIRDPARTY_PROVIDER1_IPADDRESSES:-1.2.0.0/16,::1}
```

The above configuration will allow all addresses starting with 1.2 as well as the IPv6 loopback address.

Wildcards may also be used:

```
org.apache.unomi.thirdparty.provider1.ipAddresses=${env:UNOMI_THIRDPARTY_PROVIDER1_IPADDRESSES:-1.2.*.*,::1}
```

The above configuration is exactly the same as the previous one.

More advanced ranges and subnets can be used as well, please refer to the [IPAddress](#) library documentation for details on how to format them.

If you want to add another provider you will need to add them manually in the following file (and make sure you maintain the changes when upgrading) :

```
$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,updateProperties
```

## 3.5. INSTALLING THE MAXMIND GEOIPLITE2 IP LOOKUP DATABASE

Apache Unomi 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.

## 3.6. INSTALLING GEONAMES DATABASE

Apache Unomi 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/unomi.custom.system.properties` and set `org.apache.unomi.geonames.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.

## 3.7. REST API SECURITY

The Apache Unomi 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/unomi.custom.system.properties` file :

```
org.ops4j.pax.web.ssl.keystore=${env:UNOMI_SSL_KEYSTORE:-${karaf.etc}/keystore}
org.ops4j.pax.web.ssl.password=${env:UNOMI_SSL_PASSWORD:-changeme}
org.ops4j.pax.web.ssl.keypassword=${env:UNOMI_SSL_KEYPASSWORD:-changeme}
```

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 `org.apache.unomi.security.root.password` in the `$MY_KARAF_HOME/etc/unomi.custom.system.properties` file

## 3.8. SCRIPTING SECURITY

### 3.8.1. MULTI-LAYER SCRIPTING FILTERING SYSTEM

The scripting security system is multi-layered.

For requests coming in through the `/cxs/context.json` endpoint, the following flow is used to secure incoming requests:



Conditions submitted through the context.json public endpoint are first sanitized, meaning that any scripting directly injected is removed. However, as conditions can use sub conditions that include scripting, only the first directly injected layer of scripts are removed.

The second layer is the expression filtering system, that uses an allow-listing mechanism to only accept pre-vetted expressions (through configuration and deployment on the server side). Any unrecognized expression will not be accepted.

Finally, once the script starts executing in the scripting engine, a filtering class loader will only let the script access classes that have been allowed.

This multi-layered approach makes it possible to retain a high level of security even if one layer is poorly configured or abused.

For requests coming in through the secure APIs such as rules, only the condition sanitizing step is skipped, otherwise the rest of the filtering system is the same.

### 3.8.2. SCRIPTS AND EXPRESSIONS

Apache Unomi allows using different types of expressions in the following subsystems:

- context.json filters and personalization queries
- rule conditions and actions parameters

Apache Unomi uses two integrated scripting languages to provide this functionality: OGNL and MVEL. OGNL is deprecated and is now disabled by default since 1.5.2 as it is little used (and replaced by better performing hardcoded property lookups). MVEL is more commonly used in rule actions as in the following example:

From <https://github.com/apache/unomi/blob/unomi-1.5.x/plugins/baseplugin/src/main/resources/META-INF/cxs/rules/sessionAssigned.json>:

```
{
  "metadata": {
    "id": "_ajhg9u2s5_sessionAssigned",
    "name": "Session assigned to a profile",
    "description": "Update profile visit information",
    "readOnly": true
  },
  "condition": {
    "type": "booleanCondition",
    "parameterValues": {
      "subConditions": [
        {
          "type": "eventTypeCondition",
          "parameterValues": {
            "eventId": "sessionCreated"
          }
        }
      ]
    }
  },
}
```



```

    {
      "type": "eventTypeCondition",
      "parameterValues": {
        "eventId": "sessionReassigned"
      }
    }
  ],
  "operator": "or"
}
},
"actions": [
  {
    "parameterValues": {
      "setPropertyName": "properties.previousVisit",
      "setPropertyValue": "profileProperty::lastVisit",
      "storeInSession": false
    },
    "type": "setPropertyAction"
  },
  {
    "parameterValues": {
      "setPropertyName": "properties.lastVisit",
      "setPropertyValue": "now",
      "storeInSession": false
    },
    "type": "setPropertyAction"
  },
  {
    "parameterValues": {
      "setPropertyName": "properties.nbOfVisits",
      "setPropertyValue": "script::profile.properties.?nbOfVisits != null ? (profile.properties.nbOfVisits
+ 1) : 1",
      "storeInSession": false
    },
    "type": "setPropertyAction"
  }
]
}

```

As we see in the above example, we use an MVEL script with the `setPropertyAction` to set a property value. Starting with version 1.5.2, any expression use in rules MUST be allow-listed.

OGNL was previously used wherever a parameter could be used, but MVEL could only be used with a “`script::`” prefix. Starting with version 1.5.2 OGNL will no longer be allowed and is replaced by a compatible “hardcoded” property lookup system, while MVEL requires allow-listing the scripts that are to be used.

By default, Apache Unomi comes with some built-in allowed expressions that cover all the internal uses cases.

Default allowed MVEL expressions (from <https://github.com/apache/unomi/blob/unomi-1.5.x/plugins/baseplugin/src/main/resources/META-INF/cxs/expressions/mvel.json>):

```
[
  "\\Q'systemProperties.goals.'+goalId+'TargetReached"\\E",
  "\\Q'now-'+since+'d"\\E",
  "\\Q'scores.'+scoringPlanId\\E",
  "\\QminimumDuration*1000\\E",
  "\\QmaximumDuration*1000\\E",
  "\\Qprofile.properties.?nbOfVisits != null ? (profile.properties.nbOfVisits + 1) : 1\\E",
  "\\Qsession != null ? session.size + 1 : 0\\E",
  "\\Q'properties.optimizationTest_'+event.target.itemId\\E",
  "\\Qevent.target.properties.variantId\\E",
  "\\Qprofile.properties.?systemProperties.goals.\\E[\\w\\_]*\\QReached != null ?
(profile.properties.systemProperties.goals.\\E[\\w\\_]*\\QReached) : 'now'\\E",
  "\\Qprofile.properties.?systemProperties.campaigns.\\E[\\w\\_]*\\QEngaged != null ?
(profile.properties.systemProperties.campaigns.\\E[\\w\\_]*\\QEngaged) : 'now'\\E"
]
```

If you require or are already using custom expressions, you should add a plugin to Apache Unomi to allow for this. The choice of a plugin was to make sure only system administrators and solution developers could provide such a list, avoiding the possibility to provide it through an API call or another security sensitive deployment mechanism.

There is another way of allow-listing expressions through configuration, see the “scripting configuration parameters” section below.

Procedure to add allowed expressions:

1. Create a new Apache Unomi plugin project.
2. Create a JSON file in `src/main/resources/META-INF/cxs/expressions/mvel.json` with an array of regular expressions that will contain the allowed expressions.
3. Build the project and deploy it to Apache Unomi

Warning: Do not make regular expressions too general. They should actually be as specific as possible to avoid potential injection of malicious code.

### 3.8.3. SCRIPTING EXPRESSION FILTERING CONFIGURATION PARAMETERS

Alongside with the allow-listing technology, there are new configuration parameters to control the security of the scripting engines:

```

# These parameters control the list of classes that are allowed or forbidden when executing
expressions.
org.apache.unomi.scripting.allow=${env:UNOMI_ALLOW_SCRIPTING_CLASSES:-
org.apache.unomi.api.Event,org.apache.unomi.api.Profile,org.apache.unomi.api.Session,org.apache
.unomi.api.Item,org.apache.unomi.api.CustomItem,ognl.*,java.lang.Object,java.util.Map,java.util.Ha
shMap,java.lang.Integer,org.mvel2.*}
org.apache.unomi.scripting.forbid=${env:UNOMI_FORBID_SCRIPTING_CLASSES:-}

# This parameter controls the whole expression filtering system. It is not recommended to turn it
off. The main reason to turn it off would be to check if it is interfering with something, but it should
always be active in production.
org.apache.unomi.scripting.filter.activated=${env:UNOMI_SCRIPTING_FILTER_ACTIVATED:-true}

# The following parameters control the filtering using regular expressions for each scripting sub-
system.
# The "collections" parameter tells the expression filtering system which configurations to expect. By
default only MVEL and/or OGNL are accepted values, but in the future these might be replaced by
new scripting sub-systems.
org.apache.unomi.scripting.filter.collections=${env:UNOMI_SCRIPTING_FILTER_COLLECTIONS:-
mvel,ognl}

# For each scripting sub-system, there is an allow and a forbid property that reference a .json files,
# you can either edit this files or reference your own file directly in the following config.
# Note: You can add new expressions to the "allow" file, although it is better to add them inside any
plugins you may be adding.
# This configuration is only designed to compensate for the cases where something was not
properly designed or to deal with compatibility issues.
# Just be VERY careful to make your patterns AS SPECIFIC AS POSSIBLE in order to avoid
introducing a way to abuse the expression filtering.
# Note: It is NOT recommended to change the built-in "forbid" value unless you are having issues
with its value.
# Note: mvel-allow.json contains an empty array: [], this mean nothing is allowed, so far.
# If you want to allow all expression, just remove the property
org.apache.unomi.scripting.filter.mvel.allow, but this is not recommended
# It's better to list your expressions, and provide them in the mvel-allow.json file
# example: ["\\Qsession.size + 1\\E"]
org.apache.unomi.scripting.filter.mvel.allow=${env:UNOMI_SCRIPTING_FILTER_MVEL_ALLOW:-
${karaf.etc}/mvel-allow.json}
org.apache.unomi.scripting.filter.mvel.forbid=${env:UNOMI_SCRIPTING_FILTER_MVEL_FORBID:-
${karaf.etc}/mvel-forbid.json}
org.apache.unomi.scripting.filter.ognl.allow=${env:UNOMI_SCRIPTING_FILTER_OGNL_ALLOW:-
${karaf.etc}/ognl-allow.json}
org.apache.unomi.scripting.filter.ognl.forbid=${env:UNOMI_SCRIPTING_FILTER_OGNL_FORBID:-
${karaf.etc}/ognl-forbid.json}

# This parameter controls whether OGNL scripting is allowed in expressions. Because of security
reasons it is deactivated by default. If you run into compatibility issues you could reactivate it but it
is at your own risk.
org.apache.unomi.security.properties.useOGNLScripting=${env:UNOMI_SCRIPTING_USE_OGNL:-
false}

# This parameter controls the condition sanitizing done on the ContextServlet (/cxs/context.json). If
will remove any expressions that start with "script:". It is not recommended to change this value,
unless you run into compatibility issues.
org.apache.unomi.security.personalization.sanitizeConditions=${env:UNOMI_SECURITY_SANITIZEP
ERSONALIZATIONCONDITIONS:-true}

```

### 3.8.4. GROOVY ACTIONS

Groovy actions offer the ability to define a set of actions and action types (aka action descriptors) purely from Groovy scripts defined at runtime.

Initially submitted to Unomi through a purpose-built REST API endpoint, Groovy actions are then stored in Elasticsearch. When an event matches a rule configured to execute an action, the corresponding action is fetched from Elasticsearch and executed.

#### ANATOMY OF A GROOVY ACTION

To be valid, a Groovy action must follow a particular convention which is divided in two parts:

- An annotation used to define the associated action type
- The function to be executed

Placed right before the function, the “@Action” annotation contains a set of parameter detailing how the action should be triggered.

Table 2. @Action annotation

Field name	Type	Required	Description
id	String	YES	Id of the action
actionExecutor	String	YES	Action executor contains the name of the script to call for the action type and must be prefixed with “ <b>groovy:</b> ”. The prefix indicates to Unomi which dispatcher to use when processing the action.
name	String		Action name
hidden	Boolean		Define if the action is hidden or not. It is usually used to hide objects in a UI.
parameters	List<Parameter>		The parameters of the actions, also defined by annotations

Field name	Type	Required	Description
systemTags	List<String>		A (reserved) list of tags for the associated object. This is usually populated through JSON descriptors and is not meant to be modified by end users. These tags may include values that help classify associated objects.

The function contained within the Groovy Action must be called `execute()` and its last instruction must be an integer.

This integer serves as an indication whether the values of the session and profile should be persisted. In general, the codes used are defined in the [EventService interface](#).

Each groovy actions extends by default a Base script [defined here](#)

## REST API

Actions can be deployed/updated/deleted via the dedicated [/cxs/groovyActions](#) rest endpoint.

Deploy/update an Action:

```
curl -X POST 'http://localhost:8181/cxs/groovyActions' \
--user karaf:karaf \
--form 'file=@"<file location>"'
```

A Groovy Action can be updated by submitting another Action with the same id.

Delete an Action:

```
curl -X DELETE 'http://localhost:8181/cxs/groovyActions/<Action id>' \
--user karaf:karaf
```

Note that when a groovy action is deleted by the API, the action type associated with this action will also be deleted.

## HELLO WORLD!

In this short example, we're going to create a Groovy Action that will be adding "Hello world!" to the logs whenever a new view event is triggered.

The first step consists in creating the groovy script on your filesystem, start by creating the file [hello-](#)

world.groovy:

```
@Action(id = "helloWorldGroovyAction",
        actionExecutor = "groovy:helloWorldAction")
    parameters = [@Parameter(id = "location", type = "string", multivalued = false)])
def execute() {
    logger.info("Hello {}", action.getParameterValues().get("location"))
    EventService.NO_CHANGE
}
```

As the last instruction of the script is `EventService.NO_CHANGE`, data will not be persisted.

Once the action has been created you need to submit it to Unomi (from the same folder as `hello-world.groovy`).

```
curl -X POST 'http://localhost:8181/cxs/groovyActions' \
--user karaf:karaf \
--form 'file=@hello-world.groovy'
```

Finally, register a rule to trigger execution of the groovy action:

```
curl -X POST 'http://localhost:8181/cxs/rules' \
--user karaf:karaf \
--header 'Content-Type: application/json' \
--data-raw '{
"metadata": {
    "id": "scriptGroovyActionRule",
    "name": "Test Groovy Action Rule",
    "description": "A sample rule to test Groovy actions"
},
"condition": {
    "type": "eventTypeCondition",
    "parameterValues": {
        "eventType": "view"
    }
},
"actions": [
    {
        "parameterValues": {
            "location": "world!"
        },
        "type": "helloWorldAction"
    }
]
}'
```

Note that this rule contains a “location” parameter, with the value “world!”, which is then used in the log message triggered by the action.

You can now use unomi to trigger a “view” event and see the corresponding message in the Unomi logs.

Once you're done with the Hello World! action, it can be deleted using the following command:

```
curl -X DELETE 'http://localhost:8181/cxs/groovyActions/helloWorldGroovyAction' \  
--user karaf:karaf
```

And the corresponding rule can be deleted using the following command:

```
curl -X DELETE 'http://localhost:8181/cxs/rules/scriptGroovyActionRule' \  
--user karaf:karaf
```

### 3.8.5. SCRIPTING ROADMAP

Scripting will probably undergo major changes in future versions of Apache Unomi, with the likely retirement of MVEL in favor of Groovy Actions detailed above.

These changes will not happen on maintenance versions of Apache Unomi, only in the next major version. Maintenance versions will of course maintain compatibility with existing scripting solutions.

## 3.9. AUTOMATIC PROFILE MERGING

Apache Unomi 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.

## 3.10. 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/cxs/context.js>

- <http://localhost:8181/cxs/eventcollector>

All other ports should not be accessible to the world.

For your Apache Unomi 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 Apache Unomi 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. Changing the default Karaf password can be done by modifying the [org.apache.unomi.security.root.password](#) in the `$MY_KARAF_HOME/etc/unomi.custom.system.properties` file

Or if you want to also change the user name you could modify the following file :

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

But you will also need to change the following property in the `$MY_KARAF_HOME/etc/unomi.custom.system.properties` :

```
karaf.local.user = karaf
```

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.

## 3.11. INTEGRATING WITH AN APACHE HTTP WEB SERVER

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

In your Unomi package directory, in `$MY_KARAF_HOME/etc/unomi.custom.system.properties` setup the public address for the hostname [unomi.apache.org](https://unomi.apache.org):

```
org.apache.unomi.cluster.public.address=https://unomi.apache.org/  
org.apache.unomi.cluster.internal.address=http://192.168.1.1:8181
```

and you will also need to change the cookie domain in the same file:

```
org.apache.unomi.profile.cookie.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/

```

## 3.12. 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 modify the default location settings using the centralized configuration file (`$MY_KARAF_HOME/etc/unomi.custom.system.properties`).

Here are the default values for the location settings :

```
# The following settings represent the default position that is used for localhost requests
org.apache.unomi.ip.database.location=${env:UNOMI_IP_DB:-${karaf.etc}/GeoLite2-City.mmdb}
org.apache.unomi.ip.default.countryCode=${env:UNOMI_IP_DEFAULT_COUNTRYCODE:-CH}
org.apache.unomi.ip.default.countryName=${env:UNOMI_IP_DEFAULT_COUNTRYNAME:-
Switzerland}
org.apache.unomi.ip.default.city=${env:UNOMI_IP_DEFAULT_CITY:-Geneva}
org.apache.unomi.ip.default.subdiv1=${env:UNOMI_IP_DEFAULT_SUBDIV1:-2660645}
org.apache.unomi.ip.default.subdiv2=${env:UNOMI_IP_DEFAULT_SUBDIV2:-6458783}
org.apache.unomi.ip.default.isp=${env:UNOMI_IP_DEFAULT_ISP:-Cablecom}
org.apache.unomi.ip.default.latitude=${env:UNOMI_IP_DEFAULT_LATITUDE:-46.1884341}
org.apache.unomi.ip.default.longitude=${env:UNOMI_IP_DEFAULT_LONGITUDE:-6.1282508}
```

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

## 3.13. 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. You can find the list of Apache Unomi shell commands in the "Shell commands" section of the documentation.

## 3.14. ELASTICSEARCH AUTHENTICATION AND SECURITY

With Elasticsearch 7, it's possible to secure the access to your data. (<https://www.elastic.co/guide/en/elasticsearch/reference/7.5/secure-cluster.html>)

Depending on your Elasticsearch license you may need to install Kibana and enable xpack security: <https://www.elastic.co/guide/en/elasticsearch/reference/7.5/configuring-security.html>

### 3.14.1. USER AUTHENTICATION !

If your Elasticsearch have been configured to be only accessible by authenticated users (<https://www.elastic.co/guide/en/elasticsearch/reference/7.5/setting-up-authentication.html>)

Just edit `etc/org.apache.unomi.persistence.elasticsearch.cfg` to add the following settings:

