

Developers Guide

Copyright (c) 2015-2016 The OpenNMS Group, Inc.

OpenNMS Horizon 2016.1.1
Last updated 2017-02-02 11:42:03 -05:00

Table of Contents

1. CORS Support	1
1.1. Why do I need CORS support?	1
1.2. How can I enable CORS support?	1
1.3. How can I configure CORS support?	1
2. ReST API	2
2.1. ReST URL	2
2.2. Authentication	2
2.3. Data format	2
2.4. Standard Parameters	2
2.5. Standard filter examples	3
2.6. HTTP Return Codes	3
2.7. Identifying Resources	4
2.8. Currently Implemented Interfaces	4
2.8.1. Acknowledgements	4
2.8.2. Alarm Statistics	5
2.8.3. Alarms	5
2.8.4. Events	6
2.8.5. Foreign Sources	7
2.8.6. Groups	8
2.8.7. Heatmap	9
2.8.8. KSC Reports	10
2.8.9. Maps	10
2.8.10. Measurements API	11
2.8.11. Nodes	14
2.8.12. Notifications	15
2.8.13. Outage Timelines	16
2.8.14. Outages	16
2.8.15. Requisitions	16
2.8.16. Resources API	18
2.8.17. Realtime Console data	21
2.8.18. Scheduled Outages	22
2.8.19. SNMP Configuration	23
2.8.20. Users	31
2.8.21. SNMP Trap Northbounder Interface Configuration	31
2.8.22. Email Northbounder Interface Configuration	32
2.8.23. Javamail Configuration	33
2.8.24. Syslog Northbounder Interface Configuration	34
2.9. ReST API Examples	35
2.9.1. Getting Graph data	35
2.9.2. provision.pl examples and notes	36
2.9.3. Debian (Lenny) Notes	37
2.9.4. Windows Powershell ReST	37
3. Develop Documentation	39
3.1. File Structure in opennms-doc	39
3.2. Writing	39
3.2.1. Conventions for text formatting	39
3.2.2. Gotchas	40

3.3. Headings and document structure	40
3.4. Links	41
3.5. Admonitions and useful notes	42
3.6. Attributes	43
3.7. Comments	43
3.8. Tables	43
3.9. Include images	44
3.10. Code Snippets	46
3.10.1. Explicitly defined in the document	46
3.10.2. Included from an example file	47
3.10.3. Include parts of a file	47
3.11. Cheat Sheets and additional hints	48
4. AMQP Integration	49
4.1. Event Forwarder	49
4.1.1. Setup	49
4.1.2. Debugging	50
4.2. Event Receiver	50
4.2.1. Setup	50
4.2.2. Debugging	51
4.3. Alarm Northbounder	51
4.3.1. Setup	51
4.3.2. Debugging	51
4.4. Custom Processors	52

Chapter 1. CORS Support

1.1. Why do I need CORS support?

By default, many browsers implement a *same origin policy* which prevents making requests to a resource, on an origin that's different from the source origin.

For example, a request originating from a page served from <http://www.opennms.org> to a resource on <http://www.adventuresinoss.com> would be considered a cross origin request.

CORS (Cross Origin Resource Sharing) is a standard mechanism used to enable cross origin requests.

For further details, see:

- [Mozilla's HTTP access control \(CORS\)](#)
- [W3C's CORS Spec](#)

1.2. How can I enable CORS support?

CORS support for the REST interface (or any other part of the Web UI) can be enabled as follows:

1. Open '\$OPENNMS_HOME/jetty-webapps/opennms/WEB-INF/web.xml' for editing.
2. Apply the CORS filter to the '/rest/' path by removing the comments around the **<filter-mapping>** definition. The result should look like:

```
<!-- Uncomment this to enable CORS support -->
<filter-mapping>
  <filter-name>CORS Filter</filter-name>
  <url-pattern>/rest/*</url-pattern>
</filter-mapping>
```

3. Restart OpenNMS Horizon

1.3. How can I configure CORS support?

CORS support is provided by the **org.ebaysf.web.cors.CORSFilter** servlet filter.

Parameters can be configured by modifying the filter definition in the 'web.xml' file referenced above.

By default, the allowed origins parameter is set to '*'.

The complete list of parameters supported are available from:

- <https://github.com/ebay/cors-filter>

Chapter 2. ReST API

A RESTful interface is a web service conforming to the REST architectural style as described in the book [RESTful Web Services](#). This page describes the RESTful interface for OpenNMS Horizon.

2.1. ReST URL

The base URL for Rest Calls is: <http://opennmserver:8980/opennms/rest/>

For instance, <http://localhost:8980/opennms/rest/alarms/> will give you the current alarms in the system.

2.2. Authentication

Use HTTP Basic authentication to provide a valid username and password. By default you will not receive a challenge, so you must configure your ReST client library to send basic authentication proactively.

2.3. Data format

Jersey allows ReST calls to be made using either XML or JSON. By default a request to the API is returned in XML. To get JSON encoded responses one has to send the following header with the request: **Accept: application/json**.

2.4. Standard Parameters

The following are standard params which are available on most resources (noted below)

Table 1. ReST standard parameter for resources

Parameter	Description
limit	integer, limiting the number of results. This is particularly handy on events and notifications, where an accidental call with no limit could result in many thousands of results being returned, killing either the client or the server. If set to 0, then no limit applied
offset	integer, being the numeric offset into the result set from which results should start being returned. E.g., if there are 100 result entries, offset is 15, and limit is 10, then entries 15-24 will be returned. Used for pagination
Filtering: All properties of the entity being accessed can be specified as parameters in either the <i>URL</i> (for <i>GET</i>) or the form value (for <i>PUT</i> and <i>POST</i>). If so, the value will be used to add a filter to the result. By default, the operation is equality, unless the comparator parameter is sent, in which case it applies to all comparisons in the filter. Multiple properties will result in an AND operation between the filter elements. Available comparators are:	
eq	Checks for equality
ne	Checks for non-equality
ilike	Case-insensitive wildcarding (% is the wildcard)
like	Case-sensitive wildcarding (% is the wildcard)
gt	Greater than
lt	Less than
ge	Greater than or equal
le	Less than or equal

If the value `null` is passed for a given property, then the obvious operation will occur (comparator will be ignored for that property). `nonnull` is handled similarly.

- **Ordering:** If the parameter `orderBy` is specified, results will be ordered by the named property. Default is ascending, unless the `order` parameter is set to `desc` (any other value will default to ascending)
- **Raw where clause:** If there is a `query` parameter, it will be used as a raw where clause (SQL, not HQL), and added to any other filters created by other parameters

2.5. Standard filter examples

Take `/events` as an example.

Resource	Description
<code>/events?eventUei=uei.opennms.org/international/rtc/subscribe</code>	would return the first 10 events with the rtc subscribe UEI, (10 being the default limit for events)
<code>/events?eventUei=uei.opennms.org/international/rtc/subscribe&limit=0</code>	would return all the rtc subscribe events (potentially quite a few)
<code>/events?id=100&comparator=gt</code>	would return the first 10 events with an id greater than 100
<code>/events?eventAckTime=nonnull</code>	would return the first 10 events that have a non-null Ack time (i.e. those that have been acknowledged)
<code>/events?eventAckTime=nonnull&id=100&comparator=gt&limit=20</code>	would return the first 20 events that have a non-null Ack time and an id greater than 100. Note that the <code>nonnull</code> value causes the comparator to be ignored for <code>eventAckTime</code>
<code>/events?eventAckTime=2008-07-28T04:41:30.530%2B12:00&id=100&comparator=gt&limit=20</code>	would return the first 20 events that have were acknowledged after 28th July 2008 at 4:41am (+12:00), and an id greater than 100. Note that the same comparator applies to both property comparisons. Also note that you must URL encode the plus sign when using GET.
<code>/events?orderBy=id&order=desc</code>	would return the 10 latest events inserted (probably, unless you've been messing with the id's)

2.6. HTTP Return Codes

The following apply for OpenNMS Horizon 18 and newer.

- All the DELETE requests are going to return a 204 (NO_CONTENT) on success.
- All the PUT requests are going to return a 204 (NO_CONTENT) on success.
- All the POST requests that can either add or update an entity are going to return a 204 (NO_CONTENT) on success.
- All the POST associated to resource addition are going to return a 201 (CREATED) on success.
- All the POST requests where it is required to return an object will return a 200 (OK).
- All the requests excepts GET for the Requisitions end-point and the Foreign Sources Definitions end-point will return 202 (ACCEPTED). This is because all the requests are actually executed asynchronously and there is no way to know the status of the execution, or wait until the processing is done.
- If a resource is not modified during a PUT request, a NOT_MODIFIED will be returned. A NO_CONTENT will be returned only on a success operation.

- All GET requests are going to return 200 (OK) on success.
- All GET requests are going to return 404 (NOT_FOUND) when a single resource doesn't exist; but will return 400 (BAD_REQUEST), if an intermediate resource doesn't exist. For example, if a specific IP doesn't exist on a valid node, return 404. But, if the IP is valid and the node is not valid, because the node is an intermediate resource, a 400 will be returned.
- If something not expected is received from the Service/DAO Layer when processing any HTTP request, like an exception, a 500 (INTERNAL_SERVER_ERROR) will be returned.
- Any problem related with the incoming parameters, like validations, will generate a 400 (BAD_REQUEST).

2.7. Identifying Resources

Some endpoints deal in resources, which are identified by *Resource IDs*. Since every resource is ultimately parented under some node, identifying the node which contains a resource is the first step in constructing a resource ID. Two styles are available for identifying the node in a resource ID:

Style	Description	Example
<code>node[ID]</code>	Identifies a node by its database ID, which is always an integer	<code>node[42]</code>
<code>node[FS:FID]</code>	Identifies a node by its foreign-source name and foreign-ID, joined by a single colon	<code>node[Servers:115da833-0957-4471-b496-a731928c27dd]</code>

The node identifier is followed by a period, then a resource-type name and instance name. The instance name's characteristics may vary from one resource-type to the next. A few examples:

Value	Description
<code>nodeSnmp[]</code>	Node-level (scalar) performance data for the node in question. This type is the only one where the instance identifier is empty.
<code>interfaceSnmp[eth0-04013f75f101]</code>	A layer-two interface as represented by a row in the SNMP <code>ifTable</code> . The instance identifier is composed of the interface's <code>ifName</code> and its <code>ifPhysAddress</code> (if it has one).
<code>dskIndex[_root_fs]</code>	The root filesystem of a node running the Net-SNMP management agent.

Putting it all together, here are a few well-formed resource IDs:

- `node[1].nodeSnmp[]`
- `node[42].interfaceSnmp[eth0-04013f75f101]`
- `node[Servers:115da833-0957-4471-b496-a731928c27dd].dskIndex[_root_fs]`

2.8. Currently Implemented Interfaces

2.8.1. Acknowledgements

NOTE | the default offset is 0, the default limit is 10 results. To get all results, use `limit=0` as a parameter on the URL (ie, `GET /acks?limit=0`).

GETs (Reading Data)

Resource	Description
<code>/acks</code>	Get a list of acknowledgements.
<code>/acks/count</code>	Get the number of acknowledgements. (Returns plaintext, rather than <i>XML</i> or <i>JSON</i> .)
<code>/acks/{id}</code>	Get the acknowledgement specified by the given <i>ID</i> .

