
FogLAMP Documentation

Dianomic Systems

Sep 16, 2022

CONTENTS

1	Introduction to FogLAMP	1
1.1	Typical Use Cases	1
1.2	Architectural Overview	2
1.3	No-code/Low-code Development	2
2	Quick Start Guide	3
2.1	Installing FogLAMP	3
2.2	Starting and stopping FogLAMP	6
2.3	Troubleshooting FogLAMP	6
2.4	Running the FogLAMP GUI	6
2.5	Managing Data Sources	8
2.6	Viewing Data	11
2.7	Sending Data to Other Systems	15
2.8	PI Web API OMF Endpoint	16
2.9	Edge Data Store OMF Endpoint	19
2.10	OSIsoft Cloud Services OMF Endpoint	22
2.11	PI Connector Relay	24
2.12	Number Format Hints	31
2.13	Integer Format Hints	31
2.14	Type Name Hints	32
2.15	Type Hint	32
2.16	Tag Name Hint	32
2.17	Datapoint Specific Hint	32
2.18	Asset Framework Location Hint	32
2.19	Adding OMF Hints	33
2.20	Backing up and Restoring FogLAMP	33
2.21	Troubleshooting and Support Information	34
2.22	Package Uninstallation	35
3	Processing Data	37
3.1	Why Use Filters?	37
3.2	What Can Be Done?	37
3.3	Where Can it Be Done?	38
3.4	Some Useful Filters	48
4	FogLAMP Architecture	49
4.1	FogLAMP Core	50
4.2	Storage Layer	50
4.3	South Microservices	50
4.4	North Microservices	50

4.5	Filters	51
4.6	Event Service	52
4.7	Set Point Control Service	52
4.8	REST API	52
4.9	Graphical User Interface	52
5	Buffering & Storage	53
5.1	Configuring The Storage Plugin	54
5.2	Installing A PostgreSQL server	56
5.3	SQLite Plugin Configuration	56
5.4	Storage Management	58
6	Additional Services	61
6.1	Bucket Storage Service	61
6.2	Notifications Service	66
7	FogLAMP Control Features	83
7.1	Control Functions	83
7.2	Control Paths	83
7.3	Control Dispatcher Service	95
8	Plugin Documentation	113
8.1	FogLAMP South Plugins	113
8.2	FogLAMP North Plugins	239
8.3	FogLAMP Filter Plugins	279
8.4	FogLAMP Notification Rule Plugins	341
8.5	FogLAMP Notification Delivery Plugins	350
9	Developing Data Pipelines	383
9.1	Best Practices	383
10	Securing FogLAMP	395
10.1	Enabling HTTPS Encryption	395
10.2	Requiring User Login	397
10.3	User Management	401
10.4	Certificate Store	404
11	Tuning FogLAMP	407
11.1	South Service Advanced Configuration	407
11.2	North Advanced Configuration	409
11.3	Health Monitoring	410
11.4	Scheduler	411
11.5	Storage	411
12	Troubleshooting the PI Server integration	417
12.1	Log files	417
12.2	How to check the PI Web API is installed and running	418
12.3	Commands to check the PI WEB API	419
12.4	Error messages and causes	422
12.5	OMF Plugin Data	422
12.6	Possible solutions to common problems	426
13	Plugin Developer Guide	429
13.1	Plugins	429
13.2	Representing Data	432

13.3	Writing and Using Plugins	432
13.4	South Plugins	441
13.5	South Plugins in C	452
13.6	C++ Support Classes	463
13.7	Hybrid Plugins	472
13.8	North Plugins	473
13.9	Storage Plugins	481
13.10	Filter Plugins	484
13.11	Notification Delivery Plugins	498
13.12	Plugin Packaging	502
13.13	Testing Your Plugin	507
13.14	Developing with Windows Subsystem for Linux (WSL2)	515
14	REST API Developers Guide	521
14.1	The FogLAMP REST API	521
14.2	Administration API Reference	522
14.3	User API Reference	540
15	Building FogLAMP	545
15.1	Building Developers Guide	545
16	Version History	555
16.1	FogLAMP v2	555
16.2	FogLAMP v1	562
17	Downloads	589
17.1	Packages	589
18	OMF Kerberos Authentication	591
18.1	Introduction	591
18.2	PI Server as the North endpoint	591
18.3	North plugin	591
18.4	FogLAMP server configuration	592
19	FogLAMP Plugins	595
19.1	South Plugins	595
19.2	North Plugins	597
19.3	Filter Plugins	598
19.4	Notification Rule Plugins	599
19.5	Notification Delivery Plugins	600
20	Glossary	603
	Index	607

INTRODUCTION TO FOG LAMP

FogLAMP is an open Industrial IoT system designed to make collecting, filtering, processing and using operational data simpler and more open. Core to FogLAMP is an extensible microservice based architecture enabling any data to be read, processed and sent to any system. Coupled with this extensibility FogLAMP's Apache 2 license and community of developers results in an ever growing choice of components that can be used to solve your OT data needs well into the future.