```
username=USER
password=PASSWORD
```

### 3.14.2. SSL COMMUNICATION

By default Unomi will communicate with ElasticSearch using [http](#) but you can configure your ElasticSearch server(s) to allow encrypted request using [https](#).

You can follow this documentation to enable SSL on your ElasticSearch server(s):

- [Full documentation](#)
- [Configure certificates](#)
- [Encrypt HTTP communications](#)

If your ElasticSearch is correctly configure to encrypt communications on [https](#):

Just edit `etc/org.apache.unomi.persistence.elasticsearch.cfg` to add the following settings:

```
sslEnable=true
```

By default, certificates will have to be configured on the Apache Unomi server to be able to trust the identity of the ElasticSearch server(s). But if you need to trust all certificates automatically, you can use this setting:

```
sslTrustAllCertificates=true
```

## 4. MIGRATIONS

This section contains information and steps to migrate between major Unomi versions.

### 4.1. FROM VERSION 1.4 TO 1.5

#### 4.1.1. DATA MODEL AND ELASTICSEARCH 7

Since Apache Unomi version 1.5.0 we decided to upgrade the supported ElasticSearch version to the latest 7.4.2.

To be able to do so, we had to rework the way the data was stored inside ElasticSearch.

Previously every items was stored inside the same ElasticSearch index but this is not allowed anymore in recent ElasticSearch versions.

Since Apache Unomi version 1.5.0 every type of items (see section: [Items](#)) is now stored in a dedicated separated index.

## 4.1.2. API CHANGES

To be able to handle the multiple indices the Persistence API implementation ([ElasticSearchPersistenceServiceImpl](#)) have been adapted and simplified.

The good news is that there is no API changes, the persistence API interface didn't changed.

Any custom Apache Unomi plugins or extensions should continue to work on Apache Unomi 1.5.0.

The only notable changes are located at the [ElasticSearchPersistenceServiceImpl Java class](#). This class should not be use directly, instead you should use OSGI service dependency injection using the interface [PersistenceService](#).

But if you are interested in the implementation changes:

1. The property [index.name](#) have been renamed to [index.prefix](#). Previously used for the single one index name, now every index is prefixed using this property. ([context-](#) by default)
2. We removed the property [index.names](#) originally used to create additional indices (used by the geonames DB for exemple). This property is not needed anymore because the index is automatically created by the persistence service when the mapping configuration is loaded. Example of mapping configuration file: ([geoname index mapping](#))

Because of this changes the geonames DB index name is now respecting the index naming with prefix like any other item type. Previously named: [geonames](#) is now using the index name [context-geonameentry](#) (see: [Documentation about geonames extension](#)).

## 4.1.3. MIGRATION STEPS

In order to migrate the data from ElasticSearch 5 to 7, Unomi provides a migration tool that is directly integrated.

In this migration the following is assumed:

- the ElasticSearch 5 cluster installation is referred to as the [source](#)
- the ElasticSearch 7 cluster installation is referred to as the [target](#)
- the Unomi 1.4 cluster installation is completely stopped
- the Unomi 1.5 cluster installation has never been started (just uncompressed)
- the Unomi 1.5 cluster installation has been configured to connect to the [target](#) (ElasticSearch 7) cluster

It is HIGHLY RECOMMENDED to perform a full cluster backup/snapshot of the [source](#) clusters (including ElasticSearch and Unomi clusters), and ideally to perform the migration on a restored snapshot of the [source](#) cluster. For more information on ElasticSearch 5 snapshots and restore you can find it here:

<https://www.elastic.co/guide/en/elasticsearch/reference/5.6/modules-snapshots.html>

The way the migration works is that both ElasticSearch 5 AND an ElasticSearch 7 clusters (or just single nodes) will be started at the same time, and data will be migrated from the ES 5 to the ES 7 cluster. Note that it is possible to use a single node for both the [source](#) and the [target](#) clusters to - for example - perform the migration on a single machine. If you choose to do that you will have to adjust port numbers on either the [source](#) or [target](#) cluster node. Changing ports requires a restart of the ES cluster you are modifying. In this example we will illustrate how to migrate by modifying the [source](#) cluster node ports.

So in the [source](#) 's ElasticSearch 5 [config/elasticsearch.yml](#) file we have modified the default ports to:

```
transport.tcp.port: 9310
http.port: 9210
```

Make SURE you change the ports out of the default 9200-9205 and 9300-9305 range (or whatever your cluster uses) otherwise both clusters will attempt to merge!

On the [target](#) ElasticSearch 7 cluster configuration you will need to add the following setting in the [config/elasticsearch.yml](#):

```
reindex.remote.whitelist: "localhost:9210"
```

Replace "localhost:9210" which whatever location your [source](#) cluster is available at. Restart or start your [target](#) ElasticSearch 7 cluster.

Important: Make sure you haven't started Apache Unomi before (using the [unomi:start](#) command or the autostart command line parameter) otherwise you will need to restart your Apache Unomi installation from scratch. The best way to be sure of that is to start a new Unomi install by uncompressing the archive and not launching it.

You can then start both instances of ElasticSearch 5 and ElasticSearch 7 and finally start Apache Unomi using:

```
./karaf
```

Once in the console launch the migration using the following command:

```
migrate 1.4.0
```

Note: the 1.4.0 version is the starting version. If you are starting from a different version (for example a fork), make sure that you know what official version of Apache Unomi it corresponds to and you can use the official version number as a start version for the migration.

Follow the instructions and answer the prompts. If you used the above configuration as an example you can simply use the default values.

Be careful because the first address that the tool will ask for is the [target](#) (ElasticSearch 7) cluster, not the ES 5 one.

Note that it is also possible to change the index prefix to be different from the default [context](#) value so that you could host multiple Apache Unomi instances on the same ElasticSearch cluster.

Important note: only the data that Apache Unomi manages will be migrated. If you have any other data (for example Kibana or ElasticSearch monitoring indices) they will not be migrated by this migration tool.

Once the migration has completed, you can start the new Unomi instance using:

```
unomi:start
```

You should then validate that all the data has been properly migrated. For example you could issue a command to list the profiles:

```
profile-list
```

## 4.2. IMPORTANT CHANGES IN PUBLIC SERVLETS SINCE VERSION 1.5.5 AND 2.0.0

What used to be dedicated servlets are now part of the REST endpoints. Prior to version 1.5.5 the following servlets were used:

- /context.js /context.json
- /eventcollector
- /client

In version 2.0.0 and 1.5.5 and later you have to use the new [cxs](#) REST endpoints:

- /cxs/context.js /cxs/context.json
- /cxs/eventcollector
- /cxs/client

The old servlets have been deprecated and will be removed in a future major version, so make sure to update your client applications.

## 5. QUERIES AND AGGREGATIONS

Apache Unomi contains a [query](#) endpoint that is quite powerful. It provides ways to perform queries that can quickly get result counts, apply metrics such as sum/min/max/avg or even use powerful aggregations.



In this section we will show examples of requests that may be built using this API.

## 5.1. QUERY COUNTS

Query counts are highly optimized queries that will count the number of objects that match a certain condition without retrieving the results. This can be used for example to quickly figure out how many objects will match a given condition before actually retrieving the results. It uses Elasticsearch/Lucene optimizations to avoid the cost of loading all the resulting objects.

Here's an example of a query:

```
curl -X POST http://localhost:8181/cxs/query/profile/count \
--user karaf:karaf \
-H "Content-Type: application/json" \
-d @- <<'EOF'
{
  "parameterValues": {
    "subConditions": [
      {
        "type": "profilePropertyCondition",
        "parameterValues": {
          "propertyName": "systemProperties.isAnonymousProfile",
          "comparisonOperator": "missing"
        }
      },
      {
        "type": "profilePropertyCondition",
        "parameterValues": {
          "propertyName": "properties.nbOfVisits",
          "comparisonOperator": "equals",
          "propertyValueInteger": 1
        }
      }
    ],
    "operator": "and"
  },
  "type": "booleanCondition"
}
EOF
```

The above result will return the profile count of all the profiles

Result will be something like this:

```
2084
```

## 5.2. METRICS

Metric queries make it possible to apply functions to the resulting property. The supported metrics are:

- sum
- avg
- min
- max

It is also possible to request more than one metric in a single request by concatenating them with a "/" in the URL. Here's an example request that uses the [sum](#) and [avg](#) metrics:

```
curl -X POST http://localhost:8181/cxs/query/session/profile.properties.nbOfVisits/sum/avg \
--user karaf:karaf \
-H "Content-Type: application/json" \
-d @- <<'EOF'
{
  "parameterValues": {
    "subConditions": [
      {
        "type": "sessionPropertyCondition",
        "parameterValues": {
          "comparisonOperator": "equals",
          "propertyName": "scope",
          "propertyValue": "digitall"
        }
      }
    ],
    "operator": "and"
  },
  "type": "booleanCondition"
}
EOF
```

The result will look something like this:

```
{
  "_avg":1.0,
  "_sum":9.0
}
```

## 5.3. AGGREGATIONS

Aggregations are a very powerful way to build queries in Apache Unomi that will collect and aggregate data by filtering on certain conditions.

Aggregations are composed of : - an object type and a property on which to aggregate - an aggregation setup (how data will be aggregated, by date, by numeric range, date range or ip range) - a condition (used to filter the data set that will be aggregated)

### 5.3.1. AGGREGATION TYPES

Aggregations may be of different types. They are listed here below.

#### DATE

Date aggregations make it possible to automatically generate "buckets" by time periods. For more information about the format, it is directly inherited from Elasticsearch and you may find it here: <https://www.elastic.co/guide/en/elasticsearch/reference/5.6/search-aggregations-bucket-datehistogram-aggregation.html>

Here's an example of a request to retrieve a histogram of by day of all the session that have been create by newcomers (nbOfVisits=1)

```
curl -X POST http://localhost:8181/cxs/query/session/timeStamp \  
--user karaf:karaf \  
-H "Content-Type: application/json" \  
-d @- <<'EOF'  
{  
  "aggregate": {  
    "type": "date",  
    "parameters": {  
      "interval": "1d",  
      "format": "yyyy-MM-dd"  
    }  
  },  
  "condition": {  
    "type": "booleanCondition",  
    "parameterValues": {  
      "operator": "and",  
      "subConditions": [  
        {  
          "type": "sessionPropertyCondition",  
          "parameterValues": {  
            "propertyName": "scope",  
            "comparisonOperator": "equals",  
            "propertyValue": "acme"  
          }  
        },  
        {  
          "type": "sessionPropertyCondition",  
          "parameterValues": {  
            "propertyName": "profile.properties.nbOfVisits",  
            "comparisonOperator": "equals",  
            "propertyValueInteger": 1  
          }  
        }  
      ]  
    }  
  }  
}  
EOF
```

The above request will produce a similar that looks like this:

```
{
  "_all": 8062,
  "_filtered": 4085,
  "2018-10-02": 3,
  "2018-10-03": 17,
  "2018-10-04": 18,
  "2018-10-05": 19,
  "2018-10-06": 23,
  "2018-10-07": 18,
  "2018-10-08": 20
}
```

You can see that we retrieve the count of newcomers aggregated by day.

### **DATE RANGE**

Date ranges make it possible to "bucket" dates, for example to regroup profiles by their birth date as in the example below:

```
curl -X POST http://localhost:8181/cxs/query/profile/properties.birthDate \  
--user karaf:karaf \  
-H "Content-Type: application/json" \  
-d @- <<'EOF'  
{  
  "aggregate": {  
    "property": "properties.birthDate",  
    "type": "dateRange",  
    "dateRanges": [  
      {  
        "key": "After 2009",  
        "from": "now-10y/y",  
        "to": null  
      },  
      {  
        "key": "Between 1999 and 2009",  
        "from": "now-20y/y",  
        "to": "now-10y/y"  
      },  
      {  
        "key": "Between 1989 and 1999",  
        "from": "now-30y/y",  
        "to": "now-20y/y"  
      },  
      {  
        "key": "Between 1979 and 1989",  
        "from": "now-40y/y",  
        "to": "now-30y/y"  
      },  
      {  
        "key": "Between 1969 and 1979",  
        "from": "now-50y/y",  
        "to": "now-40y/y"  
      },  
      {  
        "key": "Before 1969",  
        "from": null,  
        "to": "now-50y/y"  
      }  
    ]  
  },  
  "condition": {  
    "type": "matchAllCondition",  
    "parameterValues": {}  
  }  
}  
EOF
```

The resulting JSON response will look something like this:

```
{
  "_all":4095,
  "_filtered":4095,
  "Before 1969":2517,
  "Between 1969 and 1979":353,
  "Between 1979 and 1989":336,
  "Between 1989 and 1999":337,
  "Between 1999 and 2009":35,
  "After 2009":0,
  "_missing":517
}
```

You can find more information about the date range formats here: <https://www.elastic.co/guide/en/elasticsearch/reference/5.6/search-aggregations-bucket-daterange-aggregation.html>

## NUMERIC RANGE

Numeric ranges make it possible to use "buckets" for the various ranges you want to classify.

Here's an example of a using numeric range to regroup profiles by number of visits:

```
curl -X POST http://localhost:8181/cxs/query/profile/properties.nbOfVisits \  
--user karaf:karaf \  
-H "Content-Type: application/json" \  
-d @- <<'EOF'  
{  
  "aggregate": {  
    "property": "properties.nbOfVisits",  
    "type": "numericRange",  
    "numericRanges": [  
      {  
        "key": "Less than 5",  
        "from": null,  
        "to": 5  
      },  
      {  
        "key": "Between 5 and 10",  
        "from": 5,  
        "to": 10  
      },  
      {  
        "key": "Between 10 and 20",  
        "from": 10,  
        "to": 20  
      },  
      {  
        "key": "Between 20 and 40",  
        "from": 20,  
        "to": 40  
      },  
      {  
        "key": "Between 40 and 80",  
        "from": 40,  
        "to": 80  
      },  
      {  
        "key": "Greater than 100",  
        "from": 100,  
        "to": null  
      }  
    ]  
  },  
  "condition": {  
    "type": "matchAllCondition",  
    "parameterValues": {}  
  }  
}  
EOF
```

This will produce an output that looks like this:

```
{
  "_all":4095,
  "_filtered":4095,
  "Less than 5":3855,
  "Between 5 and 10":233,
  "Between 10 and 20":7,
  "Between 20 and 40":0,
  "Between 40 and 80":0,
  "Greater than 100":0
}
```

## 6. PROFILE IMPORT & EXPORT

The profile import and export feature in Apache Unomi is based on configurations and consumes or produces CSV files that contain profiles to be imported and exported.

### 6.1. IMPORTING PROFILES

Only `ftp`, `sftp`, `ftps` and `file` are supported in the source path. For example:

```
file:///tmp/?fileName=profiles.csv&move=.done&consumer.delay=25s
```

Where:

- `fileName` Can be a pattern, for example `include=*.csv` instead of `fileName=...` to consume all CSV files. By default the processed files are moved to `.camel` folder you can change it using the `move` option.
- `consumer.delay` Is the frequency of polling in milliseconds. For example, 20000 milliseconds is 20 seconds. This frequency can also be 20s. Other possible format are: 2h30m10s = 2 hours and 30 minutes and 10 seconds.

See <http://camel.apache.org/ftp.html> and <http://camel.apache.org/file2.html> to build more complex source path. Also be careful with FTP configuration as most servers no longer support plain text FTP and you should use SFTP or FTPS instead, but they are a little more difficult to configure properly. It is recommended to test the connection with an FTP client first before setting up these source paths to ensure that everything works properly. Also on FTP connections most servers require PASSIVE mode so you can specify that in the path using the `passiveMode=true` parameter.

Here are some examples of FTPS and SFTP source paths:

```
sftp://USER@HOST/PATH?password=PASSWORD&include=*.csv
ftps://USER@HOST?password=PASSWORD&fileName=profiles.csv&passiveMode=true
```

Where:



- **USER** is the user name of the SFTP/FTPS user account to login with
- **PASSWORD** is the password for the user account
- **HOST** is the host name (or IP address) of the host server that provides the SFTP / FTPS server
- **PATH** is a path to a directory inside the user's account where the file will be retrieved.

### 6.1.1. IMPORT API

Apache Unomi provides REST endpoints to manage import configurations:

```
GET /cxs/importConfiguration
GET /cxs/importConfiguration/{configId}
POST /cxs/importConfiguration
DELETE /cxs/importConfiguration/{configId}
```

This is how a oneshot import configuration looks like:

```
{
  "itemId": "importConfigId",
  "itemType": "importConfig",
  "name": "Import Config Sample",
  "description": "Sample description",
  "configType": "oneshot",    //Config type can be 'oneshot' or 'recurrent'
  "properties": {
    "mapping": {
      "email": 0,           //<Apache Unomi Property Id> : <Column Index In the CSV>
      "firstName": 2,
      ...
    }
  },
  "columnSeparator": ",",    //Character used to separate columns
  "lineSeparator": "\\n",    //Character used to separate lines (\n or \r)
  "multiValueSeparator": ";", //Character used to separate values for multivalued columns
  "multiValueDelimiter": "[]", //Character used to wrap values for multivalued columns
  "status": "SUCCESS",      //Status of last execution
  "executions": [          //(RETURN) Last executions by default only last 5 are returned
    ...
  ],
  "mergingProperty": "email", //Apache Unomi Property Id used to check duplicates
  "overwriteExistingProfiles": true, //Overwrite profiles that have duplicates
  "propertiesToOverwrite": "firstName, lastName, ...", //If last is set to true, which property to
  overwrite, 'null' means overwrite all
  "hasHeader": true,        //CSV file to import contains a header line
  "hasDeleteColumn": false //CSV file to import doesn't contain a TO DELETE column (if it
  contains, will be the last column)
}
```

A recurrent import configuration is similar to the previous one with some specific information to add to the JSON like:

```

{
  ...
  "configType": "recurrent",
  "properties": {
    "source":
"ftp://USER@SERVER[:PORT]/PATH?password=xxx&fileName=profiles.csv&move=.done&consumer.
delay=20000",
    // Only 'ftp', 'sftp', 'ftps' and 'file' are supported in the 'source' path
    // eg. file:///tmp/?fileName=profiles.csv&move=.done&consumer.delay=25s
    // 'fileName' can be a pattern eg 'include=*.csv' instead of 'fileName=...' to consume all CSV
files
    // By default the processed files are moved to '.camel' folder you can change it using the
'move' option
    // 'consumer.delay' is the frequency of polling. '20000' (in milliseconds) means 20 seconds.
Can be also '20s'
    // Other possible format are: '2h30m10s' = 2 hours and 30 minutes and 10 seconds
    "mapping": {
      ...
    }
  },
  ...
  "active": true, //If true the polling will start according to the 'source' configured above
  ...
}

```

## 6.2. EXPORTING PROFILES

Only `ftp`, `sftp`, `ftps` and `file` are supported in the source path. For example:

```

file:///tmp/?fileName=profiles-export- $\{date:now:yyyyMMddHHmm\}$ .csv&fileExist=Append)
sftp://USER@HOST/PATH?password=PASSWORD&binary=true&fileName=profiles-export-
 $\{date:now:yyyyMMddHHmm\}$ .csv&fileExist=Append
ftps://USER@HOST?password=PASSWORD&binary=true&fileName=profiles-export-
 $\{date:now:yyyyMMddHHmm\}$ .csv&fileExist=Append&passiveMode=true

```

As you can see in the examples above, you can inject variables in the produced file name  `$\{date:now:yyyyMMddHHmm\}$`  is the current date formatted with the pattern `yyyyMMddHHmm`. `fileExist` option put as `Append` will tell the file writer to append to the same file for each execution of the export configuration. You can omit this option to write a profile per file.

See <http://camel.apache.org/ftp.html> and <http://camel.apache.org/file2.html> to build more complex destination path.

### 6.2.1. EXPORT API

Apache Unomi provides REST endpoints to manage export configurations:

```
GET /cxs/exportConfiguration
GET /cxs/exportConfiguration/{configId}
POST /cxs/exportConfiguration
DELETE /cxs/exportConfiguration/{configId}
```

This is how a oneshot export configuration looks like:

```
{
  "itemId": "exportConfigId",
  "itemType": "exportConfig",
  "name": "Export configuration sample",
  "description": "Sample description",
  "configType": "oneshot",
  "properties": {
    "period": "2m30s",
    "segment": "contacts",
    "mapping": {
      "0": "firstName",
      "1": "lastName",
      ...
    }
  },
  "columnSeparator": ",",
  "lineSeparator": "\\n",
  "multiValueSeparator": ";",
  "multiValueDelimiter": "[]",
  "status": "RUNNING",
  "executions": [
    ...
  ]
}
```

A recurrent export configuration is similar to the previous one with some specific information to add to the JSON like:

```

{
  ...
  "configType": "recurrent",
  "properties": {
    "destination": "sftp://USER@SERVER:PORT/PATH?password=XXX&fileName=profiles-export-
    ${date.now:yyyyMMddHHmm}.csv&fileExist=Append",
    "period": "2m30s", //Same as 'consumer.delay' option in the import source path
    "segment": "contacts", //Segment ID to use to collect profiles to export
    "mapping": {
      ...
    }
  },
  ...
  "active": true, //If true the configuration will start polling upon save until the user deactivate
  it
  ...
}

```

## 6.3. CONFIGURATION IN DETAILS

First configuration you need to change would be the configuration type of your import / export feature (code name router) in the `etc/unomi.custom.system.properties` file (creating it if necessary):

```

#Configuration Type values {'nobroker', 'kafka'}
org.apache.unomi.router.config.type=nobroker

```

By default the feature is configured (as above) to use no external broker, which means to handle import/export data it will use in memory queues (In the same JVM as Apache Unomi). If you are clustering Apache Unomi, most important thing to know about this type of configuration is that each Apache Unomi will handle the import/export task by itself without the help of other nodes (No Load-Distribution).

Changing this property to kafka means you have to provide the Apache Kafka configuration, and in the opposite of the nobroker option import/export data will be handled using an external broker (Apache Kafka), this will lighten the burden on the Apache Unomi machines.

You may use several Apache Kafka instance, 1 per N Apache Unomi nodes for better application scaling.

To enable using Apache Kafka you need to configure the feature as follows:

```

#Configuration Type values {'nobroker', 'kafka'}
org.apache.unomi.router.config.type=kafka

```

Uncomment and update Kafka settings to use Kafka as a broker

```
#Kafka
org.apache.unomi.router.kafka.host=localhost
org.apache.unomi.router.kafka.port=9092
org.apache.unomi.router.kafka.import.topic=import-deposit
org.apache.unomi.router.kafka.export.topic=export-deposit
org.apache.unomi.router.kafka.import.groupId=unomi-import-group
org.apache.unomi.router.kafka.export.groupId=unomi-import-group
org.apache.unomi.router.kafka.consumerCount=10
org.apache.unomi.router.kafka.autoCommit=true
```

There is couple of properties you may want to change to fit your needs, one of them is the **import.oneshot.uploadDir** which will tell Apache Unomi where to store temporarily the CSV files to import in Oneshot mode, it's a technical property to allow the choice of the convenient disk space where to store the files to be imported. It defaults to the following path under the Apache Unomi Karaf (It is recommended to change the path to a more convenient one).

```
#Import One Shot upload directory
org.apache.unomi.router.import.oneshot.uploadDir=${karaf.data}/tmp/unomi_oneshot_import_configs/
```

Next two properties are max sizes for executions history and error reports, for some reason you don't want Apache Unomi to report all the executions history and error reports generated by the executions of an import/export configuration. To change this you have to change the default values of these properties.

```
#Import/Export executions history size
org.apache.unomi.router.executionsHistory.size=5
```

```
#errors report size
org.apache.unomi.router.executions.error.report.size=200
```

Final one is about the allowed endpoints you can use when building the source or destination path, as mentioned above we can have a path of type [file](#), [ftp](#), [ftps](#), [sftp](#). You can make it less if you want to omit some endpoints (eg. you don't want to permit the use of non secure FTP).

```
#Allowed source endpoints
org.apache.unomi.router.config.allowedEndpoints=file,ftp,sftp,ftps
```

## 7. CONSENT MANAGEMENT

### 7.1. CONSENT API

Starting with Apache Unomi 1.3, a new API for consent management is now available. This API is designed to be able to store/retrieve/update visitor consents in order to comply with new privacy

regulations such as the [GDPR](#).

### 7.1.1. PROFILES WITH CONSENTS

Visitor profiles now contain a new Consent object that contains the following information:

- a scope
- a type identifier for the consent. This can be any key to reference a consent. Note that Unomi does not manage consent definitions, it only stores/retrieves consents for each profile based on this type
- a status : GRANT, DENY or REVOKED
- a status date (the date at which the status was updated)
- a revocation date, in order to comply with GDPR this is usually set at two years

Consents are stored as a sub-structure inside a profile. To retrieve the consents of a profile you can simply retrieve a profile with the following request:

```
curl -X POST http://localhost:8181/cxs/context.json?sessionId=1234 \  
-H "Content-Type: application/json" \  
-d @- <<'EOF'  
{  
  "source": {  
    "itemId":"homepage",  
    "itemType":"page",  
    "scope":"example"  
  }  
}  
EOF
```

Here is an example of a response with a Profile with a consent attached to it:

```
{
  "profileId": "18afb5e3-48cf-4f8b-96c4-854cfaadf889",
  "sessionId": "1234",
  "profileProperties": null,
  "sessionProperties": null,
  "profileSegments": null,
  "filteringResults": null,
  "personalizations": null,
  "trackedConditions": [],
  "anonymousBrowsing": false,
  "consents": {
    "example/newsletter": {
      "scope": "example",
      "typeIdentifier": "newsletter",
      "status": "GRANTED",
      "statusDate": "2018-05-22T09:27:09Z",
      "revokeDate": "2020-05-21T09:27:09Z"
    }
  }
}
```

It is of course possible to have multiple consents defined for a single visitor profile.

### 7.1.2. CONSENT TYPE DEFINITIONS

Apache Unomi does not manage consent definitions, it leaves that to an external system (for example a CMS) so that it can handle user-facing UIs to create, update, internationalize and present consent definitions to end users.

The only thing that is import to Apache Unomi to manage visitor consents is a globally unique key, that is called the consent type.

### 7.1.3. CREATING / UPDATE A VISITOR CONSENT

A new built-in event type called "modifyConsent" can be sent to Apache Unomi to update a consent for the current profile.

Here is an example of such an event:

```
{
  "events": [
    {
      "scope": "example",
      "eventType": "modifyConsent",
      "source": {
        "itemType": "page",
        "scope": "example",
        "itemId": "anItemId"
      },
      "target": {
        "itemType": "anyType",
        "scope": "example",
        "itemId": "anyItemId"
      },
      "properties": {
        "consent": {
          "typeIdentifier": "newsletter",
          "scope": "example",
          "status": "GRANTED",
          "statusDate": "2018-05-22T09:27:09.473Z",
          "revokeDate": "2020-05-21T09:27:09.473Z"
        }
      }
    }
  ]
}
```

You could send it using the following curl request:



```
curl -X POST http://localhost:8181/cxs/context.json?sessionId=1234 \  
-H "Content-Type: application/json" \  
-d @- <<'EOF'  
{  
  "source":{  
    "itemId":"homepage",  
    "itemType":"page",  
    "scope":"example"  
  },  
  "events": [  
    {  
      "scope":"example",  
      "eventType":"modifyConsent",  
      "source":{  
        "itemType":"page",  
        "scope":"example",  
        "itemId":"anItemId"  
      },  
      "target":{  
        "itemType":"anyType",  
        "scope":"example",  
        "itemId":"anyItemId"},  
      "properties":{  
        "consent":{  
          "typeIdentifier":"newsletter",  
          "scope":"example",  
          "status":"GRANTED",  
          "statusDate":"2018-05-22T09:27:09.473Z",  
          "revokeDate":"2020-05-21T09:27:09.473Z"  
        }  
      }  
    }  
  ]  
}  
EOF
```

#### 7.1.4. HOW IT WORKS (INTERNALLY)

Upon receiving this event, Apache Unomi will trigger the modifyAnyConsent rule that has the following definition:

```

{
  "metadata" : {
    "id": "modifyAnyConsent",
    "name": "Modify any consent",
    "description" : "Modify any consent and sets the consent in the profile",
    "readOnly":true
  },

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

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

```

As we can see this rule is pretty simple it will simply execute the `modifyConsentAction` that is implemented by the [ModifyConsentAction Java class](#)

This class will update the current visitor profile to add/update/revoke any consents that are included in the event.

## 8. PRIVACY MANAGEMENT

Apache Unomi provides an endpoint to manage visitor privacy. You will find in this section information about what it includes as well as how to use it.

### 8.1. SETTING UP ACCESS TO THE PRIVACY ENDPOINT

The privacy endpoint is a bit special, because despite being protected by basic authentication as the rest of the REST API is actually designed to be available to end-users.

So in effect it should usually be proxied so that public internet users can access the endpoint but the proxy should also check if the profile ID wasn't manipulated in some way.

Apache Unomi doesn't provide (for the moment) such a proxy, but basically it should do the following:

1. check for potential attack activity (could be based on IDS policies or even rate detection), and at the minimum check that the profile ID cookie seems authentic (for example by checking that it is often coming from the same IP or the same geographic location)

2. proxy to /cxs/privacy

## 8.2. ANONYMIZING A PROFILE

It is possible to anonymize a profile, meaning it will remove all "identifying" property values from the profile. Basically all properties with the tag `personalIdentifierProperties` will be purged from the profile.

Here's an example of a request to anonymize a profile:

```
curl -X POST http://localhost:8181/cxs/profiles/{profileID}/anonymize?scope=ASCOPE
```

where `{profileID}` must be replaced by the actual identifier of a profile and `ASCOPE` must be replaced by a scope identifier.

## 8.3. DOWNLOADING PROFILE DATA

It is possible to download the profile data of a user. This will only download the profile for a user using the specified ID as a cookie value.

Warning: this operation can also be sensitive so it would be better to protected with a proxy that can perform some validation on the requests to make sure no one is trying to download a profile using some kind of "guessing" of profile IDs.

```
curl -X GET http://localhost:8181/cxs/client/myprofile.[json,csv,yaml,text] \  
--cookie "context-profile-id=PROFILE-ID"
```

where `PROFILE-ID` is the profile identifier for which to download the profile.

## 8.4. DELETING A PROFILE

It is possible to delete a profile, but this works a little differently than you might expect. In all cases the data contained in the profile will be completely erased. If the `withData` optional flag is set to true, all past event and session data will also be detached from the current profile and anonymized.

```
curl -X DELETE http://localhost:8181/cxs/profiles/{profileID}?withData=false --user karaf:karaf
```

where `{profileID}` must be replaced by the actual identifier of a profile and the `withData` specifies whether the data associated with the profile must be anonymized or not

## 8.5. RELATED

You might also be interested in the [Consent API](#) section that describe how to manage profile consents.

# 9. CLUSTER SETUP

## 9.1. CLUSTER SETUP

Apache Karaf relies on Apache Karaf Cellar, which in turn uses Hazelcast to discover and configure its cluster.

You can control most of the important clustering settings through the centralized configuration file at

```
etc/unomi.custom.system.properties
```

And notably using the following properties:

```
org.apache.unomi.hazelcast.group.name=${env:UNOMI_HAZELCAST_GROUP_NAME:-cellar}
org.apache.unomi.hazelcast.group.password=${env:UNOMI_HAZELCAST_GROUP_PASSWORD:-pass}
# This list can be comma separated and use ranges such as 192.168.1.0-7,192.168.1.21
org.apache.unomi.hazelcast.tcp-ip.members=${env:UNOMI_HAZELCAST_TCPIP_MEMBERS:-
127.0.0.1}
org.apache.unomi.hazelcast.tcp-ip.interface=${env:UNOMI_HAZELCAST_TCPIP_INTERFACE:-
127.0.0.1}
org.apache.unomi.hazelcast.network.port=${env:UNOMI_HAZELCAST_NETWORK_PORT:-5701}
```

If you need more fine-grained control over the Hazelcast configuration you could also edit the following file:

```
etc/hazelcast.xml
```

Note that it would be best to keep all configuration in the centralized custom configuration, for example by adding placeholders in the hazelcast.xml file if need be and adding the properties to the centralized configuration file.

# 10. REFERENCE

## 10.1. USEFUL APACHE UNOMI URLS

In this section we will list some useful URLs that can be used to quickly access parts of Apache Unomi that can help you understand or diagnose what is going on in the system.

You can of course find more information about the REST API in the [related section](#) in the Apache Unomi website.

For these requests it can be nice to use a browser (such as Firefox) that understands JSON to make it easier to view the results as the returned JSON is not beautified (another possibility is a tool such as Postman).

Important : all URLs are relative to the private Apache Unomi URL, by default: <https://localhost:9443>

Table 3. Useful URLs

Path	Method	Description
/cxs/profiles/properties	GET	Listing deployed properties
/cxs/definitions/conditions	GET	Listing deployed conditions
/cxs/definitions/actions	GET	Listing deployed actions
/cxs/profiles/PROFILE_ID	GET	Dumping a profile in JSON
/cxs/profiles/PROFILE_ID/sessions	GET	Listing sessions for a profile
/cxs/profiles/sessions/SESSION_ID	GET	Dumping a session in JSON
/cxs/profiles/sessions/SESSION_ID /events	GET	Listing events for a session. This query can have additional such as eventTypes, q (query), offset, size, sort. See the <a href="#">related section in the REST API</a> for details.
/cxs/events/search	POST	Listing events for a profile. You will need to provide a query in the body of the request that looks something like this (and <a href="#">documentation is available in the REST API</a> ) : { "offset" : 0, "limit" : 20, "condition" : { "type": "eventPropertyCondition", "parameterValues" : { "propertyName" : "profileId", "comparisonOperator" : "equals", "propertyValue" : "PROFILE_ID" } } } where PROFILE_ID is a profile identifier. This will indeed retrieve all the events for a given profile.
/cxs/rules/statistics	GET	Get all rule execution statistics
/cxs/rules/statistics	DELETE	Reset all rule execution statistics to 0

## 10.2. HOW PROFILE TRACKING WORKS

In this section you will learn how Apache Unomi keeps track of visitors.

### 10.2.1. STEPS


1. A visitor comes to a website

2. The web server resolves a previous request session ID if it exists, or if it doesn't it create a new sessionID
3. A request to Apache Unomi's `/cxs/context.json` servlet is made passing the web server session ID as a query parameter
4. Unomi uses the sessionID and tries to load an existing session, if none is found a new session is created with the ID passed by the web server
5. If a session was found, the profile ID is extracted from the session and if it not found, Unomi looks for a cookie called `context-profile-id` to read the profileID. If no profileID is found or if the session didn't exist, a new profile ID is created by Apache Unomi
6. If the profile ID existed, the corresponding profile is loaded by Apache Unomi, otherwise a new profile is created
7. If events were passed along with the request to the `context.json` endpoint, they are processed against the profile
8. The updated profile is sent back as a response to the `context.json` request. Along with the response

It is important to note that the profileID is always server-generated. Injecting a custom cookie with a non-valid profile ID will result in failure to load the profile. Profile ID are UUIDs, which make them (pretty) safe from brute- forcing.

## 10.3. CONTEXT REQUEST FLOW

Here is an overview of how Unomi processes incoming requests to the [ContextServlet](#).

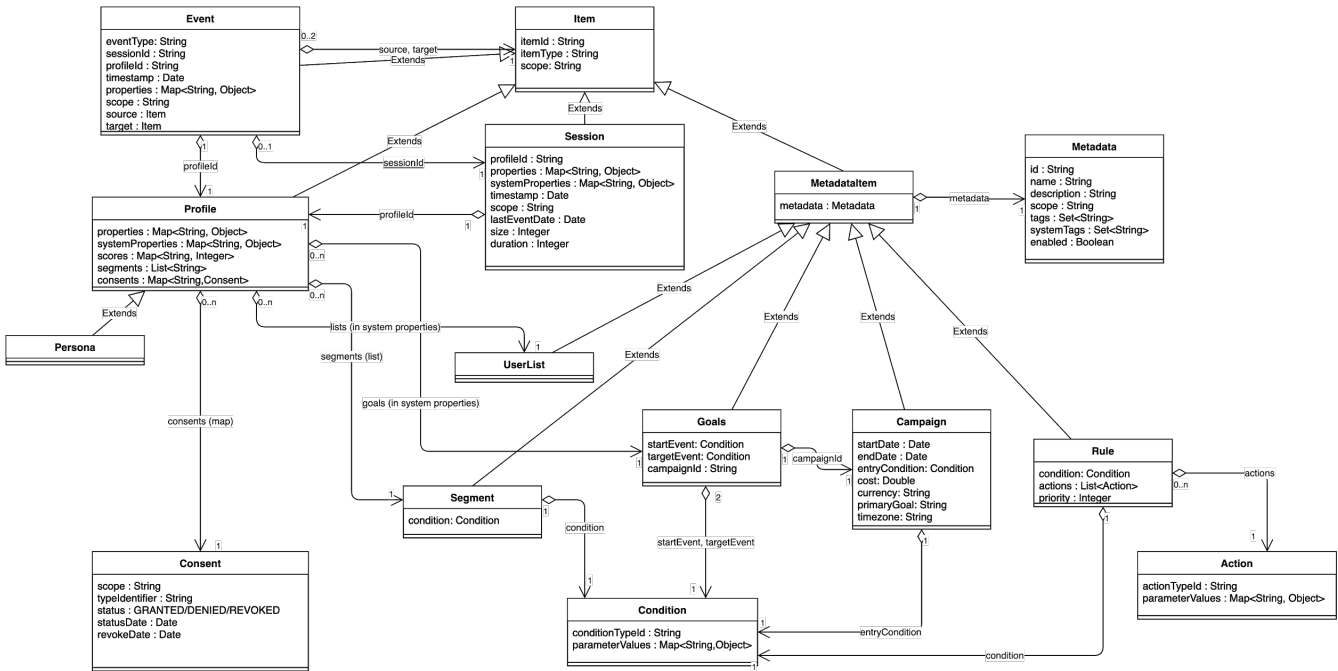


## 10.4. DATA MODEL OVERVIEW

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.

The following data model only contains the classes and properties directly related to the most important objects of Apache Unomi. There are other classes that are less central to the functionality but all the

major ones are represented in the diagram below:



We will detail many of these classes in the document below.

## 10.5. SCOPE

Scopes are not actually an object in the system, but simply unique strings that are used to "classify" objects. For example, when using scopes with a web content management system, a scope could be associated with a site identifier or even a host name.

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

### 10.5.1. EXAMPLE

In the following example, the scope uses the unique identifier of a web site called "digitall".