POSTs (Setting Data)

Resource	Description
<code>/acks</code>	Creates or modifies an acknowledgement for the given alarm <i>ID</i> or notification <i>ID</i> . To affect an alarm, set an <code>alarmId</code> parameter in the URL-encoded <i>POST</i> body; to affect a notification, set <code>notifyId</code> instead. An <code>action</code> parameter is also required, and may be one of <code>ack</code> , <code>unack</code> , <code>clear</code> , or <code>esc</code> (escalate).

Usage examples with curl

Acknowledge notification #3

```
curl -u 'admin:admin' -X POST -d notifId=3 -d action=ack http://localhost:8980/opennms/rest/acks
```

Escalate alarm #42

```
curl -u 'admin:admin' -X POST -d alarmId=42 -d action=esc http://localhost:8980/opennms/rest/acks
```

2.8.2. Alarm Statistics

It is possible to get some basic statistics on alarms, including the number of acknowledged alarms, total alarms, and the newest and oldest of acknowledged and unacknowledged alarms.

GETs (Reading Data)

Resource	Description
<code>/stats/alarms</code>	Returns statistics related to alarms. Accepts the same Hibernate parameters that you can pass to the <code>/alarms</code> ReST service.
<code>/stats/alarms/by-severity</code>	Returns the statistics related to alarms, one per severity. You can optionally pass a list of severities to the <code>severities</code> query parameter to limit it to the specified severities. (eg, <code>GET /opennms/rest/stats/alarms/by-severity?severities=MAJOR,CRITICAL</code>).

2.8.3. Alarms

NOTE

the default offset is 0, the default limit is 10 results. To get all results, use `limit=0` as a parameter on the URL (ie, `GET /events?limit=0`).

GETs (Reading Data)

Resource	Description
<code>/alarms</code>	Get a list of alarms.

Resource	Description
<code>/alarms/count</code>	Get the number of alarms. (Returns plaintext, rather than <i>XML</i> or <i>JSON</i> .)
<code>/alarms/{id}</code>	Get the alarms specified by the given <i>ID</i> .

Note that you can also query by severity, like so:

Resource	Description
<code>/alarms?comparator=ge&severity=MINOR</code>	Get the alarms with a severity greater than or equal to <i>MINOR</i> .

PUTs (Modifying Data)

PUT requires form data using `application/x-www-form-urlencoded` as a Content-Type.

Resource	Description
<code>/alarms/{id}?ack='(true;false)'</code>	Acknowledges (or unacknowledges) an alarm.
<code>/alarms?x=y& &ack='(true;false)'</code>	Acknowledges (or unacknowledges) alarms matching the additional query parameters. eg, <code>/alarms?node.id=4&ack=true</code>

New in OpenNMS 1.11.0

In OpenNMS 1.11.0, some additional features are supported in the alarm ack API:

Resource	Description
<code>/alarms/{id}?clear=true</code>	Clears an alarm.
<code>/alarms/{id}?escalate=true</code>	Escalates an alarm. eg, <code>NORMAL MINOR MAJOR CRITICAL</code> , etc.
<code>/alarms?x=y& &clear=true</code>	Clears alarms matching the additional query parameters.
<code>/alarms?x=y& &escalate=true</code>	Escalates alarms matching the additional query parameters.

Additionally, when acknowledging alarms (`ack=true`) you can now specify an `ackUser` parameter. You will only be allowed to `ack` as a different user IF you are PUTting as an authenticated user who is in the *admin* role.

Queries

As noted above, it is possible to pass a raw `query` parameter when doing ReST queries. In the case of alarms, it is possible to pass severity names when querying by severity, rather than having to know the number that the severity enum maps to. For example:

```
/alarms?query=lastEventTime%20%3E%20'2011-08-19T11%3A11%3A11.000-07%3A00'%20AND%20severity%20%3E%20MAJOR%20AND%20alarmAckUser%20IS%20NULL
```

This will get any alarms where the last event associated with the alarm is newer than August 19th, 2011 11:11:11, the severity is greater than *MAJOR*, and the alarm is not acknowledged (`alarmAckUser` is null). You should be able to use any column in the *alarm*, *event*, *node*, *ipinterface*, or *snmpinterface* tables.

2.8.4. Events

GETs (Reading Data)

Resource	Description
<code>/events</code>	Get a list of events. The default for offset is 0, and the default for limit is 10. To get all results, use <code>limit=0</code> as a parameter on the URL (ie, <code>GET /events?limit=0</code>).
<code>/events/count</code>	Get the number of events. (Returns plaintext, rather than <i>XML</i> or <i>JSON</i> .)
<code>/events/{id}</code>	Get the event specified by the given <i>ID</i> .

PUTs (Modifying Data)

PUT requires form data using `application/x-www-form-urlencoded` as a Content-Type.

Resource	Description
<code>/events/{id}?ack='(true;false)</code>	Acknowledges (or unacknowledges) an event.
<code>/events?x=y&&ack='(true;false)</code>	Acknowledges (or unacknowledges) the matching events.

POSTs (Adding Data)

POST requires XML (`application/xml`) or JSON (`application/json`) as its Content-Type.

TIP | See `${OPENNMS_HOME}/share/xsds/event.xsd` for the reference schema.

Resource	Description
<code>/events</code>	Publish an event on the event bus.

2.8.5. Foreign Sources

RESTful service to the OpenNMS Horizon Provisioning Foreign Source definitions. Foreign source definitions are used to control the scanning (service detection) of services for SLA monitoring as well as the data collection settings for physical interfaces (resources).

This API supports CRUD operations for managing the Provisioner's foreign source definitions. Foreign source definitions are *POSTed* and will be deployed when the corresponding requisition gets imported/synchronized by Provisiond.

If a request says that it gets the "active" foreign source, that means it returns the pending foreign source (being edited for deployment) if there is one, otherwise it returns the deployed foreign source.

GETs (Reading Data)

Resource	Description
<code>/foreignSources</code>	Get all active foreign sources.
<code>/foreignSources/default</code>	Get the active default foreign source.
<code>/foreignSources/deployed</code>	Get the list of all deployed (active) foreign sources.
<code>/foreignSources/deployed/count</code>	Get the number of deployed foreign sources. (Returns plaintext, rather than <i>XML</i> or <i>JSON</i> .)
<code>/foreignSources/{name}</code>	Get the active foreign source named {name}.

Resource	Description
<code>/foreignSources/{name}/detectors</code>	Get the configured detectors for the foreign source named {name}.
<code>/foreignSources/{name}/detectors/{detector}</code>	Get the specified detector for the foreign source named {name}.
<code>/foreignSources/{name}/policies</code>	Get the configured policies for the foreign source named {name}.
<code>/foreignSources/{name}/policies/{policy}</code>	Get the specified policy for the foreign source named {name}.

POSTs (Adding Data)

POST requires XML using application/xml as its Content-Type.

Resource	Description
<code>/foreignSources</code>	Add a foreign source.
<code>/foreignSources/{name}/detectors</code>	Add a detector to the named foreign source.
<code>/foreignSources/{name}/policies</code>	Add a policy to the named foreign source.

PUTs (Modifying Data)

PUT requires form data using application/x-www-form-urlencoded as a Content-Type.

Resource	Description
<code>/foreignSources/{name}</code>	Modify a foreign source with the given name.

DELETEs (Removing Data)

Resource	Description
<code>/foreignSources/{name}</code>	Delete the named foreign source.
<code>/foreignSources/{name}/detectors/{detector}</code>	Delete the specified detector from the named foreign source.
<code>/foreignSources/{name}/policies/{policy}</code>	Delete the specified policy from the named foreign source.

2.8.6. Groups

Like users, groups have a simplified interface as well.

GETs (Reading Data)

Resource	Description
<code>/groups</code>	Get a list of groups.
<code>/groups/{groupname}</code>	Get a specific group, given a group name.
<code>/groups/{groupname}/users</code>	Get the users for a group, given a group name. (new in OpenNMS 14)
<code>/groups/{groupname}/categories</code>	Get the categories associated with a group, given a group name. (new in OpenNMS 14)

POSTs (Adding Data)

Resource	Description
<code>/groups</code>	Add a new group.

PUTs (Modifying Data)

Resource	Description
<code>/groups/{groupname}</code>	Update the metadata of a group (eg, change the <code>comments</code> field).
<code>/groups/{groupname}/users/{username}</code>	Add a user to the group, given a group name and username. (new in OpenNMS 14)
<code>/groups/{groupname}/categories/{categoryname}</code>	Associate a category with the group, given a group name and category name. (new in OpenNMS 14)

DELETEs (Removing Data)

Resource	Description
<code>/groups/{groupname}</code>	Delete a group.
<code>/groups/{groupname}/users/{username}</code>	Remove a user from the group. (new in OpenNMS 14)
<code>/groups/{groupname}/categories/{categoryname}</code>	Disassociate a category from a group, given a group name and category name. (new in OpenNMS 14)

2.8.7. Heatmap

GETs (Reading Data)

Resource	Description
<code>/heatmap/outages/categories</code>	Sizes and color codes based on outages for nodes grouped by <i>Surveillance Categories</i>
<code>/heatmap/outages/foreignSources</code>	Sizes and color codes based on outages for nodes grouped by <i>Foreign Source</i>
<code>/heatmap/outages/monitoredServices</code>	Sizes and color codes based on outages for nodes grouped by monitored services
<code>/heatmap/outages/nodesByCategory/{category}</code>	Sizes and color codes based on outages for nodes associated with a specific <i>Surveillance Category</i>
<code>/heatmap/outages/nodesByForeignSource/{foreignSource}</code>	Sizes and color codes based on outages for nodes associated with a specific <i>Foreign Source</i>
<code>/heatmap/outages/nodesByMonitoredService/{monitoredService}</code>	Sizes and color codes based on outages for nodes providing a specific monitored service

Resource	Description
<code>/heatmap/alarms/categories</code>	Sizes and color codes based on alarms for nodes grouped by <i>Surveillance Categories</i>
<code>/heatmap/alarms/foreignSources</code>	Sizes and color codes based on alarms for nodes grouped by <i>Foreign Source</i>
<code>/heatmap/alarms/monitoredServices</code>	Sizes and color codes based on alarms for nodes grouped by monitored services
<code>/heatmap/alarms/nodesByCategory/{category}</code>	Sizes and color codes based on alarms for nodes associated with a specific <i>Surveillance Category</i>

Resource	Description
<code>/heatmap/alarms/nodesByForeignSource/{foreignSource}</code>	Sizes and color codes based on alarms for nodes associated with a specific <i>Foreign Source</i>
<code>/heatmap/alarms/nodesByMonitoredService/{monitoredService}</code>	Sizes and color codes based on alarms for nodes providing a specific monitored service

2.8.8. KSC Reports

GETs (Reading Data)

Resource	Description
<code>/ksc</code>	Get a list of all KSC reports, this includes ID and label.
<code>/ksc/{reportId}</code>	Get a specific KSC report, by ID.
<code>/ksc/count</code>	Get a count of all KSC reports.