FogLAMP provides a scalable, secure, robust infrastructure for collecting data from sensors, processing data at the edge using intelligent data pipelines and transporting data to historian and other management systems. FogLAMP also allows for edge based event detection and notification and control flows as a result of events, stimulus from upstream systems or user action. FogLAMP can operate over the unreliable, intermittent and low bandwidth connections often found in industrial or rugged environments.

1.1 Typical Use Cases

The depth and breadth of Industrial IoT use cases is considerable. FogLAMP is designed to address them. Below are some examples of typical FogLAMP deployments.

Unified data collection The industrial edge is one of the more challenging in computing. Today there are over 100 different protocols, no standards in machine data definitions, different types of data (time-series, vibration, array, image, radiometric, transactional, etc.), sensors producing bytes/hr to gigs/hr all in environments with network, power and environmental challenges. This diversity creates pain in managing, scaling, securing and orchestrating industrial data. Ultimately resulting in silos of data with competing context. FogLAMP is designed to eliminate those silos by providing a very flexible data collections and distribution mechanism all using the same APIs, features and functions.

Specialized Analytical Environments With the advent of cloud systems and sophisticated analytic tools it may no longer be possible to have a single system that is both your system of record and the place on which the analytics takes place. FogLAMP allows you to distribute your data to multiple systems, either in part or as a whole. This allows you to get just the data you need to the systems that need it without compromising your system of record.

Resilience FogLAMP provides mechanisms to store and forward your data. Data is no longer lost if a connection to some key system is unavailable.

Edge processing Using the FogLAMP intelligent data pipelines concept, FogLAMP allows for your data to be processed close to where it is gathered. This can save both network bandwidth and reduce costs when high bandwidth sensors such as vibration monitors or image capture is used. In addition it reduces the latency when timely action is required compared with shipping and processing data in the cloud or at some centralized IT location.

No code/Low code solutions FogLAMP provides tools that allow the OT engineer to create solutions by use of existing processing elements that can be combined and augmented with little or no coding required. This allows the OT organization to be able to quickly and independently obtain the data they need for their specific requirements.

Process Optimization & Operational Efficiency The FogLAMP intelligent pipelines, with their prebuilt processing elements and through use of machine learning techniques can be used to improve operational efficiency by giving operators immediate feedback on the state of the process of product being produced without remote analytics and the associated delays involved.

1.2 Architectural Overview

FogLAMP is implemented as a collection of microservices which include:

- Core services, including security, monitoring, and storage
- Data transformation and alerting services
- South services: Collect data from sensors and other FogLAMP systems
- North services: Transmit and integrate data to historians and other systems
- Edge data processing applications
- Event detection and notification
- Set point control

Services can easily be developed and incorporated into the FogLAMP framework. FogLAMP services may also be customized by creating new plugins, written in C/C++ or Python, for data collection, processing, export, rule evaluation and event notification. The describe how to do this.

More detail on the FogLAMP architecture can be found in the section .

1.3 No-code/Low-code Development

FogLAMP can be extended by writing code to add new plugins. Additionally, it is easily tailored by combining pre-written data processing filters applied in linear pipelines to data as it comes into or goes out of the FogLAMP system. A number of filters exist that can be customized with small snippets of code written in the Python scripting language. These snippets of code allow the end user to produce custom processing without the need to develop more complex plugins or other code. The environment also allows them to experiment with these code snippets to obtain the results desired.

Data may be processed on the way into FogLAMP or on the way out. Processing on the way in allows the data to be manipulated to the way the organization wants it. Processing on the way out allows the data to be manipulate to suit the up stream system that will use the data without impacting the data that might go to another up stream system.

See the section .