```

{
  ... other fields of an object type ...
  "scope": "digitall"
  ... other fields of an object type ...
}

```

## 10.6. ITEM

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.

Items are a generic object, that is common to many objects in the data model. It contains the following fields, that are inherited by other objects that inherit from it.

### 10.6.1. STRUCTURE DEFINITION

Inherits all the fields from: n/a

Field	Type	Description
itemId	String	This field contains a unique identifier (usually a UUID) that uniquely identifies the item in the whole system. It should be unique to a Unomi installation
itemType	String	A string containing the subtype of this item. Examples are : event, profile, session, ... any class that inherits from the Item class will have a unique and different itemType value.
scope	String (optional)	If present, this will contain a scope identifier. A scope is just a way to regroup objects notably for administrative purposes. For example, when integrating with a CMS a scope could be mapped to a website. The "system" scope value is reserved for values that are used internally by Apache Unomi

## 10.7. METADATA

The Metadata object is an object that contains additional information about an object. It is usually associated with an Item object (see MetadataItem below).

### 10.7.1. STRUCTURE DEFINITION

Inherits all the fields from: n/a

Field	Type	Description
id	String	This field contains a unique identifier (UUID) for the object the metadata object is attached to. It is usually a copy of the itemId field on an Item object.
name	String	A name for the associated object. Usually, this name will be displayed on the user interface
description	String (optional)	A description of the associated object. Will also usually be used in user interfaces
scope	String	The scope for the associated object.
tags	String array	A list of tags for the associated object, this list may be edited through a UI.
systemTags	String array	A (reserved) list of tags for the associated object. This is usually populated through JSON descriptors and is not meant to be modified by end users. These tags may include values such as “profileProperties” that help classify associated objects.
enabled	Boolean	Indicates whether the associated is enabled or not. For example, a rule may be disabled using this field.

Field	Type	Description
missingPlugins	Boolean	This is used for associated objects that require plugins to be deployed to work. If the plugin is not deployed, this object will not perform its function. For example if a rule is registered but the condition or actions it needs are not installed, the rule will not be used.
hidden	Boolean	Specifies whether the associated object should be visible in UIs or not
readOnly	Boolean	Specifies whether editing of the associated object should be allowed or not.

## 10.7.2. EXAMPLE

This example of a Metadata object structure was taken from a List associated object. See the [MetadataItem](#) to understand how the two fit together.

```
{
  "id": "firstListId",
  "name": "First list",
  "description": "Description of the first list.",
  "scope": "digitall",
  "tags": [],
  "systemTags": [],
  "enabled": true,
  "missingPlugins": false,
  "hidden": false,
  "readOnly": false
}
```

## 10.8. METADATAITEM

### 10.8.1. STRUCTURE DEFINITION

Inherits all the fields from: [Item](#)

Field	Type	Description
metadata	Metadata	This object contains just one field, of type Metadata as define just before this object type.

## 10.8.2. EXAMPLE

The following example is actually the definition of a [List](#) object, which is simply a [MetadataItem](#) sub-type with no additional fields. We can see here the “itemId” and “itemType” fields that come from the Item parent class and the “metadata” field that contains the object structure coming from the Metadata object type.

```
{
  "itemId": "userListId",
  "itemType": "userList",
  "metadata": {
    "id": "userListId",
    "name": "First list",
    "description": "Description of the first list.",
    "scope": "digital",
    "tags": [],
    "systemTags": [],
    "enabled": true,
    "missingPlugins": false,
    "hidden": false,
    "readOnly": false
  }
}
```

## 10.9. EVENT

Events represent something that is happening at a specific time (they are timestamped). They can be used to track visitor behavior, or even for back-channel system-to-system (as for example for a login) communication. Examples of events may include a click on a link on a web page, a login, a form submission, a page view or any other time-stamped action that needs to be tracked.

Events are persisted and immutable, and may be queried or aggregated to produce powerful reports.

Events can also be triggered as part of Unomi’s internal processes for example when a rule is triggered.

### 10.9.1. FIELDS

Inherits all the fields from: [Item](#)

Field	Type	Description
eventType	String	Contains an identifier for the field type, which may be any value as Apache Unomi does not come with strict event type definitions and accepts custom events types. The system comes with built-in event types such as “view”, “form”, “login”, “updateProperties” but additional event types may of course be used by developers integrating with Apache Unomi.
sessionId	String	The unique identifier of a Session object
profileId	String	The unique identifier of a Profile object
timestamp	Date	The precise date at which the Event was received by Unomi. This date is in the <a href="#">ISO 8601</a> format.
scope	String	(Optional, event type specific) An identifier for a scope
source	<a href="#">Item</a>	An Item that is the source of the event. For example a web site, an application name, a web page
target	<a href="#">Item</a>	An Item that is the target of the event. For example a button, a link, a file or a page
properties	Map<String,Object>	Properties for the event. These will change depending on the event type.

## 10.9.2. EVENT TYPES

Event types are completely open, and any new event type will be accepted by Apache Unomi.

Apache Unomi also comes with an extensive list of [built-in event types](#) you can find in the reference section of this manual.

## 10.10. PROFILE

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.

### 10.10.1. STRUCTURE DEFINITION

Inherits all the fields from: [Item](#)

Field name	Type	Description
properties	Map<String,Object>	All the (user-facing) properties for the profile
systemProperties	Map<String,Object>	Internal properties used to track things such as goals reached, merges with other profiles, lists the profile belongs to.
segments	String set	A set of Segment identifiers that profile is (currently) associated with
scores	Map<String,Integer>	A map of scores with the score identifier as the key and the score total value as the value.
mergedWith	String	If merged with another profile, the profile identifier to the master profile is stored here
consents	Map<String,Consent>	The consents for the profile, as a map with the consent identifier as a key and the Consent object type as a value.

### 10.10.2. EXAMPLE

In the example below, a profile for a visitor called “Bill Galileo” is detailed. A lot of user properties (such

as first name, last name, gender, job title and more) were copied over from the CMS upon initial login. The profile is also part of 4 segments (leads, contacts, gender\_male, age\_60\_70) and has a lot of different scores as well. It is also part of a list (systemProperties.lists), and has granted two consents for receiving newsletters. It has also been engaged in some goals (systemProperties.goals.\*StartReached) and completed some goals (systemProperties.goals.\*TargetReached)

Visitor profile

Properties Analytics Profiling

Filter properties by name

**Basic info**

First name: Bill  
Last name: Galileo  
Nationality: American  
Gender: Male

**Contact**

Email: bill@acme.com  
Phone number: +1-123-555-12345  
Address: Hotel Flamingo  
City: Las Vegas  
Country ISO code: United States  
ZIP Code: 89109

Cancel Save

```
{
  "itemId": "f7d1f1b9-4415-4ff1-8fee-407b109364f7",
  "itemType": "profile",
  "properties": {
    "lastName": "Galileo",
    "preferredLanguage": "en",
    "nbOfVisits": 2,
    "gender": "male",
    "jobTitle": "Vice President",
    "lastVisit": "2020-01-31T08:41:22Z",
    "j:title": "mister",
    "j:about": "<p> Lorem Ipsum dolor sit amet,consectetur adipisicing elit, sed do eiusmod tempor
incidunt ut laboreet dolore magna aliqua. Ut enim adminim veniam, quis nostrudexercitation
ullamco laboris nisi utaliquip ex ea commodo consequat.Duis aute irure dolor inreprehenderit in
coluptate velit essillum dolore eu fugiat nulla pariatur.Excepteur sint occaecat cupidatnon
proident, sunt in culpa qui officia deserunt mollit anim id estlaborum.</p> ",
    "firstName": "Bill",
    "pageViewCount": {
      "digitall": 19
    }
  },
  "emailNotificationsDisabled": "true",
  "company": "Acme Space",
}
```

```

    "j:nodename": "bill",
    "j:publicProperties":
  "j:about,j:firstName,j:function,j:gender,j:lastName,j:organization,j:picture,j:title",
    "firstVisit": "2020-01-30T21:18:12Z",
    "phoneNumber": "+1-123-555-12345",
    "countryName": "US",
    "city": "Las Vegas",
    "address": "Hotel Flamingo",
    "zipCode": "89109",
    "email": "bill@acme.com",
    "maritalStatus": "Married",
    "birthDate": "1959-08-12T23:00:00.000Z",
    "kids": 2,
    "age": 60,
    "income": 1000000,
    "facebookId": "billgalileo",
    "twitterId": "billgalileo",
    "linkedInId": "billgalileo",
    "leadAssignedTo": "Important Manager",
    "nationality": "American"
  },
  "systemProperties": {
    "mergeIdentifier": "bill",
    "lists": [
      "userListId"
    ],
  },
  "goals": {
    "viewLanguagePageGoalTargetReached": "2020-02-10T19:30:31Z",
    "downloadGoalExampleTargetReached": "2020-02-10T15:22:41Z",
    "viewLandingPageGoalStartReached": "2020-02-10T19:30:27Z",
    "downloadGoalExampleStartReached": "2020-02-10T19:30:27Z",
    "optimizationTestGoalStartReached": "2020-02-10T19:30:27Z"
  }
},
"segments": [
  "leads",
  "age_60_70",
  "gender_male",
  "contacts"
],
"scores": {
  "scoring_9": 10,
  "scoring_8": 0,
  "scoring_1": 10,
  "scoring_0": 10,
  "_s02s6220m": 0,
  "scoring_3": 10,
  "_27ir92oa2": 0,
  "scoring_2": 10,
  "scoring_5": 10,
  "scoring_4": 10,
  "scoring_7": 10,
  "scoring_6": 10,
  "_86igp9j1f": 1,
  "_ps8d573on": 0
},
"mergedWith": null,

```



```

"consents": {
  "digitall/newsletter1": {
    "scope": "digitall",
    "typeIdentifier": "newsletter1",
    "status": "GRANTED",
    "statusDate": "2019-05-15T14:47:28Z",
    "revokeDate": "2021-05-14T14:47:28Z"
  },
  "digitall/newsletter2": {
    "scope": "digitall",
    "typeIdentifier": "newsletter2",
    "status": "GRANTED",
    "statusDate": "2019-05-15T14:47:28Z",
    "revokeDate": "2021-05-14T14:47:28Z"
  }
}
}
}

```

## 10.11. PERSONA

A persona is a specialized version of a [Profile](#) object. It basically represents a "typical" profile and can be used notably to simulate personalized for a type of profiles. Usually personas are created from Profile data and then edited to represent a specific marketing persona.

### 10.11.1. STRUCTURE DEFINITION

Inherits all the fields from: [Profile](#)

There are no fields specific to a Persona.

### 10.11.2. EXAMPLE

In the following example a Persona represents a visitor from Europe, that can be used to match by location.

```

{
  "itemId": "europeanVisitor",
  "itemType": "persona",
  "properties": {
    "description": "Represents a visitor browsing from Europe",
    "firstName": "European",
    "lastName": "Visitor",
    "continent": "Europe"
  },
  "systemProperties": {},
  "segments": [],
  "scores": null,
  "mergedWith": null,
  "consents": {}
}

```

## 10.12. CONSENT

A consent represents a single instance of a consent granted/refused or revoked by a profile. A profile will contain multiple instances of consent identified by unique identifiers.

### 10.12.1. STRUCTURE DEFINITION

Inherits all the fields from: n/a

Field name	Type	Description
scope	String	The scope this consent is associated with. In the case of a website this might be the unique identifier for the site.
typeIdentifier	String	This is a unique consent type identifier, basically a unique name for the consent. Example of such types might include: “newsletter”, “personalization”, “tracking”.
status	GRANTED / DENIED / REVOKED	A copy of the profile associated with the session
statusDate	Date	The date (in ISO 8601 format) at which the current status was set
revokeDate	Date	The date (in ISO 8106 format) at which time the current status is automatically revoked.

### 10.12.2. EXAMPLE

In this example, the consent called “newsletter” was given on the “digitall” website.

```
{
  "scope": "digitall",
  "typeIdentifier": "newsletter",
  "status": "GRANTED",
  "statusDate": "2019-05-15T14:47:28Z",
  "revokeDate": "2021-05-14T14:47:28Z"
}
```

## 10.13. SESSION

A session represents a period of time during which a visitor/profile has been active. It makes it possible to gather data and then use it for reporting and further analysis by regrouping all the events that occurred during the session.

## 10.13.1. STRUCTURE DEFINITION

Inherits all the fields from: [Item](#)

Field name	Type	Description
properties	Map<String,Object>	All the properties for the session. These contain information such as the browser, operating system and device used, as well as information about the location of the visitor.
systemProperties	Map<String,Object>	Not used (empty)
profileId	String	The identifier of the profile that generated the session
profile	<a href="#">Profile</a>	A copy of the profile associated with the session
size	Integer	The number of view event types received during this session
duration	Integer	The duration of the session in milliseconds
lastEventDate	Date	The date of the last event that occurred in the session, in <a href="#">ISO 8601</a> format.

## 10.13.2. EXAMPLE

In this example the session contains a copy of the profile of the visitor. It is a visitor that has previously authenticated in a CMS and whose information was copied at the time of login from the CMS user account to the profile. You can also notice that the session contains the information coming from the browser's user agent which contains the browser type, version as well as the operating system used. The visitor's location is also resolved based on the IP address that was used to send events.

```
{
  "itemId": "4dcb5b74-6923-45ae-861a-6399ef88a209",
  "itemType": "session",
  "scope": "digital",
  "profileId": "f7d1f1b9-4415-4ff1-8fee-407b109364f7",
  "profile": {
    "itemId": "f7d1f1b9-4415-4ff1-8fee-407b109364f7",
    "itemType": "profile",
    "properties": {
      "preferredLanguage": "en",
      "nbOfVisits": 2,
      "gender": "male",
      "jobTitle": "Vice President",
      "lastVisit": "2020-01-31T08:41:22Z",
      "j:title": "mister",
```

```

    "j:about": "<p> Lorem Ipsum dolor sit amet,consectetur adipiscing elit, sed do eiusmod tempor
    incididunt ut laboreet dolore magna aliqua. Ut enim adminim veniam, quis nostrudexercitation
    ullamco laboris nisi utaliquip ex ea commodo consequat.Duis aute irure dolor inreprehenderit in
    coluptate velit esseillum dolore eu fugiat nulla pariatur.Excepteur sint occaecat cupidatatnon
    proident, sunt in culpa qui officia deserunt mollit anim id estlaborum.</p> ",
    "pageViewCount": {
      "digitall": 19
    },
    "emailNotificationsDisabled": "true",
    "company": "Acme Space",
    "j:publicProperties":
  "j:about,j:firstName,j:function,j:gender,j:lastName,j:organization,j:picture,j:title",
    "firstVisit": "2020-01-30T21:18:12Z",
    "countryName": "US",
    "city": "Las Vegas",
    "zipCode": "89109",
    "maritalStatus": "Married",
    "birthDate": "1959-08-12T23:00:00.000Z",
    "kids": 25,
    "age": 60,
    "income": 1000000,
    "leadAssignedTo": "Important Manager"
  },
  "systemProperties": {
    "mergeIdentifier": "bill",
    "lists": [
      "_xb2bcm4wl"
    ]
  },
  "segments": [
    "leads",
    "age_60_70",
    "gender_male",
    "contacts"
  ],
  "scores": {
    "scoring_9": 10,
    "scoring_8": 0,
    "scoring_1": 10,
    "scoring_0": 10,
    "_s02s6220m": 0,
    "scoring_3": 10,
    "_27ir92oa2": 0,
    "scoring_2": 10,
    "scoring_5": 10,
    "scoring_4": 10,
    "scoring_7": 10,
    "scoring_6": 10,
    "_86igp9j1f": 1,
    "_ps8d573on": 0
  },
  "mergedWith": null,
  "consents": {}
},
"properties": {
  "sessionCity": "Geneva",
  "operatingSystemFamily": "Desktop",

```

```

"userAgentNameAndVersion": "Firefox@@72.0",
"countryAndCity": "Switzerland@@Geneva@@2660645@@6458783",
"userAgent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10.15; rv:72.0) Gecko/20100101
Firefox/72.0",
"userAgentName": "Firefox",
"sessionCountryCode": "CH",
"deviceName": null,
"sessionCountryName": "Switzerland",
"referringURL": "null",
"deviceCategory": "Apple Macintosh",
"pageReferringURL": "http://localhost:8080/sites/digital/home/corporate-responsibility.html",
"userAgentVersion": "72.0",
"sessionAdminSubDiv2": 6458783,
"sessionAdminSubDiv1": 2660645,
"location": {
  "lon": 6.1282508,
  "lat": 46.1884341
},
"sessionIsp": "Cablecom",
"operatingSystemName": "Mac OS X",
"deviceBrand": "Apple"
},
"systemProperties": {},
"timeStamp": "2020-01-31T08:41:22Z",
"lastEventDate": "2020-01-31T08:53:32Z",
"size": 19,
"duration": 730317
}

```

## 10.14. SEGMENT

Segments are used to group profiles together, and are based on conditions that are executed on profiles to determine if they are part of a segment or not.

This also means that a profile may enter or leave a segment based on changes in their properties, making segments a highly dynamic concept.

### 10.14.1. STRUCTURE DEFINITION

Inherits all the fields from: [MetadataItem](#)

Field name	Type	Description
condition	<a href="#">Condition</a>	The root condition for the segment. Conditions may be composed by using built-in condition types such as <a href="#">booleanCondition</a> that can accept sub-conditions.

## 10.14.2. EXAMPLE