PUTs (Modifying Data)

Resource	Description
<code>/ksc/{reportId}</code>	Modify a report with the given ID.

POSTs (Creating Data)

Documentation incomplete see issue: [NMS-7162](#)

DELETEs (Removing Data)

Documentation incomplete see issue: [NMS-7162](#)

2.8.9. Maps

The *SVG maps* use *ReST* to populate their data. This is the interface for doing that.

GETs (Reading Data)

Resource	Description
<code>/maps</code>	Get the list of maps.
<code>/maps/{id}</code>	Get the map with the given <i>ID</i> .
<code>/maps/{id}/mapElements</code>	Get the elements (<i>nodes, links, etc.</i>) for the map with the given <i>ID</i> .

POSTs (Adding Data)

Resource	Description
<code>/maps</code>	Add a map.

PUTs (Modifying Data)

Resource	Description
<code>/maps/{id}</code>	Update the properties of the map with the given <i>ID</i> .

DELETES (Removing Data)

Resource	Description
<code>/maps/{id}</code>	Delete the map with the given <i>ID</i> .

2.8.10. Measurements API

The *Measurements API* can be used to retrieve collected values stored in *RRD* (or *JRB*) files. Note that all units of time are expressed in milliseconds.

GETs (Reading Data)

Resource	Description
<code>/measurements/{resourceId}/{attribute}</code>	Retrieve the measurements for a single attribute

The following table shows all supported query string parameters and their default values.

name	default	comment
start	-14400000	Timestamp in milliseconds. If < 0, the effective value will be (end + start).
end	0	Timestamp in milliseconds. If 0, the effective value will be the current timestamp.
step	300000	Requested time interval between rows. Actual step may differ. Set to 1 for maximum accuracy.
maxrows	0	When using the measurements to render a graph, this should be set to the graph's pixel width.
interval	null	Duration in milliseconds, used by strategies that implement late aggregation.
heartbeat	null	Duration in milliseconds, used by strategies that implement late aggregation.
aggregation	AVERAGE	Consolidation function used. Can typically be AVERAGE , MIN or MAX . Depends on RRA definitions.
fallback-attribute		Secondary attribute that will be queried in the case the primary attribute does not exist.

Usage examples with curl

Retrieve CPU counter metrics over the last 2 hours for node 1

```
curl -u admin:admin
"http://127.0.0.1:8980/opennms/rest/measurements/node%5B1%5D.nodeSnmp%5B%5D/CpuRawUser?start=-
7200000&maxrows=30&aggregation=AVERAGE"
```

Response

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<query-response end="1425588138256" start="1425580938256" step="300000">
  <columns>
    <values>159.5957271523179</values>
    <values>158.08531037527592</values>
    <values>158.45835584842285</values>
    ...
  </columns>
  <labels>CpuRawUser</labels>
  <timestamps>1425581100000</timestamps>
  <timestamps>1425581400000</timestamps>
  <timestamps>1425581700000</timestamps>
  ...
</query-response>
```

POSTs (Reading Data)

Resource	Description
/measurements	Retrieve the measurements for one or more attributes, possibly spanning multiple resources, with support for JEXL expressions.

Here we use a POST instead of a GET to retrieve the measurements, which allows us to perform complex queries which are difficult to express in a query string. These requests cannot be used to update or create new metrics.

An example of the POST body is available bellow.

Usage examples with curl

Retrieve bits in and bits out metrics for a particular interface. Perform calculations on bits out, and only return the derived values.

```
curl -X POST -H "Accept: application/json" -H "Content-Type: application/json" -u admin:admin -d
@report.json http://127.0.0.1:8980/opennms/rest/measurements
```

Contents of report.json

```
{
  "start": 1425563626316,
  "end": 1425585226316,
  "step": 10000,
  "maxrows": 1600,
  "source": [
    {
      "aggregation": "AVERAGE",
      "attribute": "ifHCInOctets",
      "label": "ifHCInOctets",
      "resourceId": "nodeSource[Servers:1424038123222].interfaceSnmpp[eth0-04013f75f101]",
      "transient": "false"
    },
    {
      "aggregation": "AVERAGE",
      "attribute": "ifHCOutOctets",
      "label": "ifHCOutOctets",
      "resourceId": "nodeSource[Servers:1424038123222].interfaceSnmpp[eth0-04013f75f101]",
      "transient": "true"
    }
  ],
  "expression": [
    {
      "label": "ifHCOutOctetsNeg",
      "value": "-1.0 * ifHCOutOctets",
      "transient": "false"
    }
  ]
}
```


Response

```
{
  "step": 300000,
  "start": 1425563626316,
  "end": 1425585226316,
  "timestamps": [
    1425563700000,
    1425564000000,
    1425564300000,
    ...
  ],
  "labels": [
    "ifHCInOctets",
    "ifHCOutOctetsNeg"
  ],
  "columns": [
    {
      "values": [
        139.94817275747508,
        199.0062569213732,
        162.6264894795127,
        ...
      ]
    },
    {
      "values": [
        -151.66179401993355,
        -214.7415503875969,
        -184.9012624584718,
        ...
      ]
    }
  ]
}
```

2.8.11. Nodes

Note: the default offset is 0, the default limit is 10 results. To get all results, use `limit=0` as a parameter on the URL (ie, GET `/nodes?limit=0`).

Additionally, anywhere you use "id" in the queries below, you can use the foreign source and foreign ID separated by a colon instead (ie, GET `/nodes/fs:fid`).

GETs (Reading Data)

Resource	Description
<code>/nodes</code>	Get a list of nodes. This includes the ID and node label.
<code>/nodes/{id}</code>	Get a specific node, by ID.
<code>/nodes/{id}/ipinterfaces</code>	Get the list of IP interfaces associated with the given node.
<code>/nodes/{id}/ipinterfaces/{ipAddress}</code>	Get the IP interface for the given node and IP address.
<code>/nodes/{id}/ipinterfaces/{ipAddress}/services</code>	Get the list of services associated with the given node and IP interface.
<code>/nodes/{id}/ipinterfaces/{ipAddress}/services/{service}</code>	Get the requested service associated with the given node, IP interface, and service name.

Resource	Description
/nodes/{id}/snmpinterfaces	Get the list of SNMP interfaces associated with the given node.
/nodes/{id}/snmpinterfaces/{ifIndex}	Get the specific interface associated with the given node and ifIndex.
/nodes/{id}/categories	Get the list of categories associated with the given node.
/nodes/{id}/categories/{categoryName}	Get the category associated with the given node and category name.
/nodes/{id}/assetRecord	Get the asset record associated with the given node.

POSTs (Adding Data)

POST requires XML using application/xml as its Content-Type.

Resource	Description
/nodes	Add a node.
/nodes/{id}/ipinterfaces	Add an IP interface to the node.
/nodes/{id}/ipinterfaces/{ipAddress}/services	Add a service to the interface for the given node.
/nodes/{id}/snmpinterfaces	Add an SNMP interface to the node.
/nodes/{id}/categories	Add a category association to the node.

PUTs (Modifying Data)

PUT requires form data using application/x-www-form-urlencoded as a Content-Type.

Resource	Description
/nodes/{id}	Modify a node with the given ID.
/nodes/{id}/ipinterfaces/{ipAddress}	Modify the IP interface with the given node ID and IP address.
/nodes/{id}/ipinterfaces/{ipAddress}/services/{service}	Modify the service with the given node ID, IP address, and service name.
/nodes/{id}/snmpinterfaces/{ifIndex}	Modify the SNMP interface with the given node ID and ifIndex.
/nodes/{id}/categories/{categoryName}	Modify the category with the given node ID and name.

DELETES (Removing Data)

Perform a DELETE to the singleton URLs specified in [PUT](#) above to delete that object.

2.8.12. Notifications

Note: the default offset is 0, the default limit is 10 results. To get all results, use `limit=0` as a parameter on the URL (ie, [GET /events?limit=0](#)).

GETs (Reading Data)

Resource	Description
/notifications	Get a list of notifications.
/notifications/count	Get the number of notifications. (Returns plaintext, rather than <i>XML</i> or <i>JSON</i> .)
/notifications/{id}	Get the notification specified by the given <i>ID</i> .

To acknowledge or unacknowledge a notification, use the `acks` endpoint — see [Acknowledgements](#).

2.8.13. Outage Timelines

GETs (Reading Data)

Resource	Description
<code>/header/{start}/{end}/{width}</code>	Generate the timeline header
<code>/image/{nodeId}/{ipAddress}/{serviceName}/{start}/{end}/{width}</code>	Generate the timeline image
<code>/empty/{start}/{end}/{width}</code>	Generate an empty timeline for non-monitored services
<code>/html/{nodeId}/{ipAddress}/{serviceName}/{start}/{end}/{width}</code>	Generate the raw HTML for the image

2.8.14. Outages

GETs (Reading Data)

Resource	Description
<code>/outages</code>	Get a list of outages.
<code>/outages/count</code>	Get the number of outages. (Returns plaintext, rather than <i>XML</i> or <i>JSON</i> .)
<code>/outages/{id}</code>	Get the outage specified by the given <i>ID</i> .
<code>/outages/forNode/{nodeId}</code>	Get the outages that match the given node <i>ID</i> .

2.8.15. Requisitions

RESTful service to the OpenNMS Horizon Provisioning Requisitions. In this *API*, these "groups" of nodes are aptly named and treated as requisitions.

This current implementation supports *CRUD* operations for managing provisioning requisitions. Requisitions are first *POSTed* and no provisioning (import/synchronize) operations are taken. This is done so that a) the *XML* can be verified and b) so that the operations can happen at a later time. They are moved to the deployed state (put in the active requisition repository) when an import is run.

If a request says that it gets the *active* requisition, that means it returns the pending requisition (being edited for deployment) if there is one, otherwise it returns the deployed requisition. Note that anything that says it *adds/deletes/modifies* a *node*, *interface*, etc. in these instructions is referring to modifying that element from the *requisition* not from the database itself. That will happen upon import/synchronization.

You may write requisition data if the authenticated user is in the *provision*, *rest*, or *admin* roles.

GETs (Reading Data)

Resource	Description
<code>/requisitions</code>	Get all active requisitions.
<code>/requisitions/count</code>	Get the number of active requisitions. (Returns plaintext, rather than <i>XML</i> or <i>JSON</i> .)
<code>/requisitions/deployed</code>	Get the list of all deployed (active) requisitions.

Resource	Description
<code>/requisitions/deployed/count</code>	Get the number of deployed requisitions. (Returns plaintext, rather than XML or JSON.)
<code>/requisitions/{name}</code>	Get the active requisition for the given foreign source name.
<code>/requisitions/{name}/nodes</code>	Get the list of nodes being requisitioned for the given foreign source name.
<code>/requisitions/{name}/nodes/{foreignId}</code>	Get the node with the given foreign ID for the given foreign source name.
<code>/requisitions/{name}/nodes/{foreignId}/interfaces</code>	Get the interfaces for the node with the given foreign ID and foreign source name.
<code>/requisitions/{name}/nodes/{foreignId}/interfaces/{ipAddress}</code>	Get the interface with the given IP for the node with the specified foreign ID and foreign source name.
<code>/requisitions/{name}/nodes/{foreignId}/interfaces/{ipAddress}/services</code>	Get the services for the interface with the specified IP address, foreign ID, and foreign source name.
<code>/requisitions/{name}/nodes/{foreignId}/interfaces/{ipAddress}/services/{service}</code>	Get the given service with the specified IP address, foreign ID, and foreign source name.
<code>/requisitions/{name}/nodes/{foreignId}/categories</code>	Get the categories for the node with the given foreign ID and foreign source name.
<code>/requisitions/{name}/nodes/{foreignId}/categories/{categoryName}</code>	Get the category with the given name for the node with the specified foreign ID and foreign source name.
<code>/requisitions/{name}/nodes/{foreignId}/assets</code>	Get the assets for the node with the given foreign ID and foreign source name.
<code>/requisitions/{name}/nodes/{foreignId}/assets/{assetName}</code>	Get the value of the asset for the given assetName for the node with the given foreign ID and foreign source name.

POSTs (Adding Data)

Resource	Description
<code>/requisitions</code>	Adds (or replaces) a requisition.
<code>/requisitions/{name}/nodes</code>	Adds (or replaces) a node in the specified requisition. This operation can be very helpful when working with [[Large Requisitions]].
<code>/requisitions/{name}/nodes/{foreignId}/interfaces</code>	Adds (or replaces) an interface for the given node in the specified requisition.
<code>/requisitions/{name}/nodes/{foreignId}/interfaces/{ipAddress}/services</code>	Adds (or replaces) a service on the given interface in the specified requisition.
<code>/requisitions/{name}/nodes/{foreignId}/categories</code>	Adds (or replaces) a category for the given node in the specified requisition.
<code>/requisitions/{name}/nodes/{foreignId}/assets</code>	Adds (or replaces) an asset for the given node in the specified requisition.

PUTs (Modifying Data)

Resource	Description
<code>/requisitions/{name}/import</code>	Performs an import/synchronize on the specified foreign source. This turns the "active" requisition into the "deployed" requisition.
<code>/requisitions/{name}/import?rescanExisting=false</code>	Performs an import/synchronize on the specified foreign source. This turns the "active" requisition into the "deployed" requisition. Existing nodes will not be scanned until the next rescan interval, only newly-added nodes will be. Useful if you're planning on making a series of changes.

Resource	Description
<code>/requisitions/{name}</code>	Update the specified foreign source.
<code>/requisitions/{name}/nodes/{foreignId}</code>	Update the specified node for the given foreign source.
<code>/requisitions/{name}/nodes/{foreignId}/interfaces/{ipAddress}</code>	Update the specified IP address for the given node and foreign source.

DELETES (Removing Data)

Resource	Description
<code>/requisitions/{name}</code>	Delete the pending requisition for the named foreign source.
<code>/requisitions/deployed/{name}</code>	Delete the active requisition for the named foreign source.
<code>/requisitions/{name}/nodes/{foreignId}</code>	Delete the node with the given foreign <i>ID</i> from the given requisition.
<code>/requisitions/{name}/nodes/{foreignId}/interfaces/{ipAddress}</code>	Delete the IP address from the requisitioned node with the given foreign <i>ID</i> and foreign source.
<code>/requisitions/{name}/nodes/{foreignId}/interfaces/{ipAddress}/services/{service}</code>	Delete the service from the requisitioned interface with the given IP address, foreign <i>ID</i> and foreign source.
<code>/requisitions/{name}/nodes/{foreignId}/categories/{category}</code>	Delete the category from the node with the given foreign <i>ID</i> and foreign source.
<code>/requisitions/{name}/nodes/{foreignId}/assets/{field}</code>	Delete the field from the requisition's nodes asset with the given foreign <i>ID</i> and foreign source.

2.8.16. Resources API

The *Resources API* can be used to list or delete resources at the node level and below. This service is especially useful in conjunction with the *Measurements API*.

GETs (Reading Data)

Resource	Description
<code>/resources</code>	Retrieve the full tree of resources in the system (expensive, use with care)
<code>/resources/{resourceId}</code>	Retrieve the tree of resources starting with the named resource ID
<code>/resources/fornode/{nodeCriteria}</code>	Retrieve the tree of resources for a node, given its database ID or <code>foreign-source:foreign-ID</code> tuple

DELETES (Removing Data)

Resource	Description
<code>/resources/{resourceId}</code>	Delete resource with the named resource ID, and all its child resources, if any

The following table shows all supported query string parameters and their default values.

name	default	comment
depth	varies	GET only. Limits the tree depth for retrieved resources. Defaults to 1 when listing all resources, or to -1 (no limit) when listing a single resource.

Usage examples with curl

Retrieve the tree of resources rooted at the node with database ID 1, by resource ID