QUICK START GUIDE

2.1 Installing FogLAMP

FogLAMP is extremely lightweight and can run on inexpensive edge devices, sensors and actuator boards. For the purposes of this manual, we assume that all services are running on a Raspberry Pi running the Raspbian operating system. Be sure your system has plenty of storage available for data readings.

If your system does not have Raspbian pre-installed, you can find instructions on downloading and installing it at <https://www.raspberrypi.org/downloads/raspbian/>. After installing Raspbian, ensure you have the latest updates by executing the following commands on your FogLAMP server:

```
sudo apt-get update
sudo apt-get upgrade
sudo apt-get update
```

You can obtain FogLAMP in two ways:

- Dianomic Systems hosts a package repository that allows the FogLAMP packages to be loaded using the system package manager. This is the recommended method for long term use of FogLAMP as it gives access to all the FogLAMP plugins and provides a route for easy upgrade of the FogLAMP packages. This also has the advantages that once the repository is configured you are able to install new plugins directly from the FogLAMP user interface without the need to resort to the Linux command line.
- Dianomic Systems offers pre-built, certified binaries of FogLAMP for Debian using either Intel or ARM architectures. This is perhaps the simplest method for users not used to Linux. You can download the complete set of packages from <http://dianomic.com/download-foglamp/>.

In general, FogLAMP installation will require the following packages:

- FogLAMP core
- FogLAMP user interface
- One or more FogLAMP South services
- One or more FogLAMP North service (OSI PI and OCS north services are included in FogLAMP core)

2.1.1 Using the package repository to install FogLAMP

If you choose to use the Dianomic Systems package repository to install the packages you will need to follow the steps outlined below for the particular platform you are using.

Ubuntu or Debian

On a Ubuntu or Debian system, including the Raspberry Pi, the package manager that is supported is *apt*. You will need to add the Dianomic Systems archive server into the configuration of *apt* on your system. The first thing that must be done is to add the key that is used to verify the package repository. To do this run the command

```
wget -q -O - http://archives.dianomic.com/KEY.gpg | sudo apt-key add -
```

Once complete you can add the repository itself into the *apt* configuration file */etc/apt/sources.list*. The simplest way to do this is the use the *add-apt-repository* command. The exact command will vary between systems;

- Raspberry Pi does not have an *apt-add-repository* command, the user must edit the *apt* sources file manually

```
sudo vi /etc/apt/sources.list
```

and add the line

```
deb http://archives.dianomic.com/foglamp/latest/buster/armv7l/ /
```

to the end of the file.

Note: Replace *buster* with *stretch* or *bullseye* based on the OS image used.

- Users with an Intel or AMD system with Ubuntu 18.04 should run

```
sudo add-apt-repository "deb http://archives.dianomic.com/foglamp/latest/  
↳ubuntu1804/x86_64/ / "
```

- Users with an Intel or AMD system with Ubuntu 20.04 should run

```
sudo add-apt-repository "deb http://archives.dianomic.com/foglamp/latest/  
↳ubuntu2004/x86_64/ / "
```

Note: We do not support the *aarch64* architecture with Ubuntu 20.04 yet.

- Users with an Arm system with Ubuntu 18.04, such as the Odroid board, should run

```
sudo add-apt-repository "deb http://archives.dianomic.com/foglamp/latest/  
↳ubuntu1804/aarch64/ / "
```

- Users of the Mendel operating system on a Google Coral create the file */etc/apt/sources.list.d/foglamp.list* and insert the following content

```
deb http://archives.dianomic.com/foglamp/latest/mendel/aarch64/ /
```

Once the repository has been added you must inform the package manager to go and fetch a list of the packages it supports. To do this run the command

```
sudo apt -y update
```

You are now ready to install the FogLAMP packages. You do this by running the command

```
sudo apt -y install *package*
```

You may also install multiple packages in a single command. To install the base foglamp package, the foglamp user interface and the sinusoid south plugin run the command

```
sudo DEBIAN_FRONTEND=noninteractive apt -y install foglamp foglamp-gui foglamp-south-
↪sinusoid
```

2.1.2 Installing FogLAMP downloaded packages

Assuming you have downloaded the packages from the download link given above. Use SSH to login to the system that will host FogLAMP services. For each FogLAMP package that you choose to install, type the following command:

```
sudo apt -y install PackageName
```

The key packages to install are the FogLAMP core and the FogLAMP User Interface:

```
sudo DEBIAN_FRONTEND=noninteractive apt -y install ./foglamp-1.8.0-armv7l.deb
sudo apt -y install ./foglamp-gui-1.8.0.deb
```

You will need to install one of more South plugins to acquire data. You can either do this now or when you are adding the data source. For example, to install the plugin for the Sense HAT sensor board, type:

```
sudo apt -y install ./foglamp-south-sensehat-1.8.0-armv7l.deb
```

You may also need to install one or more North plugins to transmit data. Support for OSIssoft PI and OCS are included with the FogLAMP core package, so you don't need to install anything more if you are sending data to only these systems.