```
{
  "itemId": "age_20_30",
  "itemType": "segment",
  "condition": {
    "parameterValues": {
      "subConditions": [
        {
          "parameterValues": {
            "propertyName": "properties.age",
            "comparisonOperator": "greaterThanOrEqualTo",
            "propertyValueInteger": 20
          },
          "type": "profilePropertyCondition"
        },
        {
          "parameterValues": {
            "propertyName": "properties.age",
            "comparisonOperator": "lessThan",
            "propertyValueInteger": 30
          },
          "type": "profilePropertyCondition"
        }
      ],
      "operator": "and"
    },
    "type": "booleanCondition"
  },
  "metadata": {
    "id": "age_20_30",
    "name": "age_20_30",
    "description": null,
    "scope": "digitall",
    "tags": [],
    "enabled": true,
    "missingPlugins": false,
    "hidden": false,
    "readOnly": false
  }
}
```

Here is an example of a simple segment definition registered using the REST API:

```

curl -X POST http://localhost:8181/cxs/segments \
--user karaf:karaf \
-H "Content-Type: application/json" \
-d @- <<'EOF'
{
  "metadata": {
    "id": "leads",
    "name": "Leads",
    "scope": "systemscope",
    "description": "You can customize the list below by editing the leads segment.",
    "readOnly":true
  },
  "condition": {
    "type": "booleanCondition",
    "parameterValues": {
      "operator": "and",
      "subConditions": [
        {
          "type": "profilePropertyCondition",
          "parameterValues": {
            "propertyName": "properties.leadAssignedTo",
            "comparisonOperator": "exists"
          }
        }
      ]
    }
  }
}
EOF

```

For more details on the conditions and how they are structured using conditions, see the next section.

## 10.15. CONDITION

Conditions are a very useful notion inside of Apache Unomi, as they are used as the basis for multiple other objects. Conditions may be used as parts of:

- Segments
- Rules
- Queries
- Campaigns
- Goals
- Profile filters (using to search for profiles)

The result of a condition is always a boolean value of true or false.

Apache Unomi provides quite a lot of built-in condition types, including boolean types that make it possible to compose conditions using operators such as **and**, **or** or **not**. Composition is an essential element of building more complex conditions.

For a more complete list of available condition types, see the [Built-in condition types](#) reference section.

### 10.15.1. STRUCTURE DEFINITION

Inherits all the fields from: n/a

Field name	Type	Description
conditionTypeId	String	A condition type identifier is a string that contains a unique identifier for a condition type. Example condition types may include <a href="#">booleanCondition</a> , <a href="#">eventTypeCondition</a> , <a href="#">eventPropertyCondition</a> , and so on. Plugins may implement new condition types that may implement any logic that may be needed.
parameterValues	Map<String,Object>	The parameter values are simply key-value pairs that may be used to configure the condition. In the case of a <a href="#">booleanCondition</a> for example one of the parameter values will be an <a href="#">operator</a> that will contain values such as <a href="#">and</a> or <a href="#">or</a> and a second parameter value called <a href="#">subConditions</a> that contains a list of conditions to evaluate with that operator.

### 10.15.2. EXAMPLE

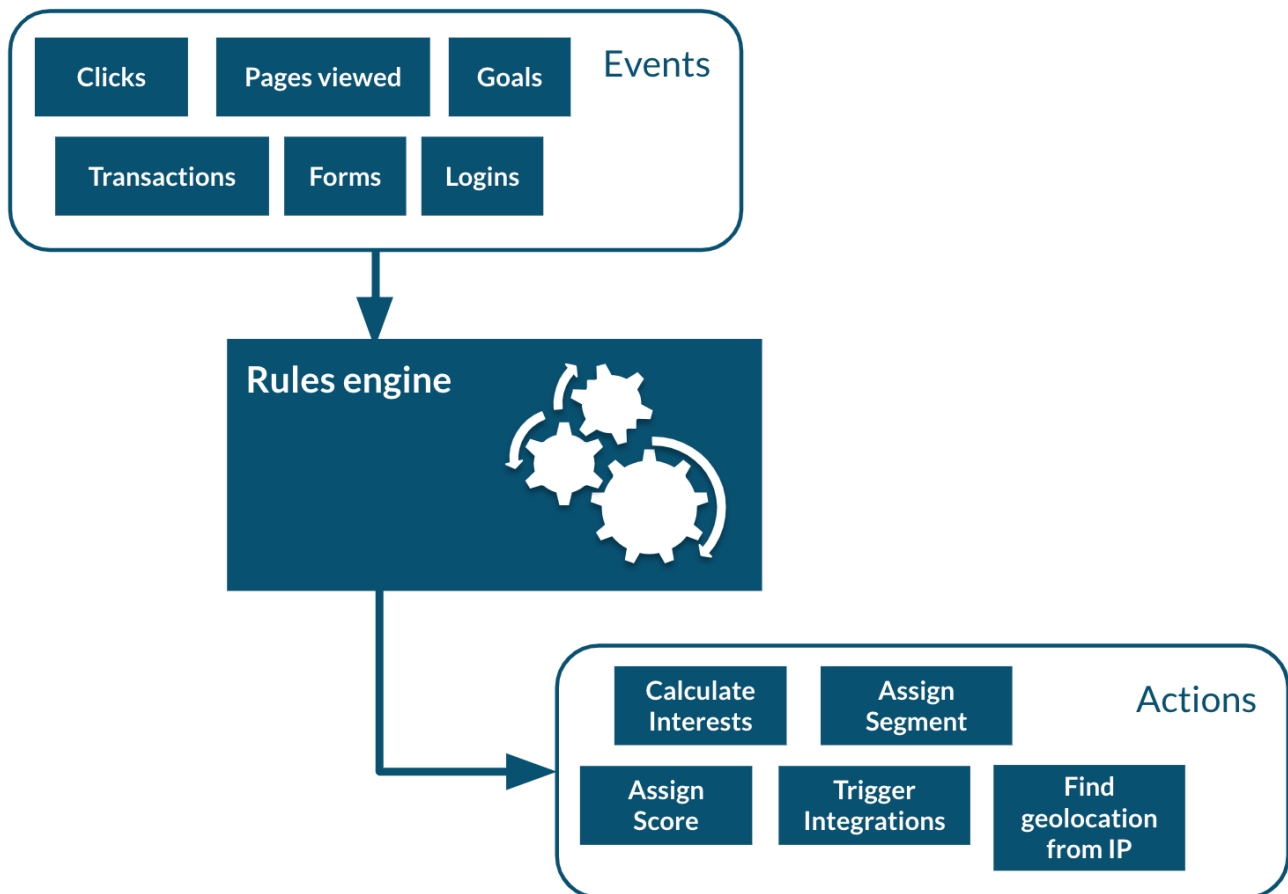
Here is an example of a complex condition:



```
{
  "condition": {
    "type": "booleanCondition",
    "parameterValues": {
      "operator": "or",
      "subConditions": [
        {
          "type": "eventTypeCondition",
          "parameterValues": {
            "eventId": "sessionCreated"
          }
        },
        {
          "type": "eventTypeCondition",
          "parameterValues": {
            "eventId": "sessionReassigned"
          }
        }
      ]
    }
  }
}
```

As we can see in the above example we use the boolean `or` condition to check if the event type is of type `sessionCreated` or `sessionReassigned`.

## 10.16. RULE



Apache Unomi has a built-in rule engine that is one of the most important components of its architecture. Every time an event is received by the server, it is evaluated against all the rules and the ones matching the incoming event will be executed. You can think of a rule as a structure that looks like this:

```

when
  conditions
then
  actions
  
```

Basically when a rule is evaluated, all the conditions in the `when` part are evaluated and if the result matches (meaning it evaluates to `true`) then the actions will be executed in sequence.

The real power of Apache Unomi comes from the fact that `conditions` and `actions` are fully pluggable and that plugins may implement new conditions and/or actions to perform any task. You can imagine conditions checking incoming event data against third-party systems or even against authentication systems, and actions actually pulling or pushing data to third-party systems.

For example the Salesforce CRM connector is simply a set of actions that pull and push data into the CRM. It is then just a matter of setting up the proper rules with the proper conditions to determine when and how the data will be pulled or pushed into the third-party system.

## 10.16.1. STRUCTURE DEFINITION

Inherits all the fields from: [MetadataItem](#)

Field name	Type	Description
condition	<a href="#">Condition</a>	The root condition for the rule. Conditions may be composed by using built-in condition types such as <a href="#">booleanCondition</a> that can accept sub-conditions.
action	<a href="#">Action</a> array	A list of <a href="#">Action</a> object that will be executed if the condition is true.
linkedItems	String array	A list of references to objects that may have generated this rule. Goals and segments dynamically generate rules to react to incoming events. It is not recommend to manipulate rules that have <code>linkedItems</code> as it may break functionality.
raiseEventOnlyOnce	Boolean	If true, the rule will only be executed once for a given event.
raiseEventOnlyOnceForProfile	Boolean	If true, the rule will only be executed once for a given profile and a matching event. Warning: this functionality has a performance impact since it looks up past events.
raiseEventOnlyOnceForSession	Boolean	If true, the rule will only be executed once for a given session and a matching event. Warning: this functionality has a performance impact since it looks up past events.
priority	Integer	The priority for the rule. The lower the priority value the higher the effective priority (they are sorted by ascending order of priority)

## 10.16.2. EXAMPLE

In this example we should the default [updateProperties](#) built-in rule that matches the [updateProperties](#) event and executes the built-in [updatePropertiesAction](#)

```

{
  "itemId": "updateProperties",
  "itemType": "rule",
  "condition": {
    "parameterValues": {},
    "type": "updatePropertiesEventCondition"
  },
  "actions": [
    {
      "parameterValues": {},
      "type": "updatePropertiesAction"
    }
  ],
  "linkedItems": null,
  "raiseEventOnlyOnceForProfile": false,
  "raiseEventOnlyOnceForSession": false,
  "priority": 0,
  "metadata": {
    "id": "updateProperties",
    "name": "Update profile/persona properties",
    "description": "Update profile/persona properties",
    "scope": "systemscope",
    "tags": [],
    "systemTags": [],
    "enabled": true,
    "missingPlugins": false,
    "hidden": false,
    "readOnly": true
  }
}

```

## 10.17. ACTION

Actions are executed by rules in a sequence, and an action is only executed once the previous action has finished executing. If an action generates an exception, it will be logged and the execution sequence will continue unless in the case of a Runtime exception (such as a `NullPointerException`).

Action use Action types that are implemented as Java classes, and as such may perform any kind of tasks that may include calling web hooks, setting profile properties, extracting data from the incoming request (such as resolving location from an IP address), or even pulling and/or pushing data to third-party systems such as a CRM server.

Apache Unomi also comes with built-in action types. You may find the list of built-in action types in the [Built-in action types](#) section.

### 10.17.1. STRUCTURE DEFINITION

Inherits all the fields from: n/a

Field name	Type	Description
actionTypeId	String	An action type identifier is a string that contains a unique identifier for a action type.
parameterValues	Map<String,Object>	The parameter values are simply key-value paris that may be used to configure the action.

### 10.17.2. EXAMPLE

In this example of an action, taking from the [form-mapping-example.json](#) rule, the [setPropertyAction](#) action is used to set the [properties.firstName](#) profile property to a value read from the event properties called [properties.firstName](#). The [setPropertyStrategy](#) is a parameter specific to this action that allows to define if existing values should be overridden or not.

```
{
  "type": "setPropertyAction",
  "parameterValues": {
    "setPropertyName": "properties(firstName)",
    "setPropertyValue": "eventProperty::properties(firstName)",
    "setPropertyStrategy": "alwaysSet"
  }
}
```

## 10.18. LIST

Lists are a “manual” way to organize profiles, whereas Segments are a dynamic way to regroup them. List objects actually only define the list in terms of name, description and other metadata but the list of members is actually not represented in the object. The profiles contain references to the lists in their “systemProperties.lists” property. This property is an array of list identifiers so in order to retrieve all the list names for a given profile, a lookup of List objects is required using the identifiers.

### 10.18.1. STRUCTURE DEFINITION

Inherits all the fields from: [MetadataItem](#)

Field name	Description
	No additional fields are present in this object type

### 10.18.2. EXAMPLE

Here’s an example of a list called “First list”, along with its description, its scope, tags, etc.. . As a List object is basically a [MetadataItem](#) sub-class it simply has all the fields defined in that parent class. Note that the List does not contain Profiles, it is Profiles that reference the Lists, not the reverse.

```

{
  "itemId": "userListId",
  "itemType": "userList",
  "metadata": {
    "id": "userListId",
    "name": "First list",
    "description": "Description of the first list.",
    "scope": "digital",
    "tags": [],
    "systemTags": [],
    "enabled": true,
    "missingPlugins": false,
    "hidden": false,
    "readOnly": false
  }
}

```

## 10.19. GOAL

A goal can be defined with two conditions: a start event condition and an target event condition. Basically the goal will be “active” when its start event condition is satisfied, and “reached” when the target event condition is true. Goals may also (optionally) be associated with Campaigns. Once a goal is “reached”, a “goal” event triggered and the profile that is currently interacting with the system will see its system properties updated to indicate which goal has been reached.

### 10.19.1. STRUCTURE DEFINITION

Inherits all the fields from: [MetadataItem](#)

Field name	Type	Description
startEvent	Condition	The condition that will be used to determine if this goal was activated by the current profile
targetEvent	Condition	The condition that will be used to determine if the current profile has reached the goal.
campaignId	String	If this goal was setup as part of a Campaign, the unique identifier for the campaign is stored in this field.

### 10.19.2. EXAMPLE

In the following example, a goal called “downloadGoalExample” is started when a new session is created (we use the “sessionCreatedEventCondition” for that) and is reached when a profile downloads a file called “ACME\_WP.pdf” (that’s what the “downloadEventCondition” means).

```

{
  "itemId": "downloadGoalExample",
  "itemType": "goal",
  "startEvent": {
    "parameterValues": {},
    "type": "sessionCreatedEventCondition"
  },
  "targetEvent": {
    "parameterValues": {
      "filePath": "/sites/digital/files/PDF/Publications/ACME_WP.pdf"
    },
    "type": "downloadEventCondition"
  },
  "campaignId": "firstCampaignExample",
  "metadata": {
    "id": "downloadGoalExample",
    "name": "downloadGoalExample",
    "description": null,
    "scope": "digital",
    "enabled": true,
    "missingPlugins": false,
    "hidden": false,
    "readOnly": false,
    "systemTags": [
      "goal",
      "downloadGoal"
    ]
  }
}

```

## 10.20. CAMPAIGN

A Campaign object represents a digital marketing campaign, along with conditions to enter the campaign and a specific duration, target and costs.

### 10.20.1. STRUCTURE DEFINITION

Inherits all the fields from: [MetadataItem](#)

Field name	Type	Description
startDate	Date	The start date of the Campaign (in ISO 8601 format)
endDate	Date	The end date of the Campaign (in ISO 8601 format)
entryCondition	<a href="#">Condition</a>	The condition that must be satisfied for a profile to become a participant in the campaign
cost	Double	An indicative cost for the campaign

Field name	Type	Description
currency	String	The currency code (3-letter) for the cost of the campaign
primaryGoal	String	A unique identifier of the primary Goal for the campaign.
timezone	String	The timezone of the campaign identified by the TZ database name (see <a href="https://en.wikipedia.org/wiki/List_of_tz_database_time_zones">https://en.wikipedia.org/wiki/List_of_tz_database_time_zones</a> )

### 10.20.2. EXAMPLE

In the following example a campaign that starts January 1st 31, 2020 at 8:38am and finished on February 29th, 2020 at the same time has the following entry condition: the session duration must be less or equal to 3000 milliseconds (3 seconds) and the profile has viewed the “about” page on the “digitall” website. The cost of the campaign is USD 1’000’000 and the timezone is Europe/Zurich. The primary goal for the campaign is the goal we should as an example in the Goal section.



```

{
  "itemId": "firstCampaignExample",
  "itemType": "campaign",
  "startDate": "2020-01-31T08:38:00Z",
  "endDate": "2020-02-29T08:38:00Z",
  "entryCondition": {
    "parameterValues": {
      "subConditions": [
        {
          "parameterValues": {
            "propertyName": "duration",
            "comparisonOperator": "lessThanOrEqualTo",
            "propertyValueInteger": 3000
          },
          "type": "sessionPropertyCondition"
        },
        {
          "parameterValues": {
            "pagePath": "/sites/digital/home/about"
          },
          "type": "pageViewEventCondition"
        }
      ],
      "operator": "and"
    },
    "type": "booleanCondition"
  },
  "cost": 1000000,
  "currency": "USD",
  "primaryGoal": "downloadGoalExample",
  "timezone": "Europe/Zurich",
  "metadata": {
    "id": "firstCampaignExample",
    "name": "firstCampaign",
    "description": "Example of a campaign",
    "scope": "digital",
    "tags": [],
    "systemTags": [
      "landing",
      "campaign"
    ],
    "enabled": true,
    "missingPlugins": false,
    "hidden": false,
    "readOnly": false
  }
}

```

## 10.21. SCORING PLAN

Scoring plans make it possible to define scores that will be tracked for profiles and use conditions to increment a score when the conditions are met. This makes it possible to then use threshold conditions on profiles when they reach a certain score.

## 10.21.1. STRUCTURE DEFINITION

Inherits all the fields from: [MetadataItem](#)

Field name	Type	Description
elements	ScoringElement array	A ScoringElement is composed of: <ul style="list-style-type: none"><li>* A Condition</li><li>* A Score increment value</li></ul> Each element defines a separate condition (tree) that will increment the defined score for this scoring plan, making it possible to have completely different conditions to augment a score

## 10.21.2. EXAMPLE

In this example a scoring plan contains a single element that will increment a score with an increment one 1 once the profile has viewed at least 3 pages (using the “hasSeenNPagesCondition” condition).