```
curl -u admin:admin "http://127.0.0.1:8980/opennms/rest/resources/node%5B1%5D"
```

Response

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<resource id="node[1]"
  label="anode"
  name="1"
  link="element/node.jsp?node=1"
  typeLabel="Node">
<children count="11" totalCount="11">
  <resource id="node[1].nodeSnmplib["
    label="Node-level Performance Data"
    name=""
    typeLabel="SNMP Node Data"
    parentId="node[1]">
<children/>
<stringPropertyAttributes/>
<externalValueAttributes/>
<rrdGraphAttributes>
  <entry>
    <key>loadavg1</key>
    <value name="loadavg1"
      relativePath="snmp/1"
      rrdFile="loadavg1.jrb"/>
  </entry>
  <key>tcpActiveOpens</key>
  <value name="tcpActiveOpens"
    relativePath="snmp/1"
    rrdFile="tcpActiveOpens.jrb"/>
  </entry>
  <entry>
    <key>memTotalFree</key>
    <value name="memTotalFree"
      relativePath="snmp/1"
      rrdFile="memTotalFree.jrb"/>
  </entry>
  ...
</rrdGraphAttributes>
</resource>
<resource id="node[1].interfaceSnmplib["lo]"
  label="lo (10 Mbps)"
  name="lo"
  link="element/snmpinterface.jsp?node=1&ifindex=1"
  typeLabel="SNMP Interface Data"
  parentId="node[1]">
<children/>
<stringPropertyAttributes>
  <entry>
    <key>ifName</key>
    <value>lo</value>
  </entry>
  ...
</stringPropertyAttributes>
<externalValueAttributes>
  <entry>
    <key>ifSpeed</key>
    <value>10000000</value>
  </entry>
  <entry>
    <key>ifSpeedFriendly</key>
    <value>10 Mbps</value>
  </entry>
</externalValueAttributes>
```

```

<rrdGraphAttributes>
  ...
  <entry>
    <key>ifHCInOctets</key>
    <value name="ifHCInOctets"
           relativePath="snmp/1/lo"
           rrdFile="ifHCInOctets.jrb"/>
  </entry>
  <entry>
    <key>ifHCOutOctets</key>
    <value name="ifHCOutOctets"
           relativePath="snmp/1/lo"
           rrdFile="ifHCOutOctets.jrb"/>
  </entry>
  ...
</rrdGraphAttributes>
</resource>
...
</children>
<stringPropertyAttributes/>
<externalValueAttributes/>
<rrdGraphAttributes/>
</resource>

```

Retrieve the tree of resources rooted at the node with database ID 1, without having to construct a resource ID

```
curl -u admin:admin "http://127.0.0.1:8980/opennms/rest/resources/fornode/1"
```

Retrieve the tree of resources rooted at the node with foreign-ID node42 in requisition Servers, by resource ID

```
curl -u admin:admin "http://127.0.0.1:8980/opennms/rest/resources/nodeSource%5BServers:node42%5D"
```

Retrieve the tree of resources rooted at the node with foreign-ID node42 in requisition Servers, without having to construct a resource ID

```
curl -u admin:admin "http://127.0.0.1:8980/opennms/rest/resources/fornode/Servers:node42"
```

2.8.17. Realtime Console data

The *Realtime Console (RTC)* calculates the availability for monitored services. Data provided from the *RTC* is available to the ReST API.

GETs (Reading Data)

Resource	Description
/availability/categories/{category}	Get all nodes and availability data from a given SLA category filter, i.e. Web Servers (Web+Servers)
/availability/categories/{category}/nodes	Get node availability data for each node of a given SLA category filter
/availability/categories/{category}/nodes/{nodeId}	Get detailed service availability for a given node in a given SLA category filter
/availability/nodes/{nodeId}	Get detailed availability for all services on a given node

Example

```
curl -u demo:demo http://demo.opennms.org/opennms/rest/availability/categories/Web+Servers
curl -u demo:demo http://demo.opennms.org/opennms/rest/availability/categories/nodes
curl -u demo:demo http://demo.opennms.org/opennms/rest/availability/categories/nodes/31
curl -u demo:demo http://demo.opennms.org/opennms/rest/availability/nodes/31
```

2.8.18. Scheduled Outages

GETs (Reading Data)

Parameter	Description
<code>/sched-outages</code>	to get a list of configured scheduled outages.
<code>/sched-outages/{outageName}</code>	to get the details of a specific outage.

POSTs (Setting Data)

Parameter	Description
<code>/sched-outages</code>	to add a new outage (or update an existing one).

PUTs (Modifying Data)

Parameter	Description
<code>/sched-outages/{outageName}/collectd/{package}</code>	to add a specific outage to a collectd's package.
<code>/sched-outages/{outageName}/pollerd/{package}</code>	to add a specific outage to a pollerd's package.
<code>/sched-outages/{outageName}/threshd/{package}</code>	to add a specific outage to a threshd's package.
<code>/sched-outages/{outageName}/notifd</code>	to add a specific outage to the notifications.

DELETEs (Removing Data)

Parameter	Description
<code>/sched-outages/{outageName}</code>	to delete a specific outage.
<code>/sched-outages/{outageName}/collectd/{package}</code>	to remove a specific outage from a collectd's package.
<code>/sched-outages/{outageName}/pollerd/{package}</code>	to remove a specific outage from a pollerd's package.
<code>/sched-outages/{outageName}/threshd/{package}</code>	to remove a specific outage from a threshd's package.
<code>/sched-outages/{outageName}/notifd</code>	to remove a specific outage from the notifications.

2.8.19. SNMP Configuration

You can edit the community string, SNMP version, etc. for an IP address using this interface. If you make a change that would overlap with an existing `snmp-config.xml`, it will automatically create groups of `<definition />` entries as necessary. If no `<definition />` entry is created it matches the defaults.

There are different versions of the interface (see below). The following operations are supported:

GETs (Reading Data)

Parameter	Description
<code>/snmpConfig/{ipAddress}</code>	Get the SNMP configuration for a given IP address.

PUTs (Modifying Data)

Parameter	Description
<code>/snmpConfig/{ipAddress}</code>	Add or update the SNMP configuration for a given IP address.