2.1.3 Checking package installation

To check what packages have been installed, ssh into your host system and use the dpkg command:

```
dpkg -l | grep 'foglamp'
```

2.1.4 Run with PostgreSQL

To start FogLAMP with PostgreSQL, first you need to install the PostgreSQL package explicitly. See the below links for setup

Also you need to change the value of Storage plugin. See or with below curl command

```
$ curl -sX PUT localhost:8081/foglamp/category/Storage/plugin -d '{"value": "postgres"
↪}'
{
  "description": "The main storage plugin to load",
```

(continues on next page)

(continued from previous page)

```
"type": "string",  
"order": "1",  
"displayName": "Storage Plugin",  
"default": "sqlite",  
"value": "postgres"  
}
```

Now, it's time to restart FogLAMP. Thereafter you will see FogLAMP is running with PostgreSQL.

2.2 Starting and stopping FogLAMP

FogLAMP administration is performed using the “foglamp” command line utility. You must first ssh into the host system. The FogLAMP utility is installed by default in /usr/local/foglamp/bin.

The following command options are available:

- **Start:** Start the FogLAMP system
- **Stop:** Stop the FogLAMP system
- **Status:** Lists currently running FogLAMP services and tasks
- **Reset:** Delete all data and configuration and return FogLAMP to factory settings
- **Kill:** Kill FogLAMP services that have not correctly responded to Stop
- **Help:** Describe FogLAMP options

For example, to start the FogLAMP system, open a session to the FogLAMP device and type:

```
/usr/local/foglamp/bin/foglamp start
```

2.3 Troubleshooting FogLAMP

FogLAMP logs status and error messages to syslog. To troubleshoot a FogLAMP installation using this information, open a session to the FogLAMP server and type:

```
grep -a 'foglamp' /var/log/syslog | tail -n 20
```


2.4 Running the FogLAMP GUI

FogLAMP offers an easy-to-use, browser-based GUI. To access the GUI, open your browser and enter the IP address of the FogLAMP server into the address bar. This will display the FogLAMP dashboard.

You can easily use the FogLAMP UI to monitor multiple FogLAMP servers. To view and manage a different server, click “Settings” in the left menu bar. In the “Connection Setup” pane, enter the IP address and port number for the new server you wish to manage. Click the “Set the URL & Restart” button to switch the UI to the new server.

If you are managing a very lightweight server or one that is connected via a slow network link, you may want to reduce the UI update frequency to minimize load on the server and network. You can adjust this rate in the “GUI Settings” pane of the Settings screen. While the graph rate and ping rate can be adjusted individually, in general you should set them to the same value.

2.4.1 FogLAMP Dashboard



This screen provides an overview of FogLAMP operations. You can customize the information and time frames displayed on this screen using the drop-down menus in the upper right corner. The information you select will be displayed in a series of graphs.

You can choose to view a graph of any of the sensor reading being collected by the FogLAMP system. In addition, you can view graphs of the following system-wide information:

- **Readings:** The total number of data readings collected by FogLAMP since system boot
- **Buffered:** The number of data readings currently stored by the system
- **Discarded:** Number of data readings discarded before being buffered (due to data errors, for example)
- **Unsent:** Number of data readings that were not sent successfully
- **Purged:** The total number of data readings that have been purged from the system
- **Unpurged:** The number of data readings that were purged without being sent to a North service.

2.5 Managing Data Sources

The screenshot shows the FogLAMP web interface. The top navigation bar includes the FogLAMP logo, the device name 'raspberrypi/FogLAMP', and system statistics: 'Received: 9,480', 'Sent: 0', and 'Uptime: 00:39:49'. On the left, a sidebar menu lists various functions, with 'South' currently selected. The main content area is titled 'South Services' and features an 'Add +' button. Below this is a table with the following data:

Name	Status	Assets	Readings
<u>SenseHat sensor</u>	enabled	sensehat/pressure	1,580
		sensehat/temperature	1,580
		sensehat/humidity	1,580
		sensehat/magnetometer	1,580
		sensehat/gyroscope	1,580
		sensehat/accelerometer	1,580

At the bottom of the screen, a copyright notice reads: '© 2018 DIANOMIC SYSTEMS. All Rights Reserved.'