```
{
  "itemId": "viewMoreThan3PagesId",
  "itemType": "scoring",
  "elements": [
    {
      "condition": {
        "parameterValues": {
          "value": 3,
          "scope": "digital",
          "comparisonOperator": "greaterThanOrEqualTo"
        },
        "type": "hasSeenNPagesCondition"
      },
      "value": 1
    }
  ],
  "metadata": {
    "id": "viewMoreThan3PagesId",
    "name": "Viewed more than 3 pages",
    "description": null,
    "scope": "digital",
    "tags": [],
    "systemTags": [
      "st:behavioral"
    ],
    "enabled": true,
    "missingPlugins": false,
    "hidden": false,
    "readOnly": false
  }
}
```

## 10.22. BUILT-IN EVENT TYPES

Apache Unomi comes with built-in event types, which we describe below.

### 10.22.1. LOGIN EVENT TYPE

The login event type is used to signal an authentication event has been triggered. This event should be “secured”, meaning that it should not be accepted from any location, and by default Apache Unomi will only accept this event from configured “third-party” servers (identified by their IP address and a Unomi application key).

Usually, the login event will contain information passed by the authentication server and may include user properties and any additional information. Rules may be set up to copy the information from the event into the profile, but this is not done in the default set of rules provided by Apache Unomi for security reasons. You can find an example of such a rule here: <https://github.com/apache/unomi/blob/master/samples/login-integration/src/main/resources/META-INF/cxs/rules/exampleLogin.json>

#### STRUCTURE OVERVIEW

Based on the structure of the following object: Event

Field name	Value/description
eventType	login
source	Not used (null)
target	an Item representing the user that logged in
scope	the scope in which the user has authenticated
properties	Not used (empty)

#### EXAMPLE

In this case, a user has logged into a site called “digitall”, and his user information the following properties are associated with the active user..and perhaps show his visitor profile or user information.



bill



.....

Show

Remember Me

LOGIN

```
{
  "itemId": "0b8825a6-efb8-41a6-bea5-d745b33c94cb",
  "itemType": "event",
  "scope": "digitall",
  "eventType": "login",
  "sessionId": "7b8a5f17-cdb0-4c14-b676-34c1c0de0825",
  "profileId": "f7d1f1b9-4415-4ff1-8fee-407b109364f7",
  "timeStamp": "2020-01-30T21:18:28Z",
  "properties": {},
  "source": null,
  "target": {
    "itemId": "13054a95-092d-4d7b-81f5-e4656c2ebc88",
    "itemType": "cmsUser",
    "scope": null,
    "properties": {
      "j:function": "Vice President",
      "preferredLanguage": "en",
      "j:title": "mister",
      "emailNotificationsDisabled": "true",
      "j:organization": "Acme Space",
      "j:gender": "male",
      "j:nodename": "bill",
      "j:lastName": "Galileo",
      "j:publicProperties":
    "j:about,j:firstName,j:function,j:gender,j:lastName,j:organization,j:picture,j:title",
      "j:firstName": "Bill",
      "j:about": "<p> Lorem Ipsum dolor sit amet.</p> "
    }
  }
}
```

## 10.22.2. VIEW EVENT TYPE

This event is triggered when a web page is viewed by a user. Some integrators might also want to trigger it when a single-page-application screen is displayed or when a mobile application screen is displayed.

### STRUCTURE DESCRIPTION

Based on the structure of the following object: Event

Field name	Value/description
eventType	view
source	the source for the view event, could be a web site, an application name, etc...
target	the page/screen being viewed
properties	Not used (empty)

### EXAMPLE

In this case a use has visited the home page of the digitall site. As this is the first page upon login, the destination and referring URL are the same.

```

{
  "itemId": "c75f50c2-ab55-4d95-be69-cbbee180d6b",
  "itemType": "event",
  "scope": "digitall",
  "eventType": "view",
  "sessionId": "7b8a5f17-cdb0-4c14-b676-34c1c0de0825",
  "profileId": "f7d1f1b9-4415-4ff1-8fee-407b109364f7",
  "timeStamp": "2020-01-30T21:18:32Z",
  "properties": {},
  "source": {
    "itemId": "29f5fe37-28c0-48f3-966b-5353bed87308",
    "itemType": "site",
    "scope": "digitall",
    "properties": {}
  },
  "target": {
    "itemId": "f20836ab-608f-4551-a930-9796ec991340",
    "itemType": "page",
    "scope": "digitall",
    "properties": {
      "pageInfo": {
        "templateName": "home",
        "language": "en",
        "destinationURL": "http://localhost:8080/sites/digitall/home.html",
        "categories": [],
        "pageID": "f20836ab-608f-4551-a930-9796ec991340",
        "nodeType": "jnt:page",
        "pagePath": "/sites/digitall/home",
        "pageName": "Home",
        "referringURL": "http://localhost:8080/sites/digitall/home.html",
        "tags": [],
        "isContentTemplate": false
      },
      "attributes": {},
      "consentTypes": []
    }
  }
}

```

### 10.22.3. FORM EVENT TYPE

This event type is used to track form submissions. These could range from login to survey form data captured and processed in Apache Unomi using rules.

#### STRUCTURE DESCRIPTION


Based on the structure of the following object: Event

Field name	Value/description
eventType	form
source	the page/screen on which the form was submitted

Field name	Value/description
target	the form that was submitted (there could be more than one form on a page/screen)
properties	contain the data submitted in the form

## EXAMPLE

A form exists on the digitall site, and has been submitted by a visitor. In this case it was a search form that contains fields to adjust the search parameters.



Text search	Author and date	More...
<p>Search</p> <p>all these words ▾</p> <p>test</p> <p>search in:</p> <p><input checked="" type="checkbox"/> site content <input checked="" type="checkbox"/> documents</p>	<p>Author:</p> <p><input type="text"/></p> <p>Created:</p> <p>anytime ▾</p> <p>Last editor:</p> <p><input type="text"/></p> <p>Modified:</p> <p>anytime ▾</p>	<p>Site:</p> <p>Digitall ▾</p> <p>Language:</p> <p>English ▾</p> <p>Results per page:</p> <p>10 ▾</p> <p>Page:</p> <p><input type="text"/></p> <p>Select page</p> <p>include subpages</p> <p><input checked="" type="checkbox"/></p>

```

{
  "itemId": "44177ffe-b5c8-4575-a8e5-f8aa0d4ee792",
  "itemType": "event",
  "scope": "digitall",
  "eventType": "form",
  "sessionId": "be416c08-8b9b-4611-990f-3a8bf3ed4e68",
  "profileId": "bc1e1238-a9ac-4b3a-8f63-5eec205cfd5",
  "timeStamp": "2020-01-30T21:41:22Z",
  "properties": {
    "jcrMethodToCall": "get",
    "src_originSiteKey": "digitall",
    "src_terms[0].term": "test",
    "src_terms[0].applyFilter": "true",
    "src_terms[0].match": "all_words",
    "src_terms[0].fields.siteContent": "true",
    "src_terms[0].fields.tags": "true",
    "src_terms[0].fields.files": "true",
    "src_sites.values": "digitall",
    "src_sitesForReferences.values": "systemsite",
    "src_languages.values": "en"
  },
  "source": {
    "itemId": "97e14221-33dd-4608-82ae-9724d15d4f12",
    "itemType": "page",
    "scope": "digitall",
    "properties": {
      "pageInfo": {
        "templateName": "home",
        "language": "en",
        "destinationURL": "http://localhost:8080/sites/digitall/home/search-results.html",
        "categories": [],
        "pageID": "97e14221-33dd-4608-82ae-9724d15d4f12",
        "nodeType": "jnt:page",
        "pagePath": "/sites/digitall/home/search-results",
        "pageName": "Search Results",
        "referringURL": "http://localhost:8080/cms/edit/default/en/sites/digitall/home.html",
        "tags": [],
        "isContentTemplate": false
      },
      "attributes": {},
      "consentTypes": []
    }
  },
  "target": {
    "itemId": "searchForm",
    "itemType": "form",
    "scope": "digitall",
    "properties": {}
  }
}

```

#### 10.22.4. UPDATE PROPERTIES EVENT TYPE

This event is usually used by user interfaces that make it possible to modify profile properties, for example a form where a user can edit his profile properties, or a management UI to modify.



Note that this event type is a protected event type that is only accepted from configured third-party servers.

## STRUCTURE DEFINITION

Based on the structure of the following object: Event

Field name	Value/description
eventType	updateProperties
source	the screen that has triggered the update to the profile properties
target	Not used (null)
properties	{ targetId: the identifier of the profile to update targetType: “profile” if updating a profile or “persona” for personas add/update/delete: properties to be added/updated or deleted on the target profile}

## EXAMPLE

In this example, this “updateProperties” event contains properties that must be added to the targetId profile.

The screenshot shows a web interface for managing a visitor profile. On the left, there is a profile card for 'Bill Galileo' with a blue circular logo, the name 'bill', and visit dates: 'Last visit FEBRUARY 7, 2020' and 'First visit JANUARY 30, 2020'. The main area is titled 'Visitor profile' and has tabs for 'Properties', 'Analytics', and 'Profiling'. A search bar 'Filter properties by name' is at the top. Below it, the 'Basic info' section contains fields for 'First name' (Bill), 'Last name' (Galileo), 'Nationality' (American), and 'Gender' (Male). A 'Contact' section below contains fields for 'Email' (bill@acme.com), 'Phone number' (+1-123-555-12345), 'Address' (Hotel Flamingo), 'City' (Las Vegas), 'Country ISO code' (United States), and 'ZIP Code' (89109). At the bottom right, there are 'Cancel' and 'Save' buttons.

```

{
  "itemId": "d8fec330-33cb-42bc-a4e2-bb48ea7ed29b",
  "itemType": "event",
  "scope": null,
  "eventType": "updateProperties",
  "sessionId": "66e63ec9-66bc-4fac-8a8a-febcc3d6cbb7",
  "profileId": "bc1e1238-a9ac-4b3a-8f63-5eec205cfd5",
  "timeStamp": "2020-01-31T08:51:15Z",
  "properties": {
    "targetId": "f7d1f1b9-4415-4ff1-8fee-407b109364f7",
    "targetType": "profile",
    "add": {
      "properties.phoneNumber": "+1-123-555-12345",
      "properties.countryName": "US",
      "properties.city": "Las Vegas",
      "properties.address": "Hotel Flamingo",
      "properties.zipCode": "89109",
      "properties.email": "bill@acme.com"
    }
  },
  "source": {
    "itemId": "wemProfile",
    "itemType": "wemProfile",
    "scope": "digitall",
    "properties": {}
  },
  "target": null
}

```

### 10.22.5. IDENTIFY EVENT TYPE

This event type is used to add information learned about the current profile. This could be through a form that has asked the user to provide some information about himself, or it could be information sent by another system (CRM, SSO, DMP, LiveRamp or equivalent) to augment the data for the current profile.

It should be noted that, as in the case of a login event, it might be a good idea to be careful as to who and what system are allowed to send this event. Also, in order for this event to perform any modifications, an associated rule will be needed in the Unomi system to perform modifications to a profile (there is no default rule).

Event type	Available publicly	Default rule	Targeted at back-office	Can remove/update properties
identify	yes	no	no	no
updateProperties	no	yes	yes	yes

The rule of thumb is: if you need to send profile data from public system to add information to a profile you should use the identify event type and add a rule to only process the data you want to accept. If you want to add/update/delete properties in a secure manner from a known server you could use the updateProperties but you should always check first if you can't use the identify or event form event

types with specific rules as this reduces greatly the potential for misuse.

## STRUCTURE DESCRIPTION

Based on the structure of the following object: Event

Field name	Value/description
eventType	identify
source	the site/application name that triggered the identify event
target	the user information contained in the event
properties	Not used (empty)

## EXAMPLE

In this example (coming from the Apache Unomi tracker example), an event containing additional information about the user (his nickname, favorite compiler and industry) was sent to Apache Unomi.

```

{
  "itemId": "18dfd6c7-9055-4ef0-a2eb-14c1482b4544",
  "itemType": "event",
  "scope": "myScope",
  "eventType": "identify",
  "sessionId": "928d9237-fb3d-4e53-cbee-1aeb1df7f03a",
  "profileId": "temp_023ded50-bb43-4fe2-acbc-13bfa8de16de",
  "timeStamp": "2020-01-15T14:13:25Z",
  "properties": {},
  "source": {
    "itemId": "myScope",
    "itemType": "site",
    "scope": "myScope",
    "properties": {
      "page": {
        "path": "/tracker/",
        "referrer": "http://localhost:8181/",
        "search": "",
        "title": "Apache Unomi Web Tracker Test Page",
        "url": "http://localhost:8181/tracker/"
      }
    }
  },
  "target": {
    "itemId": "null",
    "itemType": "analyticsUser",
    "scope": "myScope",
    "properties": {
      "nickname": "Amazing Grace",
      "favoriteCompiler": "A-0",
      "industry": "Computer Science"
    }
  }
}

```

### 10.22.6. SESSION CREATED EVENT TYPE

The session created event is an internal event created by Apache Unomi when a new session is created. This indicates that a new visitor has interacted with a system that is using Apache Unomi to track their behavior.

#### STRUCTURE DEFINITION

Based on the structure of the following object: Event

Field name	Value/description
eventType	sessionCreated
source	Not used (null)
target	the Session item that was created with all its fields and properties

Field name	Value/description
properties	Not used (empty)

### EXAMPLE

In this example, a new session was created for a visitor coming to the digitall website. The session contains the firstVisit property. It may be augmented over time with more information including location.

```
{
  "itemId": "b3f5486f-b317-4182-9bf4-f497271e5363",
  "itemType": "event",
  "scope": "digitall",
  "eventType": "sessionCreated",
  "sessionId": "be416c08-8b9b-4611-990f-3a8bf3ed4e68",
  "profileId": "bc1e1238-a9ac-4b3a-8f63-5eec205cfdc5",
  "timeStamp": "2020-01-30T21:13:26Z",
  "properties": {},
  "source": null,
  "target": {
    "itemId": "be416c08-8b9b-4611-990f-3a8bf3ed4e68",
    "itemType": "session",
    "scope": "digitall",
    "profileId": "bc1e1238-a9ac-4b3a-8f63-5eec205cfdc5",
    "profile": {
      "itemId": "bc1e1238-a9ac-4b3a-8f63-5eec205cfdc5",
      "itemType": "profile",
      "properties": {
        "firstVisit": "2020-01-30T21:13:26Z"
      },
      "systemProperties": {},
      "segments": [],
      "scores": null,
      "mergedWith": null,
      "consents": {}
    },
    "properties": {},
    "systemProperties": {},
    "timeStamp": "2020-01-30T21:13:26Z",
    "lastEventDate": null,
    "size": 0,
    "duration": 0
  }
}
```

### 10.22.7. GOAL EVENT TYPE

A goal event is triggered when the current profile (visitor) reaches a goal.

## STRUCTURE DEFINITION

Based on the structure of the following object: Event

Field name	Value/description
eventType	goal
source	the Event that triggered the goal completion
target	the Goal item that was reached
properties	Not used (empty)

## EXAMPLE

In this example, a visitor has reached a goal by viewing a page called “sub-home” on the site “digitall” (event source). This goal event had the goal object as a target. The goal object (see Goal object later in this document) has a start event of creating a new session and a target event of a page view on the page “sub-home”.

```
{
  "itemId": "9fa70519-382d-412b-82ea-99b5989fbd0d",
  "itemType": "event",
  "scope": "digitall",
  "eventType": "goal",
  "sessionId": "42bd3fde-5fe9-4df6-8ae6-8550b8b06a7f",
  "profileId": "3ec46b2c-fbaa-42d5-99df-54199c807fc8",
  "timeStamp": "2017-05-29T23:49:16Z",
  "properties": {},
  "source": {
    "itemId": "aadcd86c-9431-43c2-bdc3-06683ac25927",
    "itemType": "event",
    "scope": "digitall",
    "eventType": "view",
    "sessionId": "42bd3fde-5fe9-4df6-8ae6-8550b8b06a7f",
    "profileId": "3ec46b2c-fbaa-42d5-99df-54199c807fc8",
    "timeStamp": "2017-05-29T23:49:16Z",
    "properties": {},
    "source": {
      "itemId": "6d5f4ae3-30c9-4561-81f3-06f82af7da1e",
      "itemType": "site",
      "scope": "digitall",
      "properties": {}
    }
  },
  "target": {
    "itemId": "67dfc299-9b74-4264-a865-aebdc3482539",
    "itemType": "page",
    "scope": "digitall",
    "properties": {
      "pageInfo": {
        "language": "en",
        "destinationURL": "https://acme.com/home/sub-home.html",
        "pageID": "67dfc299-9b74-4264-a865-aebdc3482539",
        "pagePath": "/sites/digitall/home/sub-home",

```

```

        "pageName": "sub-home",
        "referringURL": "https://acme.com/home/perso-on-profile-past-event-page.html"
    },
    "category": {},
    "attributes": {}
}
}
},
"target": {
    "itemId": "_v4ref2mxg",
    "itemType": "goal",
    "startEvent": {
        "parameterValues": {},
        "type": "sessionCreatedEventCondition"
    },
    "targetEvent": {
        "parameterValues": {
            "pagePath": "/sites/digital/home/sub-home"
        },
        "type": "pageViewEventCondition"
    },
    "campaignId": null,
    "metadata": {
        "id": "_v4ref2mxg",
        "name": "sub-home-visit",
        "description": "",
        "scope": "digital",
        "tags": [
            "pageVisitGoal"
        ],
        "enabled": true,
        "missingPlugins": false,
        "hidden": false,
        "readOnly": false
    }
}
}
}

```

### 10.22.8. MODIFY CONSENT EVENT TYPE

Consent type modification events are used to tell Unomi that consents were modified. A built-in rule will update the current profile with the consent modifications contained in the event. Consent events may be sent directly by a current profile to update their consents on the profile.

#### STRUCTURE DEFINITION

Based on the structure of the following object: Event

Field name	Value/description
eventType	modifyConsent

Field name	Value/description
source	the page that has triggered the update the consents and that contains the different consent types the current profile could grant or deny
target	The consent that was modified
properties	The consent's new value. See the Consent object type for more information.

### EXAMPLE

In this example, a user-generated a consent modification when visiting the home page, possibly by interacting with a consent form that captured his preferences. Different consent types were present on the page and he decided to GRANT the “mailchimp” consent.

## Privacy Manager

Consents

Other settings

YES
  NO

**Allow tracking**

If approved we are allowed to track the visitor

YES
  NO

**Mailchimp**

desc

YES
  NO

**Newsletter 1**

desc

YES
  NO

**Newsletter 2**

desc

YES
  NO

**Receive newsletter**

If approved we will send newsletter.

Please select consent values to close this window

{