Determine API version

To determine the version of the API running in your OpenNMS Horizon type <http://localhost:8980/opennms/rest/snmpConfig/1.1.1.1> in your browser and have a look at the output:

- **Version 1:** If the output only have attributes `community`, `port`, `retries`, `timeout` and `version`

- **Version 2:** If there are more attributes than described before (e.g. max Repetitions)

API Version 1

In version 1 only a few attributes defined in `snmp-config.xsd` are supported. These are defined in `snmp-info.xsd`:

```
<xs:schema
  xmlns:tns="http://xmlns.opennms.org/xsd/config/snmp-info"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  version="1.0"
  targetNamespace="http://xmlns.opennms.org/xsd/config/snmp-info">
  <xs:element name="snmp-info" type="tns:snmpInfo"/>
  <xs:complexType name="snmpInfo">
    <xs:sequence>
      <xs:element name="community" type="xs:string" minOccurs="0"/>
      <xs:element name="port" type="xs:int"/>
      <xs:element name="retries" type="xs:int"/>
      <xs:element name="timeout" type="xs:int"/>
      <xs:element name="version" type="xs:string" minOccurs="0"/>
    </xs:sequence>
  </xs:complexType>
</xs:schema>
```

The following table shows all supported attributes, optional restrictions and the mapping between `snmp-info.xsd` and `snmp-config.xsd`. All parameters can be set regardless the version.

attribute snmp-info.xml	attribute snmp-config.xml	default	restricted to version	restriction
version	version	v1	-	"v1", "v2c" or "v3" are valid arguments. If an invalid or empty argument is provided "v1" is used.
port	port	161	-	Integer > 0
retries	retry	1	-	Integer > 0
timeout	timeout	3000	-	Integer > 0
community	read-community	public	-	any string with a length >= 1

Example 1:

```
curl -v -X PUT -H "Content-Type: application/xml" \
  -H "Accept: application/xml" \
  -d "<snmp-info>
    <community>yRuSonoZ</community>
    <port>161</port>
    <retries>1</retries>
    <timeout>2000</timeout>
    <version>v2c</version>
  </snmp-info>" \
  -u admin:admin http://localhost:8980/opennms/rest/snmpConfig/10.1.1.1
```

Creates or updates a `<definition/>`-entry for IP address 10.1.1.1 in `snmp-config.xml`.

Example 2:

```
curl -v -X GET -u admin:admin http://localhost:8980/opennms/rest/snmpConfig/10.1.1.1
```

Returns the SNMP configuration for IP address 10.1.1.1 as defined in example 1.

API Version 2

Since Version 2 all attributes of a `<definition />` entry defined in `snmp-config.xsd` (<http://xmlns.opennms.org/xsd/config/snmp>) can be set or get via the interface - except it is only possible to set the configuration for one IP address and not for a range of IP addresses. This may change in the future.

The interface uses *SnmInfo* objects for communication. Therefore it is possible to set for example v1 and v3 parameters in one request (e.g. `readCommunity` String and `privProtocol` String). However OpenNMS Horizon does not allow this. It is only allowed to set attributes which have no version restriction (e.g. timeout value) or the attributes which are limited to the version (e.g. `readCommunity` String if version is v1/v2c). The same is for getting data from the API, even if it is possible to store v1 and v3 parameters in one definition block in the `snmp-config.xml` manually, the *ReST API* will only return the parameters which match the version. If no version is defined, the default is assumed (both in *PUT* and *GET* requests).

The *SnmInfo* schema is defined as follows:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<xs:schema
  elementFormDefault="qualified"
  version="1.0"
  targetNamespace="http://xmlns.opennms.org/xsd/config/snmp-info"
  xmlns:tns="http://xmlns.opennms.org/xsd/config/snmp-info"
  xmlns:xs="http://www.w3.org/2001/XMLSchema">

  <xs:element name="snmp-info" type="tns:snmpInfo"/>
  <xs:complexType name="snmpInfo">
    <xs:sequence>
      <xs:element name="authPassPhrase" type="xs:string" minOccurs="0"/>
      <xs:element name="authProtocol" type="xs:string" minOccurs="0"/>
      <xs:element name="community" type="xs:string" minOccurs="0"/>
      <xs:element name="contextEngineId" type="xs:string" minOccurs="0"/>
      <xs:element name="contextName" type="xs:string" minOccurs="0"/>
      <xs:element name="engineId" type="xs:string" minOccurs="0"/>
      <xs:element name="enterpriseId" type="xs:string" minOccurs="0"/>
      <xs:element name="maxRepetitions" type="xs:int" minOccurs="0"/>
      <xs:element name="maxRequestSize" type="xs:int" minOccurs="0"/>
      <xs:element name="maxVarsPerPdu" type="xs:int" minOccurs="0"/>
      <xs:element name="port" type="xs:int" minOccurs="0"/>
      <xs:element name="privPassPhrase" type="xs:string" minOccurs="0"/>
      <xs:element name="privProtocol" type="xs:string" minOccurs="0"/>
      <xs:element name="proxyHost" type="xs:string" minOccurs="0"/>
      <xs:element name="readCommunity" type="xs:string" minOccurs="0"/>
      <xs:element name="retries" type="xs:int" minOccurs="0"/>
      <xs:element name="securityLevel" type="xs:int" minOccurs="0"/>
      <xs:element name="securityName" type="xs:string" minOccurs="0"/>
      <xs:element name="timeout" type="xs:int" minOccurs="0"/>
      <xs:element name="version" type="xs:string" minOccurs="0"/>
      <xs:element name="writeCommunity" type="xs:string" minOccurs="0"/>
    </xs:sequence>
  </xs:complexType>
</xs:schema>
```

The following table shows all supported attributes, the mapping between `snmp-info.xsd` and `snmp-config.xsd`. It also shows the version limitations, default values and the restrictions - if any.

attribute snmp-info.xml	attribute snmp-config.xml
default	restricted to version
restriction	version
version	v1
-	"v1", "v2c" or "v3" are valid arguments. If an invalid or empty argument is provided "v1" is used.
port	port
161	-
Integer > 0	retries
retry	1
-	Integer > 0
timeout	timeout
3000	-
Integer > 0	maxVarsPerPdu
max- vars-per- pdu	10
-	Integer > 0
maxRepetitions	max-repetitions
2	-
Integer > 0	maxRequestSize
max- request- size	65535
-	Integer > 0
proxyHost	proxy-host
	-
	readCommunity
read- community	public
v1, v2c	

attribute snmp-info.xml	attribute snmp-config.xml
writeCommunity	write-community
private	v1, v2c
	securityName
security-name	opennmsUser
v3	
securityLevel	security-level
noAuthNoPriv	v3

attribute snmp-info.xml	attribute snmp-config.xml
<p>Integer value, which can be null, 1, 2, or 3.</p> <ul style="list-style-type: none"> 1 means noAuthNoPriv2 means authNoPriv3 means authPriv <p>If you do not set the security level manually it is determined automatically:</p> <ul style="list-style-type: none"> if no authPassPhrase set the securityLevel is 1if a authPassPhrase and no privPassPhrase is set the security level is 2.if a authPassPhrase and a privPassPhrase is set the security level is 3. 	<p>authPassPhrase</p>
auth-passphrase	0p3nNMSv3

attribute snmp-info.xml	attribute snmp-config.xml
v3	
authProtocol	auth-protocol
MD5	v3
only MD5 or SHA are valid arguments	privPassPhrase
privacy-passphrase	0p3nNMSv3
v3	
privProtocol	privacy-protocol
DES	v3
only DES, AES, AES192 or AES256 are valid arguments.	engineId
engine-id	
v3	
contextEngineId	context-engine-id
	v3
	contextName
context-name	
v3	
enterpriseId	enterprise-id
	v3

Example 1:

```
curl -v -X PUT -H "Content-Type: application/xml" \  
-H "Accept: application/xml" \  
-d "<snmp-info\  
  <readCommunity>yRuSonoZ</readCommunity>  
  <port>161</port>  
  <retries>1</retries>  
  <timeout>2000</timeout>  
  <version>v2c</version>  
</snmp-info>" \  
-u admin:admin http://localhost:8980/opennms/rest/snmpConfig/10.1.1.1
```

Creates or updates a `<definition/>`-entry for IP address 10.1.1.1 in `snmp-config.xml`.

Example 2:

```
curl -v -X GET -u admin:admin http://localhost:8980/opennms/rest/snmpConfig/10.1.1.1
```

Returns the SNMP configuration for IP address 10.1.1.1 as defined in example 1.

Example 3:

```
curl -v -X PUT -H "Content-Type: application/xml" \  
-H "Accept: application/xml" \  
-d "<snmp-info\  
  <readCommunity>yRuSonoZ</readCommunity>  
  <port>161</port>  
  <retries>1</retries>  
  <timeout>2000</timeout>  
  <version>v1</version>  
  <securityName>secret-stuff</securityName>  
  <engineId>engineId</engineId>  
</snmp-info>" \  
-u admin:admin http://localhost:8980/opennms/rest/snmpConfig/10.1.1.1
```

Creates or updates a `<definition/>`-entry for IP address 10.1.1.1 in `snmp-config.xml` ignoring attributes `securityName` and `engineId`.

Example 4:

```
curl -v -X PUT -H "Content-Type: application/xml" \  
-H "Accept: application/xml" \  
-d "<snmp-info\  
  <readCommunity>yRuSonoZ</readCommunity>  
  <port>161</port>  
  <retries>1</retries>  
  <timeout>2000</timeout>  
  <version>v3</version>  
  <securityName>secret-stuff</securityName>  
  <engineId>engineId</engineId>  
</snmp-info>" \  
-u admin:admin http://localhost:8980/opennms/rest/snmpConfig/10.1.1.1
```

Creates or updates a `<definition/>`-entry for IP address 10.1.1.1 in `snmp-config.xml` ignoring attribute `readCommunity`.

2.8.20. Users

Since users are not currently stored in the database, the ReST interface for them is not as full-fledged as that of nodes, etc.

IMPORTANT

You cannot use hibernate criteria for filtering. You may need to touch the `$OPENNMS_HOME/etc/users.xml` file on the filesystem for any addition or modification actions to take effect (see [NMS-6469](#) for details).

GETs (Reading Data)

Parameter	Description
<code>/users</code>	Get a list of users.
<code>/users/{username}</code>	Get a specific user, by username.

POSTs (Adding Data)

Parameter	Description
<code>/users</code>	Add a user. If supplying a password it is assumed to be hashed or encrypted already, at least as of 1.12.5. To indicate that the supplied password uses the salted encryption algorithm rather than the older <i>MD5</i> based algorithm, you need to pass an element named <code>passwordSalt</code> with text <code>true</code> after the password element (or key/value pairs if using <i>JSON</i>).

PUTs (Modifying Data)

Parameter	Description
<code>/users/{username}</code>	Update an existing user's full-name, user-comments, password, passwordSalt and duty-schedule values.

2.8.21. SNMP Trap Northbender Interface Configuration

GETs (Reading Data)

Resource	Description
<code>/config/snmptrap-nbi</code>	Gets full content of the configuration.
<code>/config/snmptrap-nbi/status</code>	Gets the status of the SNMP Trap NBI (returns either true or false).
<code>/config/snmptrap-nbi/destinations</code>	Gets the name of all the existing destinations.
<code>/config/snmptrap-nbi/destinations/{name}</code>	Gets the content of the destination named {name}

PUTs (Update defaults)

On a successful request, the Syslog NBI will be notified about the configuration change.

Resource	Description
<code>/config/snmptrap-nbi/status?enabled=(true;false)</code>	Sets the status of the SNMP Trap NBI.

POSTs (Adding Data)

POST requires form data using application/x-www-form-urlencoded as a Content-Type.

On a successful request, the SNMP Trap NBI will be notified about the configuration change.

Resource	Description
<code>/config/snmptrap-nbi</code>	Updates the full content of the configuration.
<code>/config/snmptrap-nbi/destinations</code>	Adds a new or overrides an existing destination.

PUTs (Modifying Data)

PUT requires form data using application/x-www-form-urlencoded as a Content-Type.

On a successful request, the SNMP Trap NBI will be notified about the configuration change.

Resource	Description
<code>/config/snmptrap-nbi/destinations/{name}</code>	Updates the content of the destination named {name}

DELETES (Remove Data)

On a successful request, the SNMP Trap NBI will be notified about the configuration change.

Resource	Description
<code>/config/snmptrap-nbi/destinations/{name}</code>	Updates the content of the destination named {name}

2.8.22. Email Northbender Interface Configuration

GETs (Reading Data)

Resource	Description
<code>/config/email-nbi</code>	Gets full content of the configuration.
<code>/config/email-nbi/status</code>	Gets the status of the Email NBI (returns either true or false).
<code>/config/email-nbi/destinations</code>	Gets the name of all the existing destinations.
<code>/config/email-nbi/destinations/{name}</code>	Gets the content of the destination named {name}

PUTs (Update defaults)

On a successful request, the Email NBI will be notified about the configuration change.

Resource	Description
<code>/config/email-nbi/status?enabled=(true;false)</code>	Sets the status of the Email NBI.

POSTs (Adding Data)

POST requires form data using application/x-www-form-urlencoded as a Content-Type.

On a successful request, the Email NBI will be notified about the configuration change.

Resource	Description
<code>/config/email-nbi/destinations</code>	Adds a new or overrides an existing destination.

PUTs (Modifying Data)

PUT requires form data using `application/x-www-form-urlencoded` as a Content-Type.

On a successful request, the Email NBI will be notified about the configuration change.

Resource	Description
<code>/config/email-nbi/destinations/{name}</code>	Updates the content of the destination named {name}

DELETes (Remove Data)

On a successful request, the Email NBI will be notified about the configuration change.

Resource	Description
<code>/config/email-nbi/destinations/{name}</code>	Updates the content of the destination named {name}

2.8.23. Javamail Configuration

GETs (Reading Data)

Resource	Description
<code>/config/javamail/default/readmail</code>	Get the name of the default readmail config.
<code>/config/javamail/default/sendmail</code>	Get the name of the default sendmail config.
<code>/config/javamail/readmails</code>	Get the name of all the existing readmail configurations.
<code>/config/javamail/sendmails</code>	Get the name of all the existing sendmail configurations.
<code>/config/javamail/end2ends</code>	Get the name of all the existing end2end mail configurations.
<code>/config/javamail/readmails/{name}</code>	Get the content of the readmail configuration named {name}
<code>/config/javamail/sendmails/{name}</code>	Get the content of the sendmail configuration named {name}
<code>/config/javamail/end2ends/{name}</code>	Get the content of the end2end mail configuration named {name}

POSTs (Adding/Updating Data)

POST requires form data using `application/xml` or `application/json` as a Content-Type.

On a successful request, the Email NBI will be notified about the configuration change.

Resource	Description
<code>/config/javamail/readmails</code>	Adds a new or overrides an existing readmail configuration.
<code>/config/javamail/sendmails</code>	Adds a new or overrides an existing sendmail configuration.
<code>/config/javamail/end2ends</code>	Adds a new or overrides an existing end2ends mail configuration.

PUTs (Update defaults)

On a successful request, the Email NBI will be notified about the configuration change.

Resource	Description
<code>config/javamail/default/readmail/{name}</code>	Sets the readmail named {name} as the new default.
<code>config/javamail/default/sendmail/{name}</code>	Sets the sendmail named {name} as the new default.

PUTs (Modifying Data)

PUT requires form data using application/x-www-form-urlencoded as a Content-Type.

On a successful request, the Email NBI will be notified about the configuration change.

Resource	Description
<code>/config/javamail/readmails/{name}</code>	Updates the content of the readmail configuration named {name}
<code>/config/javamail/sendmails/{name}</code>	Updates the content of the sendmail configuration named {name}
<code>/config/javamail/end2ends/{name}</code>	Updates the content of the end2end mail configuration named {name}

DELETES (Remove Data)

On a successful request, the Email NBI will be notified about the configuration change.

Resource	Description
<code>/config/javamail/readmails/{name}</code>	Removes the readmail configuration named {name}
<code>/config/javamail/sendmails/{name}</code>	Removes the sendmail configuration named {name}
<code>/config/javamail/end2ends/{name}</code>	Removes the end2end mail configuration named {name}

2.8.24. Syslog Northbound Interface Configuration

GETs (Reading Data)

Resource	Description
<code>/config/syslog-nbi</code>	Gets full content of the configuration.
<code>/config/syslog-nbi/status</code>	Gets the status of the Syslog NBI (returns either true or false).
<code>/config/syslog-nbi/destinations</code>	Gets the name of all the existing destinations.
<code>/config/syslog-nbi/destinations/{name}</code>	Gets the content of the destination named {name}

PUTs (Update defaults)

On a successful request, the Syslog NBI will be notified about the configuration change.

Resource	Description
<code>/config/syslog-nbi/status?enabled=(true;false)</code>	Sets the status of the Syslog NBI.

POSTs (Adding Data)

POST requires form data using application/x-www-form-urlencoded as a Content-Type.

On a successful request, the Syslog NBI will be notified about the configuration change.

Resource	Description
<code>/config/syslog-nbi</code>	Updates the full content of the configuration.
<code>/config/syslog-nbi/destinations</code>	Adds a new or overrides an existing destination.

PUTs (Modifying Data)

PUT requires form data using application/x-www-form-urlencoded as a Content-Type.

On a successful request, the Syslog NBI will be notified about the configuration change.

Resource	Description
<code>/config/syslog-nbi/destinations/{name}</code>	Updates the content of the destination named {name}

DELETES (Remove Data)

On a successful request, the Syslog NBI will be notified about the configuration change.

Resource	Description
<code>/config/syslog-nbi/destinations/{name}</code>	Updates the content of the destination named {name}

2.9. ReST API Examples

2.9.1. Getting Graph data

While graphs aren't technically available via *ReST*, you can parse some *ReST* variables to get enough data to pull a graph. This isn't ideal because it requires multiple fetches, but depending on your use case, this may be adequate for you.

I'm in-lining some sample *PHP* code which should do this (not tested at all, cut & paste from old code I have that does not use the *ReST*- interface, and/or coded straight into the browser so *YMMV*). If you go to your *NMS* and click the resource graphs, then right click the graph you want and hit *_View Image* you will get the full *URL* that would need to be passed to pull that graph as a standalone image.

From that just take the *URL* and plug in the values you pulled from *ReST* to get a graph for whatever node you wanted.

```

function fetchit($thing, $user = "user", $pass = "pass") {
    $url = "http://localhost:8980/opennms";
    $ch = curl_init();
    curl_setopt($ch, CURLOPT_URL, $url . $thing);
    curl_setopt($ch, CURLOPT_HEADER, 0);
    curl_setopt($ch, CURLOPT_RETURNTRANSFER, 1);
    curl_setopt($ch, CURLOPT_USERAGENT, $useragent);
    curl_setopt($ch, CURLOPT_USERPWD, $user.':'.$pass);
    $data = curl_exec($ch);
    curl_close($ch);
    return $data;
}

// this assumes you already have found the nodeId via a previous REST call or some other means. Provided
// more as an example than what you might want.
function getNodeInterfaces($nodeId) {
    $data = fetchit("/rest/nodes/$nodeId/snmpinterfaces");
    return simplexml_load_string($data);
}

function fetchGraphs($nodeId) {
    $ints = getNodeInterfaces($nodeId);
    $chars = array('/', '.', ':', '-', ' ');
    $endtime = time();
    $starttime = (string)(time() - ($days * 24 * 60 * 60)) ;

    // use bcmath or a better version of PHP if you don't want this hypocrisy here.
    $endtime = $endtime . '000';
    $starttime = $starttime . '000';

    for($i=0; $i<count($ints->snmpInterfaces); $i++) {
        $ifname = $ints->snmpInterfaces[$i]->snmpInterface->ifName;
        $mac = $ints->snmpInterfaces[$i]->snmpInterface->physAddr;
        $if = str_replace($chars, "_", $ifname);
        if ( strlen(trim($mac)) < 12 ) { $mac_and_if = $if; } else { $mac_and_if = $if .'-' . $mac; };

        $image = fetchit("$url/graph/graph.png?resource=node[$nodeId].interfaceSnmp[
$mac_and_if]&report=report=mib2.HCbits&start=$starttime&end=$endtime");
        // you can pop this to a file now, or set header('Content-type: image/png'); then print
        "$image";
    }
}

```

2.9.2. provision.pl examples and notes

One way to test out the new *ReST* interface is to use `provision.pl`. If you run it you'll get a summary of the output, but it's not totally obvious how it all works.

Here is an example of adding a new node using the *ReST* interface:

```
# add a new foreign source called ubr
/usr/share/opennms/bin/provision.pl requisition add ubr
/usr/share/opennms/bin/provision.pl node add ubr 10341111 clownbox
/usr/share/opennms/bin/provision.pl node set ubr 10341111 city clownville
/usr/share/opennms/bin/provision.pl node set ubr 10341111 building clown-town-hall
/usr/share/opennms/bin/provision.pl node set ubr 10341111 parent-foreign-id 1122114
/usr/share/opennms/bin/provision.pl interface add ubr 10341111 10.1.3.4

# this is like a commit. No changes will take effect until you import a foreign source
/usr/share/opennms/bin/provision.pl requisition import ubr
```

You will probably need to specify the username/password of an admin. To do this add:

```
--username=admin --password=clovnms
```

to the command line.

2.9.3. Debian (Lenny) Notes

For Lenny, you'll need to pull a package out of backports to make everything work right. Read <http://backports.org/dokuwiki/doku.php?id=instructions> for instructions on adding it to `sources.list`.

```
# install liburi-perl from backports
sudo apt-get -t lenny-backports install liburi-perl
```

2.9.4. Windows Powershell ReST

Example of using *Windows Powershell* to fill some asset fields with *ReST*.