Data sources are managed from the South Services screen. To access this screen, click on “South” from the menu bar on the left side of any screen.

The South Services screen displays the status of all data sources in the FogLAMP system. Each data source will display its status, the data assets it is providing, and the number of readings that have been collected.

2.5.1 Adding Data Sources

To add a data source, you will first need to install the plugin for that sensor type. If you have not already done this, open a terminal session to your FogLAMP server. Download the package for the plugin and enter:

```
sudo apt -y install PackageName
```

Once the plugin is installed return to the FogLAMP GUI and click on “Add+” in the upper right of the South Services screen. FogLAMP will display a series of 3 screens to add the data source:

1. The first screen will ask you to select the plugin for the data source from the list of installed plugins. If you do not see the plugin you need, refer to the Installing FogLAMP section of this manual. In addition, this screen allows you to specify a display name for the data source.
2. The second screen allows you to configure the plugin and the data assets it will provide.

Note: Every data asset in FogLAMP must have a unique name. If you have multiple sensors using the same plugin, modify the asset names on this screen so they are unique.

Some plugins allow you to specify an asset name prefix that will apply to all the asset names for that sensor. Refer to the individual plugin documentation for descriptions of the fields on this screen.

3. If you modify any of the configuration fields, click on the “save” button to save them.
4. The final screen allows you to specify whether the service will be enabled immediately for data collection or await enabling in the future.

To modify the configuration of a data source, click on its name in the South Services screen. This will display a list of all parameters available for that data source. If you make any changes, click on the “save” button in the top panel to save the new configuration. Click on the “x” button in the upper right corner to return to the South Services screen.

2.5.3 Enabling and Disabling Data Sources

To enable or disable a data source, click on its name in the South Services screen. Under the list of data source parameters, there is a check box to enable or disable the service. If you make any changes, click on the “save” button in the bottom panel near the check box to save the new configuration.

2.6 Viewing Data

The screenshot shows the FogLAMP web interface. The top header displays system status: Received: 24,510, Sent: 0, Uptime: 01:23:54. The left sidebar contains a navigation menu with options: Dashboard, Assets & Readings (selected), South, North, Configuration, Schedules, Certificate Store, Backup & Restore, Logs, Audit, System, Tasks, Support, and Settings. The main content area shows a table titled 'Assets & Readings' with two columns: 'Asset' and 'Readings'. The table lists six assets, all with a reading of 4,080. Each row has a graph icon and a download icon. The footer shows copyright information: © 2018 DIANOMIC SYSTEMS. All Rights Reserved.

Asset	Readings
sensehat/accelerometer	4,080
sensehat/gyroscope	4,080
sensehat/humidity	4,080
sensehat/magnetometer	4,080
sensehat/pressure	4,080
sensehat/temperature	4,080


You can inspect all the data buffered by the FogLAMP system on the Assets page. To access this page, click on “Assets & Readings” from the left-side menu bar.

This screen will display a list of every data asset in the system. Alongside each asset are two icons; one to display a graph of the asset and another to download the data stored for that asset as a CSV file.