```

"scope": "digitall",
"eventType": "modifyConsent",
"source": {
  "itemType": "page",
  "scope": "digitall",
  "itemId": "f20836ab-608f-4551-a930-9796ec991340",
  "properties": {
    "pageInfo": {
      "pageID": "f20836ab-608f-4551-a930-9796ec991340",
      "nodeType": "jnt:page",
      "pageName": "Home",
      "pagePath": "/sites/digitall/home",
      "templateName": "home",
      "destinationURL": "http://localhost:8080/sites/digitall/home.html",
      "referringURL": "http://localhost:8080/cms/render/default/en/sites/digitall/home.html",
      "language": "en",
      "categories": [],
      "tags": [],
      "isContentTemplate": false
    }
  },
  "attributes": {},
  "consentTypes": [
    {
      "typeIdentifier": "tracking",
      "activated": true,
      "title": "Allow tracking",
      "description": "If approved we are allowed to track the visitor"
    },
    {
      "typeIdentifier": "mailchimp",
      "activated": true,
      "title": "Mailchimp",
      "description": "desc"
    },
    {
      "typeIdentifier": "newsletter1",
      "activated": true,
      "title": "Newsletter 1",
      "description": "desc"
    },
    {
      "typeIdentifier": "newsletter2",
      "activated": true,
      "title": "Newsletter 2",
      "description": "desc"
    },
    {
      "typeIdentifier": "newsletter",
      "activated": true,
      "title": "Receive newsletter",
      "description": "If approved we will send newsletter."
    }
  ]
}
},
"target": {
  "itemType": "consent",

```

```

    "scope": "digital",
    "itemId": "mailchimp"
  },
  "properties": {
    "consent": {
      "scope": "digital",
      "typeIdentifier": "mailchimp",
      "status": "GRANTED",
      "statusDate": "2020-01-31T20:10:00.463Z",
      "revokeDate": "2022-01-30T20:10:00.463Z"
    }
  }
}

```

## 10.23. BUILT-IN CONDITION TYPES

Apache Unomi comes with an extensive collection of built-in condition types. Instead of detailing them one by one you will find here an overview of what a JSON condition descriptor looks like:

```

{
  "metadata": {
    "id": "booleanCondition",
    "name": "booleanCondition",
    "description": "",
    "systemTags": [
      "profileTags",
      "logical",
      "condition",
      "profileCondition",
      "eventCondition",
      "sessionCondition",
      "sourceEventCondition"
    ],
    "readOnly": true
  },
  "conditionEvaluator": "booleanConditionEvaluator",
  "queryBuilder": "booleanConditionESQueryBuilder",
  "parameters": [
    {
      "id": "operator",
      "type": "String",
      "multivalued": false,
      "defaultValue": "and"
    },
    {
      "id": "subConditions",
      "type": "Condition",
      "multivalued": true
    }
  ]
}

```

Note that condition types have two important identifiers:

- conditionEvaluator
- queryBuilder

This is because condition types can either be used to build queries or to evaluate a condition in real time. When implementing a new condition type, both implementations must be provided. Here's an example of an OSGi Blueprint registration for the above condition type descriptor:

From <https://github.com/apache/unomi/blob/master/plugins/baseplugin/src/main/resources/OSGI-INF/blueprint/blueprint.xml>

```

...
<service
interface="org.apache.unomi.persistence.elasticsearch.conditions.ConditionESQueryBuilder">
  <service-properties>
    <entry key="queryBuilderId" value="booleanConditionESQueryBuilder"/>
  </service-properties>
  <bean
class="org.apache.unomi.plugins.baseplugin.conditions.BooleanConditionESQueryBuilder"/>
</service>
...
<!-- Condition evaluators -->
<service interface="org.apache.unomi.persistence.elasticsearch.conditions.ConditionEvaluator">
  <service-properties>
    <entry key="conditionEvaluatorId" value="booleanConditionEvaluator"/>
  </service-properties>
  <bean class="org.apache.unomi.plugins.baseplugin.conditions.BooleanConditionEvaluator"/>
</service>
...

```

As you can see two Java classes are used to build a single condition type. You don't need to understand all these details in order to use condition types, but this might be interesting to know if you're interested in building your own condition type implementations. For more details on building your own custom plugins/extensions, please refer to the corresponding sections.

### 10.23.1. EXISTING CONDITION TYPE DESCRIPTORS

Here is a non-exhaustive list of condition types built into Apache Unomi. Feel free to browse the source code if you want to discover more. But the list below should get you started with the most useful conditions:

- <https://github.com/apache/unomi/tree/master/plugins/baseplugin/src/main/resources/META-INF/cxs/conditions>

Of course it is also possible to build your own custom condition types by developing custom Unomi plugins/extensions.

You will also note that some condition types can re-use a [parentCondition](#). This is a way to inherit from another condition type to make them more specific.

## 10.24. BUILT-IN ACTION TYPES

Unomi comes with quite a lot of built-in action types. Instead of detailing them one by one you will find here an overview of what an action type descriptor looks like:

```
{
  "metadata": {
    "id": "UNIQUE_IDENTIFIER_STRING",
    "name": "DISPLAYABLE_ACTION_NAME",
    "description": "DISPLAYABLE_ACTION_DESCRIPTION",
    "systemTags": [
      "profileTags",
      "event",
      "availableToEndUser",
      "allowMultipleInstances"
    ],
    "readOnly": true
  },
  "actionExecutor": "ACTION_EXECUTOR_ID",
  "parameters": [
    ... parameters specific to each action ...
  ]
}
```

The ACTION\_EXECUTOR\_ID points to a OSGi Blueprint parameter that is defined when implementing the action in a plugin. Here's an example of such a registration:

From <https://github.com/apache/unomi/blob/master/plugins/mail/src/main/resources/OSGI-INF/blueprint/blueprint.xml>

```
<bean id="sendMailActionImpl" class="org.apache.unomi.plugins.mail.actions.SendMailAction">
  <!-- ... bean properties ... -->
</bean>
<service id="sendMailAction" ref="sendMailActionImpl"
interface="org.apache.unomi.api.actions.ActionExecutor">
  <service-properties>
    <entry key="actionExecutorId" value="sendMail"/>
  </service-properties>
</service>
```

In the above example the ACTION\_EXECUTOR\_ID is `sendMail`

### 10.24.1. EXISTING ACTION TYPES DESCRIPTORS

Here is a non-exhaustive list of actions built into Apache Unomi. Feel free to browse the source code if you want to discover more. But the list below should get you started with the most useful actions:

- <https://github.com/apache/unomi/tree/master/plugins/baseplugin/src/main/resources/META-INF/cxs/actions>

- <https://github.com/apache/unomi/tree/master/plugins/request/src/main/resources/META-INF/cxs/actions>
- <https://github.com/apache/unomi/tree/master/plugins/mail/src/main/resources/META-INF/cxs/actions>

Of course it is also possible to build your own custom actions by developing custom Unomi plugins/extensions.

## 10.25. UPDATING EVENTS USING THE CONTEXT SERVLET

One of the use cases that needed to be supported by Unomi is the ability to build a user profile based on Internal System events or [Change Data Capture](#) which usually transported through internal messaging queues such as Kafka.

This can easily be achieved using the [KafkaInjector](#) module built in within Unomi.

But, as streaming system usually operates in [at-least-once](#) semantics, we need to have a way to guarantee we won't have duplicate events in the system.

### 10.25.1. SOLUTION

One of the solutions to this scenario is to have the ability to control and pass in the [eventId](#) property from outside of Unomi, using an authorized 3rd party. This way whenever an event with the same [itemId](#) will be processed once again, it won't be appended to the list of events, but will be updated.

Here is an example of a request that contains the [itemId](#)

```
curl -X POST http://localhost:8181/cxs/context.json \  
-H "Content-Type: application/json" \  
-d @- <<'EOF'  
{  
  "events":[  
    {  
      "itemId": "exampleEventId",  
      "eventType":"view",  
      "scope": "example",  
      "properties" : {  
        "firstName" : "example"  
      }  
    }  
  ]  
}  
EOF
```

Make sure to use an authorized third party using [X-Unomi-Peer](#) request headers and that the [eventType](#) is in the list of allowed events

## 10.25.2. DEFINING RULES

Another use case we support is the ability to define a rule on the above mentioned events. If we have a rule that increment a property on profile level, we would want the action to be executed only once per event id. this can be achieved by adding `"raiseEventOnlyOnce": false` to the rule definition.

```
curl -X POST http://localhost:8181/cxs/context.json \  
-H "Content-Type: application/json" \  
-d @- <<'EOF'  
{  
  "metadata": {  
    "id": "updateNumberOfOrders",  
    "name": "update number of orders on orderCreated eventType",  
    "description": "update number of orders on orderCreated eventType"  
  },  
  "raiseEventOnlyOnce": false,  
  "condition": {  
    "type": "eventTypeCondition",  
    "parameterValues": {  
      "eventId": "orderCreated"  
    }  
  },  
  "actions": [  
    {  
      "parameterValues": {  
        "setPropertyName": "properties.nbOfOrders",  
        "setPropertyValue": "script::profile.properties.?nbOfOrders != null ?  
(profile.properties.nbOfOrders + 1) : 1",  
        "storeInSession": false  
      },  
      "type": "setPropertyAction"  
    }  
  ]  
}
```

# 11. INTEGRATION SAMPLES

## 11.1. SAMPLES

Apache Unomi provides the following samples:

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

## 11.2. LOGIN SAMPLE

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

### 11.2.1. WARNING !

The example code uses client-side Javascript code to send the login event. This is only done this way for the sake of samples 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.

### 11.2.2. INSTALLING THE SAMPLES

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

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

Install the login samples using the following command:

```
bundle:install mvn:org.apache.unomi/login-integration-sample/${project.version}
```

when the bundle is successfully installed 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 samples 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
```

## 11.3. TWITTER SAMPLE

### 11.3.1. 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 SAMPLES

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 SAMPLES

In standard Karaf fashion, you will need to copy the samples 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-2.0.0-SNAPSHOT.jar ../../package/target/unomi-2.0.0-SNAPSHOT/deploy
```

## TESTING THE SAMPLES

You can now go to <http://localhost:8181/twitter/index.html> to test the samples 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.

### 11.3.2. 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.

### 11.3.3. 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>/cxs/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/cxs/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 content filters>,
  "personalizations":<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.

## 11.4. EXAMPLE

### 11.4.1. HTML PAGE

The code for the HTML page with our Tweet button can be found at <https://github.com/apache/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.

### 11.4.2. 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/cxs/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: `http://localhost:8181/cxs/context.json?sessionId=sessionId`. This URI requests context from Unomi, resulting in an updated `cxs` 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`. However, if the client requests `context.js` then useful functions to interact with Unomi are added to the `cxs` 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:

```
{
  "sessionId" : <optional session identifier>,
  "source": <Item source of the context request>,
  "events": <optional array of events to trigger>,
  "requiredProfileProperties": <optional array of property identifiers>,
  "requiredSessionProperties": <optional array of property identifiers>,
  filters: <optional array of filters to evaluate>,
  "personalizations": <optional array of personalizations 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.

Here is an example of a filter request:

```
curl --location --request POST 'http://localhost:8181/cxs/context.json' \  
--header 'Content-Type: application/json' \  
--header 'Cookie: JSESSIONID=48C8AFB3E18B8E3C93C2F4D5B7BD43B7; context-profile-id=01060c4c-  
a055-4c8f-9692-8a699d0c434a' \  
--data-raw '{  
  "source": null,  
  "requireSegments": false,  
  "requiredProfileProperties": null,  
  "requiredSessionProperties": null,  
  "events": null,  
  "filters": [  
    {  
      "id": "filter1",  
      "filters": [  
        {  
          "condition": {  
            "parameterValues": {  
              "propertyName": "properties.gender",  
              "comparisonOperator": "equals",  
              "propertyValue": "male"  
            },  
            "type": "profilePropertyCondition"  
          }  
        }  
      ]  
    }  
  ],  
  "personalizations": null,  
  "profileOverrides": null,  
  "sessionPropertiesOverrides": null,  
  "sessionId": "demo-session-id"  
}'
```

And here's the result:

```
{
  "profileId": "01060c4c-a055-4c8f-9692-8a699d0c434a",
  "sessionId": "demo-session-id",
  "profileProperties": null,
  "sessionProperties": null,
  "profileSegments": null,
  "filteringResults": {
    "filter1": false
  },
  "processedEvents": 0,
  "personalizations": null,
  "trackedConditions": [],
  "anonymousBrowsing": false,
  "consents": {}
}
```

As we can see, the `filter1` filter we sent in our request, in this example, evaluated to false for the current profile, so we can use that result to perform any customization for the current profile, in this case use the fact that he is male.

## PERSONALIZATIONS

Filters make it possible to evaluate conditions against a profile in real-time, but for true personalization it is needed to have a more powerful mechanism: strategies. Sometimes we want to provide multiple variants that each have their own conditions and we want to know which is the best variant to use for the current profile. This can be achieved with the `personalizations` structure in the `ContextRequest`.

Here is an example of a `personalizations` request:

```

curl --location --request POST 'http://localhost:8181/cxs/context.json' \
--header 'Content-Type: application/json' \
--header 'Cookie: JSESSIONID=48C8AFB3E18B8E3C93C2F4D5B7BD43B7; context-profile-id=01060c4c-
a055-4c8f-9692-8a699d0c434a' \
--data-raw '{
  "source": null,
  "requireSegments": false,
  "requiredProfileProperties": null,
  "requiredSessionProperties": [ "unomiControlGroups" ],
  "events": null,
  "filters": null,
  "personalizations": [
    {
      "id": "gender-test",
      "strategy": "matching-first",
      "strategyOptions": {
        "fallback": "var2",
        "controlGroup": {
          "percentage": 10.0,
          "displayName": "Gender test control group",
          "path": "/gender-test",
          "storeInSession": true
        }
      }
    },
    {
      "id": "var1",
      "filters": [
        {
          "appliesOn": null,
          "condition": {
            "parameterValues": {
              "propertyName": "properties.gender",
              "comparisonOperator": "equals",
              "propertyValue": "male"
            },
            "type": "profilePropertyCondition"
          },
          "properties": null
        }
      ],
      "properties": null
    },
    {
      "id": "var2",
      "filters": null,
      "properties": null
    }
  ]
}
',
  "profileOverrides": null,
  "sessionPropertiesOverrides": null,
  "sessionId": "demo-session-id"
}'

```

In the above example, we basically setup two variants : `var1` and `var2` and setup the `var2` to be the fallback variant in case no variant is matched. We could of course specify more than a variant. The `strategy` indicates to the personalization service how to calculate the "winning" variant. In this case the strategy `matching-first` will return variants that match the current profile. We also use the `controlGroups` option to specify that we want to have a control group for this personalization. The `10.0` percentage value represents 10% (0.0 to 100.0) of traffic that will be assigned randomly to the control group. The control group will be stored in the profile and the session of the visitors if they were assigned to it. We also specify that we store the control group in the session (by default it is stored in the profile)

Currently the following strategies are available:

- `matching-first`: will return the variant IDs that match the current profile (using the initial content order)
- `random`: will return a shuffled list of variant IDs (ignoring any conditions)
- `score-sorted`: allows to sort the variants based on scores associated with the filtering conditions, effectively sorting them by the highest scoring condition first.

Here is the result of the above example:

```
{
  "profileId": "01060c4c-a055-4c8f-9692-8a699d0c434a",
  "sessionId": "demo-session-id",
  "profileProperties": null,
  "sessionProperties": {
    "unomiControlGroups": [
      {
        "id": "previousPerso",
        "displayName": "Previous perso",
        "path": "/home/previousPerso.html",
        "timeStamp": "2021-12-15T13:52:38Z"
      }
    ]
  },
  "profileSegments": null,
  "filteringResults": null,
  "processedEvents": 0,
  "personalizations": {
    "gender-test": [
      "var2"
    ]
  },
  "trackedConditions": [
  ],
  "anonymousBrowsing": false,
  "consents": {}
}
```

In the above example we can see the profile and session were assigned to other control groups but not the current one (the ids are different).



## 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 returns 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-samples](#)), 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-samples';
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",
  "systemTags": [
    "event"
  ],
  "parameters": []
}

```

We specify the identifier for the action type, a list of systemTags 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"
interface="org.apache.unomi.api.actions.ActionExecutor">
    <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;
}
```

## 11.5. 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.

## 11.6. ANNEX

Here is an overview of how Unomi processes incoming requests to the [ContextServlet](#).



## 11.7. WEATHER UPDATE SAMPLE

# 12. CONNECTORS

## 12.1. CONNECTORS

Apache Unomi provides the following connectors:

- [Salesforce CRM connector](#)
- [Mailchimp connector](#)

### 12.1.1. 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](#).

## 12.2. SALESFORCE CONNECTOR

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

### 12.2.1. GETTING STARTED

#### SALESFORCE ACCOUNT SETUP

1. Create a new developer account here:

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

2. Create a new Connected App, by going into Setup -> App Manager and click "Create Connected App"
3. 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](#)

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

[Consumer key](#)  
[Consumer secret \(click to see it\)](#)

5. 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", the click on "Reset my security token". You will receive an email with the security token.

#### APACHE UNOMI SETUP

1. You are now ready to configure the Apache Unomi Salesforce Connector. In the `etc/unomi.custom.system.properties` file add/change the following settings:



```
org.apache.unomi.sfdc.user.username=${env:UNOMI_SFDC_USER_USERNAME:-}  
org.apache.unomi.sfdc.user.password=${env:UNOMI_SFDC_USER_PASSWORD:-}  
org.apache.unomi.sfdc.user.securityToken=${env:UNOMI_SFDC_USER_SECURITYTOKEN:-}  
org.apache.unomi.sfdc.consumer.key=${env:UNOMI_SFDC_CONSUMER_KEY:-}  
org.apache.unomi.sfdc.consumer.secret=${env:UNOMI_SFDC_CONSUMER_SECRET:-}
```

## DEPLOYMENT FROM MAVEN REPOSITORY

In this procedure we assume you have access to a Maven repository that contains a compiled version of the Salesforce connector. If this is not the case or you prefer to deploy using a KAR bundle, see the KAR deployment instructions instead.