```

# Installdate of Windows
$wmi = Get-WmiObject -Class Win32_OperatingSystem
$dateInstalled = $wmi.ConvertToDateTime($wmi.InstallDate)

# Serialnumber and manufacturer of server
Get-WmiObject win32_bios | select SerialNumber
$wmi = Get-WmiObject -Class win32_bios
$manufacturer = $wmi.Manufacturer

# Text file with a description of the server for the comments field
$comment = Get-Content "C:\Program Files\BGInfo\Info_Description.txt" | Out-String

$user = "admin"
$pass= "admin"

$secpasswd = ConvertTo-SecureString $user -AsPlainText -Force
$cred = New-Object System.Management.Automation.PSCredential ($pass, $secpasswd)

$nodeid = Invoke-RestMethod -Uri
http://opennms.domain.nl:8980/opennms/rest/nodes?label=servername.domain.nl -Credential $cred
$nodeid = $nodeid.nodes.node.id

$uri="http://opennms.domain.nl:8980/opennms/rest/nodes/$nodeid/assetRecord"

Invoke-RestMethod -Uri
"http://opennms.massxess.nl:8980/opennms/rest/nodes/$nodeid/assetRecord/?building=133" -Credential $cred
-Method PUT
Invoke-RestMethod -Uri "$uri/?manufacturer=$manufacturer" -Credential $cred -Method PUT
Invoke-RestMethod -Uri "$uri/?dateInstalled=$dateInstalled" -Credential $cred -Method PUT
Invoke-RestMethod -Uri "$uri/?comment=$comment" -Credential $cred -Method PUT

```

Chapter 3. Develop Documentation

This document is the guideline for people who wish to contribute to writing documentation for the OpenNMS project. The OpenNMS software is free and open source, contribution of any kind is welcome. We ask that you observe the rules and guidelines outlined here to maintain consistency across the project.

Each (sub)project is represented as a section of the documentation. Each section will produce a HTML output in the file system that is generated in the `target/generated` sources folder.

The chosen file format for documentation is AsciiDoc ([AsciiDoc Homepage](#)). Document files use the `.adoc` file extension.

Note that there are different ways to contribute documentation, each suitable for the different use cases:

- Tutorials and How To's should be published on the [OpenNMS Wiki](#). For example: you want to describe how to use the Net-SNMP agent and the SNMP monitor from OpenNMS to solve a special use case with OpenNMS.
- The documentation in the source code should be formal technical documentation. The writing style should be accurate and concise. However, ensure that you explain concepts in detail and do not make omissions.

3.1. File Structure in opennms-doc

Directory	Contents
<code>guide-user/</code>	module with the guide for OpenNMS user e.g. NOC user who don't change behavior of OpenNMS.
<code>guide-admin/</code>	module with the guide for administrators configuring, optimizing and running OpenNMS
<code>guide-development/</code>	module with the guide for those who want to develop OpenNMS
<code>guide-install/</code>	module with the guide of how to install OpenNMS on different operating systems
<code>releasenotes/</code>	module with the changelog and release notes

3.2. Writing

The following rules will help you to commit correctly formatted and prepared documentation for inclusion in the OpenNMS project. It is important that we maintain a level of consistency across all of our committers and the documentation they produce.

When writing place a single sentence on each line. This makes it easy to move content around, and also easy to spot long, or fragmented, sentences. This will also allow us to assign comments on a sentence in GitHub which will facilitate easier merging.

NOTE Other than writing documentation, you can help out by providing comments on documentation, reviewing, suggesting improvements or reporting bugs. To do this head over to: [issue tracker for documentation!](#)

3.2.1. Conventions for text formatting

The following conventions are used:

- File names and path are written in ``poller-configuration.xml`` they will be rendered in: `poller-configuration.xml`;
- Names that indicate special attention, e.g. this configuration matches `*any*` entry: this is rendered as: this configuration matches **any** entry;
- `_Italics_` is rendered as *Italics* and used for emphasis and indicate internal names and abbreviations;
- `*Bold*` is rendered as **Bold** and should be used sparingly, for strong emphasis only;
- `+methodName()+` is rendered as `methodName()` and is also used for literals, (note: the content between the `+` signs *will* be parsed);
- ``command`` is rendered as `command` (typically used for command-line or parts used in configuration files), (note: the content between the ``` signs *will not* be parsed);
- ``my/path/`` is rendered as `my/path/` this is used for file names and paths;
- ```double quote"` (which is two grave accents to the left and two acute accents to the right) renders as ```double quote"`;
- ``single quote'` (which is a single grave accent to the left and a single acute accent to the right) renders as ``single quote'`.

3.2.2. Gotchas

- Always leave a blank line at the top of the documents section. It might be the title ends up in the last paragraph of the document;
- Start in line 2 setting a relative path to the images directory to picture rendering on GitHub:

```
// Allow image rendering
:imagesdir: relative/path/to/images/dir
```

- Always leave a blank line at the end of documents;
- As `{}` are used for AsciiDoc attributes, everything inside will be treated as an attribute. To avoid this you have to escape the opening brace: `\{`. If you do not escape the opening brace, the braces and the text inside them will be removed without any warning being issued!;
- Forcing line breaks can be achieved with `` + `` at the end of the line followed by a line break.

Example in source force line break

```
This is the first line +
and this a forced 2nd line
```

This is the first line
and this a forced 2nd line

3.3. Headings and document structure

Each document starts over with headings from level zero (the document title). Each document should have an id. In some cases sections in the document need to have id's as well, this depends on where they fit in the overall structure. If you wish to have a link to specific content that content has to have an id. A missing id in a mandatory place will cause the build to

fail.

To start a document:

```
[[unique-id-verbose-is-ok]]  
= The Document Title
```

If you are including the document inside another document and you need to push the headings down to the right level in the output, the `leveloffset` attribute is used.

Subsequent headings in a document should use the following syntax:

```
== Subheading  
... content here ...  
=== Subsubheading  
content here ...
```

3.4. Links

When you need to link to other parts of the manual you use the target id. To use a target id you follow this syntax:

```
<<doc-guidelines-links>>
```

This will render as: [\[doc-guidelines-links\]](#)

NOTE | To use the target id in you document simply write the target id in your text, for example:

see <<target-id>>

this should suffice for most cases.

If you need to link to another document with your own link text, then follow this procedure:

```
<<target-id, link text that fits in the context>>
```

NOTE | Having lots of linked text may work well in a web context but is a distracting in print. The documentation we are creating is intended for both mediums so be considerate of this in your usage.

If you wish to use an external link, they are are added as:

```
http://www.opennms.org/[Link text here]
```

This will render in the output as: [Link text here](#)

For short links it may be beneficial not to use accompanying link text:

```
http://www.opennms.org/
```

Which renders as: <http://www.opennms.org/>

NOTE | It is acceptable to have a period trailing after the URL, it will not render as a part of the link.

3.5. Admonitions and useful notes

These are useful for defining specific sections, such as Notes, Tips and Important information. We encourage the use of them in the documentation as long as they are used appropriately. Choose from the following:

Source template for making a note for additional hints

```
NOTE: This is my note.
```

This is how its rendered:

NOTE | This is my note.

Source for giving a tip

```
TIP: This is my tip.
```

This is how its rendered:

TIP | This is my tip.

Source for giving a important hint

```
IMPORTANT: This is my important hint.
```

This is how its rendered:

IMPORTANT | This is my important hint.

Source for giving a caution

```
CAUTION: This is my caution.
```

This is how its rendered:

CAUTION | This is my caution.

Source for giving a warning

```
WARNING: This is my warning.
```

This is how its rendered:

WARNING | This is my warning.

A multiline variation:

```
TIP: Tiptext. +
    Line 2.
```

Which is rendered as:

```
TIP | Tiptext.
      | Line 2.
```

```
NOTE | Remember to write these in full caps. There is no easy manner in which to add new admonitions, do not
      | create your own.
```

3.6. Attributes

Common attributes you can use in documents:

- `{opennms-version}` - rendered as "2016.1.1"

These can substitute part of URLs that point to, for example, API docs or source code. Note that `opennms-git-tag` also handles the case of `snapshot/master`.

Sample AsciiDoc attributes which can be used:

- `{docdir}` - root directory of the documents
- `{nbsp}` - non-breaking space

3.7. Comments

There's a separate build that includes comments. When the comments are used they show up with a yellow background. This build doesn't run by default, but after a normal build, you can use `make annotated` to create a build yourself. You can use the resulting 'annotated' page to search for content as the full manual is a single page.

To write a comment:

```
// this is a comment
```

Comments are not visible in the standard build. Comment blocks won't be included in the output of any build. The syntax for a comment block is:

```
////
Note that includes in here will still be processed, but not make it into the output.
That is, missing includes here will still break the build!
////
```

3.8. Tables

For representing structured information you can use tables. A table is constructed in the following manner:

```
[options="header, autowidth"]
|===
| Parameter      | Description                               | Required | Default value
| `myFirstParm` | my first long description                 | required | `myDefault`
| `myScndParm`  | my second long description               | required | `myDefault`
|===
```

This is rendered as:

Parameter	Description	Required	Default value
<code>myFirstParm</code>	my first long description	required	<code>myDefault</code>
<code>myScndParm</code>	my second long description	required	<code>myDefault</code>

NOTE

Please align your columns in the AsciiDoc source in order to give better readability when editing in text view. If you have a very long description, break at 120 characters and align the text to improve source readability.

```
[options="header, autowidth"]
|===
| Parameter      | Description                               | Required | Default value
| `basic-authentication` | Authentication credentials to perform basic authentication. Credentials should comply to http://www.rfc-editor.org/rfc/rfc1945.txt section 11.1, without the Base64 encoding part. That's: be a string made of the concatenation of:
| 1- the user ID; +
| 2- a colon; +
| 3- the password. +
| `header[0-9]+` | `basic-authentication` takes precedence over the `user` and `password` parameters. Additional headers to be sent along with the request. Example of valid parameter's names are `header0`, `header1` and `header180`. `header` is not a valid parameter name.
|===
```

Figure 1. Example in AsciiDoc source for very long table descriptions

this is rendered as:

Parameter	Description	Required	Default value
<code>basic-authentication</code>	Authentication credentials to perform basic authentication. Credentials should comply to RFC1945 section 11.1, without the Base64 encoding part. That's: be a string made of the concatenation of: 1- the user ID; 2- a colon; 3- the password. <code>basic-authentication</code> takes precedence over the <code>user</code> and <code>password</code> parameters.	optional	-
<code>header[0-9]+</code>	Additional headers to be sent along with the request. Example of valid parameter's names are <code>header0</code> , <code>header1</code> and <code>header180</code> . <code>header</code> is not a valid parameter name.	optional	-


3.9. Include images

When visualizing complex problems you can help the explanation and provide greater information by using an image. We

use in OpenNMS documentation modules two directories for images.

The image folder structure mirrors the text structure. In this case it is a little bit easier to locate the AsciiDoc text file where the image is included.

Example folder structure for image files



- ① This folder contains all documentation modules;
- ② The module for this documentation for target group of documentation contributors;
- ③ Indicates a source folder;
- ④ The documentation root folder;
- ⑤ Folder for images. Images should be *.png or *.jpg if included in the documentation;
- ⑥ The image used, the format is a leading <number>_ followed by a name using no spaces;
- ⑦ Some images are created from tools like *yED*, this folder should contain the editable version of the file with the same file name;
- ⑧ Editable version of the image source file, note no spaces in the name;
- ⑨ Main document file which includes all documentation parts and is rendered as `index.html` for the web;
- ⑩ AsciiDoc source file which can include images;
- ⑪ Target folder with generated HTML output after `mvn clean package` has been performed;

IMPORTANT

All images in the entire manual share the same namespace, it is therefore best practice to use unique identifiers for images.

To include an image file, make sure that it resides in the 'images/' directory relative to the document you're including it within. Then use the following syntax for inclusion in the document:

First included image

```
.This is a caption of the image  
image::docs/02_opennms-logo.png[]
```

Which is rendered as:



Figure 2. This is a caption of the image

IMPORTANT

The image path for the images you include is relative to the *.adoc source file, where you use the image.