1. Connect to the Apache Unomi Karaf Shell using :

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

2. 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/${project.version}/xml/features  
feature:install unomi-salesforce-connector-karaf-kar
```

## DEPLOYMENT USING KAR BUNDLE

If you have a KAR bundle (for example after building from source in the [extensions/salesforce-connector/karaf-kar/target](#) directory), you can follow these steps to install :

1. Ensure that Apache Karaf and Apache Unomi are started
2. Execute the following command in karaf: `feature:install unomi-salesforce-connector-karaf-kar`
3. The installation is done !

## TESTING THE CONNECTOR

1. 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 connectors, 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 connectors's defined actions in rules to push or pull data to/from the Salesforce CRM. You can find more information about rules in the [Data Model](#) and the [Getting Started](#) pages.

## 12.2.2. PROPERTIES

To define how Salesforce attributes will be mapped to Unomi profile properties, edit the following entry using the pattern below :

```
org.apache.unomi.sfdc.fields.mappings=${env:UNOMI_SFDC_FIELDS_MAPPINGS:-  
email<=>Email,firstName<=>FirstName,lastName<=>LastName,company<=>Company,phoneNumbe  
r<=>Phone,jobTitle<=>Title,city<=>City,zipCode<=>PostalCode,address<=>Street,sfdcStatus<=>Status,  
sfdcRating<=>Rating}
```

Please note that Salesforce needs the company and the last name to be set, otherwise the lead won't be created. An identifier needs to be set as well. The identifier will be used to map the Unomi profile to the Salesforce lead. By default, the email is set as the identifier, meaning that if a lead in Salesforce and a profile in Unomi have the same email, they'll be considered as the same person.

```
org.apache.unomi.sfdc.fields.mappings.identifier=${env:UNOMI_SFDC_FIELDS_MAPPINGS_IDENTIF  
IER:-email<=>Email}
```

## 12.2.3. HOT-DEPLOYING UPDATES TO THE SALESFORCE CONNECTOR (FOR DEVELOPERS)

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

1. Compile the connectors using:

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

2. Login to the Unomi Karaf Shell using:

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

3. 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
```

4. 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)

#### 12.2.4. 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
```

#### 12.2.5. 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
```

#### 12.2.6. 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 -DsfdcProperties=../test.properties
```

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

## 12.3. MAILCHIMP CONNECTOR

This extension has 3 actions:

- Add a visitor into a defined Mailchimp list.
- Remove a visitor from a defined Mailchimp list.
- Unsubscribe a visitor from a defined Mailchimp list.

### 12.3.1. GETTING STARTED

1. Create a new MailChimp account: <https://login.mailchimp.com/signup/>
2. Generate a new API Key, or get the default: <https://usX.admin.mailchimp.com/account/api/>
3. Configure the MailChimp Connector Basic in the `etc/unomi.custom.system.properties` file and add/change the following settings:

```
org.apache.unomi.mailchimp.apiKey=${env:UNOMI_MAILCHIMP_APIKEY:-yourApiKey}
org.apache.unomi.mailchimp.url.subDomain=${env:UNOMI_MAILCHIMP_URL_SUBDOMAIN:-us16}
```

4. Before starting configure the mapping between Apache Unomi profile properties and MailChimp member properties.

The mapping can't be use with multitued properties. You need to setup your MailChimp properties first in the MailChimp administration.

Go to: [lists/](#)  
Select the triggered list  
[Settings](#)

Then in the cfg file `org.apache.unomi.mailchimp.list.merge-fields.activate={Boolean}` if you like to activate the mapping feature.

This is the property to configure for the mapping, the format is as shown.  
`org.apache.unomi.mailchimp.list.merge-fields.mapping={Apache Unomi property ID}<{MailChimp Tag name}`



there is a particular format for the address `{Apache Unomi property ID}<{MailChimp Tag name}<{MailChimp tag sub entry}`

MailChimp supported type are:

- Date The format is (DD/MM/YYYY) or (MM/DD/YYYY)
- Birthday The format is (DD/MM) or (MM/DD)
- Website or Text They are text
- Number The number will be parse into a Integer
- Phone The North American format is not supported, use international
- Address



Street, City, Country and Zip are mandatory properties, otherwise the address property will be skipped.

```
address<=>ADDRESS<=>addr1,
city<=>ADDRESS<=>city,
zipCode<=>ADDRESS<=>zip,
countryName<=>ADDRESS<=>country
```

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

```
feature:repo-add mvn:org.apache.unomi/unomi-mailchimp-connector-karaf/${project.version}/xml/features
feature:install unomi-mailchimp-connector-karaf-kar
```

## 13. DEVELOPERS

### 13.1. BUILDING

#### 13.1.1. 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.

#### 13.1.2. BUILDING

1. Get the code: `git clone https://github.com/apache/unomi.git`
2. Change to the top level directory of Apache Unomi source distribution.
3. 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 clean install
```

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

4. The distributions will be available under "package/target" directory.

#### 13.1.3. 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> Apache Unomi >= 1.5 : <https://www.elastic.co/downloads/past-releases/elasticsearch-7-4-2>

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>

### 13.1.4. 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
```

You will then need to launch (only on the first Karaf start) the Apache Unomi packages using the following Apache Karaf shell command:

```
unomi:start
```

### 13.1.5. 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 4.2.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-jaxrs 3.3.4  
feature:repo-add cellar 4.1.3  
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.

### 13.1.6. 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 (3):  
  11.0.5, x86_64: "OpenJDK 11.0.5" /Library/Java/JavaVirtualMachines/openjdk-  
  11.jdk/Contents/Home  
  1.8.0_181, x86_64: "Java SE 8" /Library/Java/JavaVirtualMachines/jdk1.8.0_181.jdk/Contents/Home  
  1.7.0_80, x86_64: "Java SE 7" /Library/Java/JavaVirtualMachines/jdk1.7.0_80.jdk/Contents/Home  
  
/Library/Java/JavaVirtualMachines/openjdk-11.jdk/Contents/Home
```



You can then select the one you want using :

```
export JAVA_HOME= `/usr/libexec/java_home -v 1.8.0_181`
```

and then check that it was correctly referenced using:

```
java -version
```

which should give you a result such as this:

```
openjdk version "11.0.5" 2019-10-15
OpenJDK Runtime Environment (build 11.0.5+10)
OpenJDK 64-Bit Server VM (build 11.0.5+10, mixed mode)
```

### 13.1.7. 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 8 or more recent !

To run the tests simply activate the following profile :

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

### 13.1.8. 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 `/cxs/context.js` script, which will try to retrieving 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.

## 13.2. SSH SHELL COMMANDS

Apache Unomi provides its own Apache Karaf Shell commands to make it easy to control the application lifecycle or perform queries or modifications on the internal state of the system.

All Apache Unomi-specific commands are namespaced and use the `unomi:` namespace. You can use the

Apache Karaf Shell's autocompletion to list all the commands available.

### 13.2.1. USING THE SHELL

You can connect to the Apache Karaf SSH Shell using the following command:

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

The default username/password is karaf/karaf. You should change this as soon as possible by editing the `etc/users.properties` file.

Once connected you can simply type in :

```
unomi:
```

And hit the <tab> key to see the list of all the available Apache Unomi commands. Note that some commands are only available when the application is started.

You can also use the `help` command on any command such as in the following example:

```
karaf@root(>) help unomi:migrate
DESCRIPTION
  unomi:migrate

  This will Migrate your date in ES to be compliant with current version

SYNTAX
  unomi:migrate [fromVersionWithoutSuffix]

ARGUMENTS
  fromVersionWithoutSuffix
    Origin version without suffix/qualifier (e.g: 1.2.0)
    (defaults to 1.2.0)
```

### 13.2.2. LIFECYCLE COMMANDS

The commands control the lifecycle of the Apache Unomi server and are used to migrate, start or stop the server.

*Table 4. Table Lifecycle commands*

Command	Arguments	Description
migrate	fromVersion	This command must be used only when the Apache Unomi application is NOT STARTED. It will perform migration of the data stored in Elasticsearch using the argument fromVersion as a starting point.
stop	n/a	Shutdown the Apache Unomi application
start	n/a	Starts the Apache Unomi application. Note that this state will be remembered between Apache Karaf launches, so in general it is only needed after a first installation or after a <a href="#">migrate</a> command
version	n/a	Prints out the currently deployed version of the Apache Unomi application inside the Apache Karaf runtime.

### 13.2.3. RUNTIME COMMANDS

These commands are available once the application is running. If an argument is between brackets [] it means it is optional.

Table 5. Table Runtime commands

Command	Arguments	Description
rule-list	[maxEntries] [--csv]	Lists all the rules registered in the Apache Unomi server. The maxEntries (defaults to 100) will allow you to specify how many entries need to be retrieved. If the value is inferior to the total value, a message will display the total value of rules registered in the server. If you add the "--csv" option the list will be output as a CSV formatted table
rule-view	rule-id	Dumps a single rule in JSON. The rule-id argument can be retrieved from the <a href="#">rule-list</a> command output.

Command	Arguments	Description
rule-remove	rule-id	Removes a single rule from Apache Unomi. The <a href="#">rule-id</a> argument can be retrieved from the <a href="#">rule-list</a> command output. Warning: no confirmation is asked, be careful with this command.
rule-reset-stats	n/a	Resets the rule statistics. This is notably useful when trying to understand rule performance and impact
rule-tail	n/a	Dumps any rule that is executed by the server. Only executed rules are logged here. If you want to have more detailed information about a particular rule's condition evaluation and if it's already been raised use the <a href="#">rule-watch</a> command instead. This tail will continue until a CTRL+C key combination is pressed.
rule-watch	rule-ids	Dumps detailed evaluation and execution information about the rules that are where specified in the <a href="#">rule-ids</a> arguments (you can specify multiple rule identifiers separated by spaces). The <a href="#">Status</a> column has the following values: EVALUATE - indicates that the rule's conditions are being evaluated (but they might not be satisfied), AR PROFILE - means the rule has already been raised for the profile and will therefore not execute again for this profile, AR SESSION - means the rule has already been executed for this session and will therefore only be executed when another session for the profile is created, EXECUTE means the rule's actions are being executed.

Command	Arguments	Description
event-tail	n/a	Dumps any incoming events to the Apache Unomi server to the console. Use CTRL+C to exit tail
event-view	event-id	Dumps a single event in JSON. The <a href="#">event-id</a> can be retrieved from the event-tail command output.
event-list	[max-entries] [event-type] [--csv]	List the last events processed by Apache Unomi. The <a href="#">max-entries</a> parameter can be used to control how many events are displayed (default is 100). The <a href="#">event-type</a> makes it possible to filter the list by event type. The <a href="#">--csv</a> argument is used to output the list as a CSV list instead of an ASCII table.
event-search	profile-id [event-type] [max-entries]	This command makes it possible to search for the last events by <a href="#">profile-id</a> and by <a href="#">event-type</a> . A <a href="#">max-entries</a> parameter (with a default value of 100) is also accepted to control the number of results returned by the search.
action-list	[--csv]	Lists all the rule actions registered in the Apache Unomi server. This command is useful when developing plugins to check that everything is properly registered. If you add the <a href="#">"--csv"</a> option the list will be output as a CSV formatted table
action-view	action-id	Dumps a single action in JSON. The action-id argument can be retrieved from the <a href="#">action-list</a> command output.
condition-list	[csv]	List all the conditions registered in the server. If you add the <a href="#">"--csv"</a> option the list will be output as a CSV formatted table
condition-view	condition-id	Dumps a single condition in JSON. The condition-id can be retrieved from the <a href="#">condition-list</a> command output.

Command	Arguments	Description
profile-list	[--csv]	List the last 10 modified profiles. If you add the "--csv" option the list will be output as a CSV formatted table
profile-view	profile-id	Dumps a single profile in JSON. The profile-id argument can be retrieved from the <a href="#">profile-list</a> command output.
profile-remove	profile-id	Removes a profile identified by <a href="#">profile-id</a> argument. Warning: no confirmation is asked so be careful with this command!
segment-list	[--csv]	Lists all the segments registered in the Apache Unomi server. If you add the "--csv" option the list will be output as a CSV formatted table
segment-view	segment-id	Dumps a single segment in JSON. The segment-id argument can be retrieved from the <a href="#">segment-list</a> command output.
segment-remove	segment-id	Removes a single segment identified by the <a href="#">segment-id</a> argument. Warning: no confirmation is asked so be careful with this command!
session-list	[--csv]	Lists the last 10 sessions by last event date. If you add the "--csv" option the list will be output as a CSV formatted table
session-view	session-id	Dumps a single session in JSON. The session-id argument can be retrieved from the <a href="#">session-list</a> , <a href="#">profile-list</a> or <a href="#">event-tail</a> command output.

Command	Arguments	Description
deploy-definition	[bundleId] [type] [fileName]	This command can be used to force redeployment of definitions from bundles. By default existing definitions will not be overridden unless they come from SNAPSHOT bundles. Using this command you can override this mechanism. Here are some examples of using this command: <a href="#">unomi:deploy-definition 175 rule *</a> will redeploy all the rules provided by bundle with id 175. If you launch the command without any arguments you will get prompts for what you want to deploy from which bundle. If you want to deploy all the definitions of a bundle you can also use wildcards such as in the following example: <a href="#">deploy-definition 175 * *</a> . It is also possible to give no argument to this command and it will then interactively request the definitions you want to deploy.
undeploy-definition	[bundleId] [type] [fileName]	This command does the opposite of the <a href="#">deploy-definition</a> command and works exactly the same way in terms of arguments and interactive mode except that it undeploys definitions instead of deploying them. This command can be very useful when working on a plugin. For example to remove all the definitions deployed by a plugin you can simply use the following command: <a href="#">undeploy-definition BUNDLE_ID * *</a> when <a href="#">BUNDLE_ID</a> is the identifier of the bundle that contains your plugin.

## 13.3. WRITING PLUGINS

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

## 13.4. TYPES VS. INSTANCES

Several extension points in Unomi rely on the concept of type: a plugin 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.

So for example we have the following types vs instances:

- ConditionTypes vs Conditions
- ActionTypes vs Actions
- PropertyTypes vs Properties (for profiles and sessions)

## 13.5. 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
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.

## 13.6. EXTENSION POINTS

In this section the value types that may be used as extension points are presented. Examples of these types will be given in the next section with more details.

### 13.6.1. 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”.

### 13.6.2. 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”.

### 13.6.3. 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.

### 13.6.4. PROPERTYMERGESTRATEGYTYPE

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

### 13.6.5. 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 systemTags or file system structure, using sub-directories to organize definition files.

### 13.6.6. 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.

### 13.6.7. 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.

### 13.6.8. 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.

### 13.6.9. TAG

`Tag`s are simple labels that are used to classify all other objects inside Unomi.

### 13.6.10. VALUETYPE

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

## 13.7. CUSTOM PLUGINS

Apache Unomi is a pluggable server that may be extended in many ways. This document assumes you are familiar with the [Apache Unomi Data Model](#) . 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](#).

### 13.7.1. CREATING A PLUGIN

An plugin 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-plugins</artifactId>
    <version>${project.version}</version>
  </parent>

  <modelVersion>4.0.0</modelVersion>

  <artifactId>unomi-plugin-example</artifactId>
  <name>Apache Unomi :: Plugins :: Example</name>
  <description>A sample example of a Unomi plugin</description>
  <version>${project.version}</version>
  <packaging>bundle</packaging>

  <dependencies>
    <!-- This dependency is not required but generally used in plugins -->
    <dependency>
      <groupId>org.apache.unomi</groupId>
      <artifactId>unomi-api</artifactId>
      <version>${project.version}</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>

```

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

### 13.7.2. DEPLOYMENT AND CUSTOM DEFINITION

When you deploy a custom bundle with a custom definition (see "Predefined xxx" chapters under) for

the first time, the definition will automatically be deployed at your bundle start event **if it does not exist**. After that if you redeploy the same bundle, the definition will not be redeployed, but you can redeploy it manually using the command `unomi:deploy-definition <bundleId> <fileName>` If you need to modify an existing definition when deploying the module, see [Migration patches](#).

### 13.7.3. 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.

### 13.7.4. 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.

### 13.7.5. 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",
    "systemTags": ["properties", "profileProperties", "contactProfileProperties"]
  },
  "type": "string",
  "defaultValue": "",
  "automaticMappingsFrom": [ ],
  "rank": "304.0"
}

```

### 13.7.6. 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": "",
    "systemTags": [
      "event",
      "eventCondition"
    ],
    "readOnly": true
  },
  "parentCondition": {
    "type": "eventTypeCondition",
    "parameterValues": {
      "eventType": "profileUpdated"
    }
  },
  "parameters": [
  ]
}

```

### 13.7.7. 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"
      },
      "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.

### 13.7.8. CUSTOM ACTION TYPES

Custom action types 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 type 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": "",
    "systemTags": [
      "demographic",
      "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 interface="org.apache.unomi.api.actions.ActionExecutor">
    <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`



### 13.7.9. CUSTOM CONDITION TYPES

Custom condition types 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 type definition:

```
{
  "metadata": {
    "id": "matchAllCondition",
    "name": "matchAllCondition",
    "description": "",
    "systemTags": [
      "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 type 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](#)

## 13.8. MIGRATION PATCHES

You may provide patches on any predefined items by simply adding a JSON file in :

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

These patches will be applied when the module will be deployed the first time. They allow to modify an item, that would have been previously deployed on unomi by a previous version of the extension or by something else.

Each patch must have a unique id - unomi will use this id to remember that the patch has already been applied. It can also be used to reapply the patch when need by using the karaf command [unomi:deploy-definition](#)

A patch also need to reference the item to patch by setting [patchedItemId](#) and [patchedItemType](#), and an operation that tells what the patch should do.

[patchedItemType](#) can take one of the following value:

- condition
- action
- goal
- campaign
- persona
- propertyType
- rule
- segment
- scoring

`operation` can take one of the following value:

- patch
- override
- remove

You can apply a patch in `json-patch` format in the `data` field, and by specifying operation `patch` like in this example :

```
{
  "itemId": "firstName-patch1",
  "patchedItemId": "firstName",
  "patchedItemType": "propertyType",
  "operation": "patch",
  "data": [
    {
      "op": "replace", "path": "/defaultValue", "value": "foo"
    }
  ]
}
```

If you need to completely redeploy a definition, you can use the `override` operation and put the definition in `data`

```
{
  "itemId": "gender-patch1",
  "patchedItemId": "gender",
  "patchedItemType": "propertyType",
  "operation": "override",
  "data": {
    "metadata": {
      "id": "gender",
      "name": "Gender",
      "systemTags": [
        "properties",
        "profileProperties"
      ]
    },
    "type": "string",
    "defaultValue": "foo",
    "automaticMappingsFrom": [ ],
    "rank": "105.0"
  }
}
```

It is also possible to simply remove an item by using the operation [remove](#) :

```
{
  "itemId": "firstName-patch2",
  "patchedItemId": "firstName",
  "patchedItemType": "propertyType",
  "operation": "remove"
}
```

Patches can also be deployed at runtime by using the REST endpoint `/patch/apply` .