3.10. Code Snippets

You can include code snippets, configuration- or source code files in the documentation. You can enable syntax highlighting by providing the given language parameter, this will work on source code or configuration.

3.10.1. Explicitly defined in the document

WARNING

be careful to use this kind of code snippets as sparsely as possible. Code becomes obsolete very quickly, archaic usage practices are detrimental.

if you do wish to include snippets use the following method:

This is a sample configuration explicitly in the documentation

```
<service name="DNS" interval="300000" user-defined="false" status="on">  
  <parameter key="retry" value="2" />  
  <parameter key="timeout" value="5000" />  
  <parameter key="port" value="53" />  
  <parameter key="lookup" value="localhost" />  
  <parameter key="fatal-response-codes" value="2,3,5" /><!-- ServFail, NXDomain, Refused -->  
  <parameter key="rrd-repository" value="/opt/opennms/share/rrd/response" />  
  <parameter key="rrd-base-name" value="dns" />  
  <parameter key="ds-name" value="dns" />  
</service>
```

If there's no suitable syntax highlighter for the code used just omit the language: [source].

Currently the following syntax highlighters are enabled:

- Bash
- Groovy
- Java
- JavaScript

- Python
- XML

For other highlighters that could be added see <https://code.google.com/p/google-code-prettify/>.

3.10.2. Included from an example file

You can include source or configuration from an external file. In this way you can provide a working example configuration maintaining doc and example at the same time. The procedure and rules are the same as with images, the path is relative to the *.adoc file where the file to be used is included.

Include complete external file

```
[source,xml]
----
include::../configs/wmi-config.xml[]
----
```

This is how it's rendered:

```
<?xml version="1.0"?>
<wmi-config retry="2" timeout="1500"
username="Administrator" domain="WORKGROUP" password="password">
</wmi-config>
```

3.10.3. Include parts of a file

If you want to include just a specific segment of a large configuration file, you can assign tags that indicate to AsciiDoc the section that is to be included. In this example just the service definition of the *ICMP monitor* should be included.

In the 'poller-configuration.xml' tag the section in the following manner:

```
...
<rrd step="300">
  <rra>RRA:AVERAGE:0.5:1:2016</rra>
  <rra>RRA:AVERAGE:0.5:12:1488</rra>
  <rra>RRA:AVERAGE:0.5:288:366</rra>
  <rra>RRA:MAX:0.5:288:366</rra>
  <rra>RRA:MIN:0.5:288:366</rra>
</rrd>
<!-- # tag::IcmpServiceConfig[] -->
<service name="ICMP" interval="300000" user-defined="false" status="on">
  <parameter key="retry" value="2" />
  <parameter key="timeout" value="3000" />
  <parameter key="rrd-repository" value="/opt/opennms/share/rrd/response" />
  <parameter key="rrd-base-name" value="icmp" />
  <parameter key="ds-name" value="icmp" />
</service>
<!-- # end::IcmpServiceConfig[] -->
<service name="DNS" interval="300000" user-defined="false" status="on">
  <parameter key="retry" value="2" />
  <parameter key="timeout" value="5000" />
  <parameter key="port" value="53" />
...

```

Include this tagged part in the documentation using the tag parameter

```
[source,xml]
----
include::../configs/poller-configuration.xml[tags=IcmpServiceConfig]
----
```

This is how it rendered

```
<service name="ICMP" interval="300000" user-defined="false" status="on">
  <parameter key="retry" value="2" />
  <parameter key="timeout" value="3000" />
  <parameter key="rrd-repository" value="/opt/opennms/share/rrd/response" />
  <parameter key="rrd-base-name" value="icmp" />
  <parameter key="ds-name" value="icmp" />
</service>
```

NOTE | Spaces and tabs are taken from the original file.

3.11. Cheat Sheets and additional hints

For instructions on how to build your own version of the manual:

- [readme](#)

The documentation uses the AsciiDoc format. There are a number of guides that will help you to get started with using AsciiDoc:

- [Aciidoc Reference](#)
- [AsciiDoc FAQ](#)
- [AsciiDoc cheatsheet](#)
- [AsciiDoc Cheatsheet](#)

For other resources, to gain familiarity with AsciiDoc, you can visit:

- [AsciiDoc User Manual](#)
- [AsciiDoc Maven Plugin](#)
- [AsciiDoc discussion list](#)
- [AsciiDoc issue tracker](#)
- [Docbook to AsciiDoc](#)
- [How to create handsome PDF documents without frustration](#)

Chapter 4. AMQP Integration

The AMQP Integration allows external systems to communicate with the event bus of OpenNMS Horizon and receive alarms via the AMQP protocol.

NOTE | *AMQP* is standard messaging protocol supported by a number of brokers including *ActiveMQ* and *QPID*.

The integration is written using Camel + OSGi and has the following components:

- Event Forwarder
- Event Receiver
- Alarm Northbounder

Custom filtering (i.e. which events to forward) and transformations (i.e. how the events are represented in the messages) can be used in each of the components. Generic implementations

The integration is written using Camel + OSGi and exposes interfaces through which events and alarms can be filtered and/or transformed. The features are described in detail below.

NOTE | Each component can be configured and setup independently, i.e. you can choose to only forward alarms.

4.1. Event Forwarder

The event forwarder listens for *all* events on the internal event bus of OpenNMS Horizon. Events from the bus are sent to a Camel processor, which can filter or transform these, before being sent to the AMQP endpoint.

The event forwarder exposes the following properties via the `org.opennms.features.amqp.eventforwarder` pid:

Property	Default	Description
connectionUrl	amqp://guest:guest@onms/test?brokerlist='tcp://127.0.0.1:5672'	Used by the AMQPConnectionFactory. See ConnectionURL for a full list of options.
destination	amqp:OpenNMS-Exchange/opennms-routing-key	Target queue or topic. See AMQP for details.
processorName	default-event-forwarder-processor	Name <code>org.apache.camel.Processor</code> used to filter and/or format the events.

The default processor, the `default-event-forwarder-processor`, marshalls events to XML and does not perform any filtering. This means that when enabled, all events will be forwarded to the AMQP destination with XML strings as the message body.

4.1.1. Setup

Start by logging into a Karaf shell.

Update the properties with your deployment specific values:

```
config:edit org.opennms.features.amqp.eventforwarder
propset connectionUrl amqp://guest:guest@onms/test?brokerlist='tcp://127.0.0.1:5672\'
propset destination amqp:OpenNMS-Exchange/opennms-routing-key
propset processorName default-event-forwarder-processor
config:update
```

Install the feature:

```
features:install opennms-amqp-event-forwarder
```

4.1.2. Debugging

You can get detailed information on the Camel route using:

```
camel:route-info forwardEvent
```

4.2. Event Receiver

The event receiver listens for messages from an AMQP target and forwards them onto the internal event bus of OpenNMS Horizon. Messages are sent to a Camel processor, which can filter or transform these, before being sent onto the event bus.

The event forwarder exposes the following properties via the `org.opennms.features.amqp.eventreceiver` pid:

Property	Default	Description
connectionUrl	amqp://guest:guest@onms/test?brokerlist='tcp://127.0.0.1:5672'	Used by the AMQPConnectionFactory. See ConnectionURL for a full list of options.
source	amqp:OpenNMS-Queue	Source queue or topic. See AMQP for details.
processorName	default-event-receiver-processor	Name <code>org.apache.camel.Processor</code> used to filter and/or format the events.

The default processor, the `default-event-receiver-processor`, expects the message bodies to contain XML strings which are it unmarshalls to events.

4.2.1. Setup

Start by logging into a Karaf shell.

Update the properties with your deployment specific values:

```
config:edit org.opennms.features.amqp.eventreceiver
propset connectionUrl amqp://guest:guest@onms/test?brokerlist='tcp://127.0.0.1:5672\'
propset source amqp:OpenNMS-Queue
propset processorName default-event-receiver-processor
config:update
```

Install the feature:

```
features:install opennms-amqp-event-receiver
```

4.2.2. Debugging

You can get detailed information on the Camel route using:

```
camel:route-info receiveEvent
```

4.3. Alarm Northbounder

The alarm northbounder listens for *all* northbound alarms. Alarms are sent to a Camel processor, which can filter or transform these, before being sent to the AMQP endpoint.

The alarm northbounder exposes the following properties via the `org.opennms.features.amqp.alarmnorthbounder` pid:

Property	Default	Description
connectionUrl	amqp://guest:guest@onms/test?brokerlist='tcp://127.0.0.1:5672'	Used by the AMQPConnectionFactory. See ConnectionURL for a full list of options.
destination	amqp:OpenNMS-Exchange/opennms-routing-key	Target queue or topic. See AMQP for details.
processorName	default-alarm-northbounder-processor	Name <code>org.apache.camel.Processor</code> used to filter and/or format the events.

The default processor, the `default-alarm-northbounder-processor`, converts the alarms to a string and does not perform any filtering. This means that when enabled, all alarms will be forwarded to the AMQP destination with strings as the message body.

4.3.1. Setup

Start by logging into a Karaf shell.

Update the properties with your deployment specific values:

```
config:edit org.opennms.features.amqp.alarmnorthbounder
propset connectionUrl amqp://guest:guest@onms/test?brokerlist='tcp://127.0.0.1:5672\'
propset destination amqp:OpenNMS-Exchange/opennms-routing-key
propset processorName default-alarm-northbounder-processor
config:update
```

Install the feature:

```
features:install opennms-amqp-alarm-northbounder
```

4.3.2. Debugging

You can get detailed information on the Camel route using:

```
camel:route-info forwardAlarm
```

4.4. Custom Processors

If your integration requires specific filtering and or formatting, you can write your own processor by implementing the `org.apache.camel.Processor` interface.

For example, we can implement a custom processor used for event forwarding:

```
import org.apache.camel.Exchange;
import org.apache.camel.Processor;

public class MyEventProcessor implements Processor {
    @Override
    public void process(final Exchange exchange) throws Exception {
        final Event event = exchange.getIn().getBody(Event.class);

        // Filtering
        if (!shouldForward(event)) {
            exchange.setProperty(Exchange.ROUTE_STOP, Boolean.TRUE);
            return;
        }

        // Transforming
        MyDTO eventAsDTO = toDTO(event);
        exchange.getIn().setBody(eventAsDTO, MyDTO.class);
    }
}
```

In order to use the processor, package it as a bundle, and expose it to the OSGi service registry using:

```
<bean id="myEventProcessor" class="org.opennms.integrations.evilcorp.MyEventProcessor" />

<service id="myEventProcessorService" ref="myEventProcessor" interface="org.apache.camel.Processor">
  <service-properties>
    <entry key="name" value="evilcorp-event-forwarder-processor"/>
  </service-properties>
</service>
```

Once your bundle is in the Karaf container, you can update the loaded you can refer to your processor with:

```
config:edit org.opennms.features.amqp.eventforwarder
propset processorName evilcorp-event-forwarder-processor
config:update
```

If the event forwarder feature was already started, it should automatically restart and start using the new processor. Otherwise, you can start the feature with:

```
feature:install opennms-amqp-event-forwarder
```