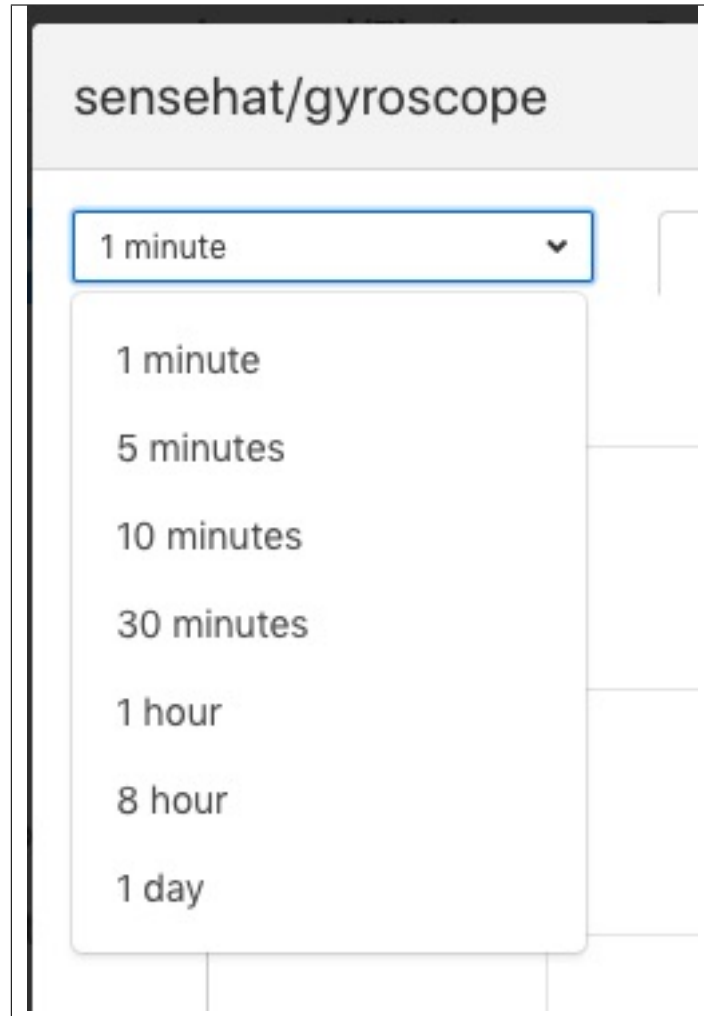
2.6.1 Display Graph




By clicking on the graph button next to each asset name, you can view a graph of individual data readings. A graph will be displayed with a plot for each data point within the asset.



It is possible to change the time period to which the graph refers by use of the plugin list in the top left of the graph.



Where an asset contains multiple data points each of these is displayed in a different colour. Graphs for particular data points can be toggled on and off by clicking on the key at the top of the graph. Those data points not should will be indicated by striking through the name of the data point.



A summary tab is also available, this will show the minimum, maximum and average values for each of the data points. Click on *Summary* to show the summary tab.



2.6.2 Download Data



By clicking on the download icon adjacent to each asset you can download the stored data for the asset. The format of the file is download is a CSV file that is designed to be loaded into a spreadsheet such as Excel, Numbers or OpenOffice Calc.

The file contains a header row with the names of the data points within the asset, the first column is always the timestamp when the reading was taken, the header for this being *timestamp*. The data is sorted in chronological order with the newest data first.

sensehat_gyroscope-readings			
timestamp	z	x	y
2020-05-04 14:30:49.145006	0.000792725	0.0010765493	0.0022465843
2020-05-04 14:30:48.145022	0.0010982286	-0.0004502609	0.000719551
2020-05-04 14:30:47.145006	0.0007928684	0.0032151192	-0.0011130939
2020-05-04 14:30:46.145008	-0.0013448559	0.0047423765	0.0001088944
2020-05-04 14:30:45.145000	-0.0004286431	0.0007723272	-0.0020291833
2020-05-04 14:30:44.144999	-0.0001233947	0.0013834909	0.0007194807
2020-05-04 14:30:43.145001	-0.000734292	-0.0001437888	0.0004143068

2.7 Sending Data to Other Systems

The screenshot displays the FogLAMP web interface. The top header shows the FogLAMP logo, the device name 'raspberrypi/FogLAMP', and system statistics: 'Received: 10,974', 'Sent: 0', and 'Uptime: 00:44:14'. The left sidebar menu has 'North' highlighted. The main panel, titled 'North Plugins', features a 'Create North Instance +' button and a table listing data sending processes. The table has three columns: 'Process', 'Status', and 'Sent'. One process, 'Plant Librarian', is listed with a status of 'enabled' and 0 readings sent. The footer of the main panel contains the copyright notice: '© 2018 DIANOMIC SYSTEMS. All Rights Reserved.'

Data destinations are managed from the North Services screen. To access this screen, click on “North” from the menu bar on the left side of any screen.

The North Services screen displays the status of all data sending processes in the FogLAMP system. Each data destination will display its status and the number of readings that have been collected.

2.7.1 Adding Data Destinations

To add a data destination, click on “Create North Instance+” in the upper right of the North Services screen. FogLAMP will display a series of 3 screens to add the data destination:

1. The first screen will ask you to select the plugin for the data destination from the list of installed plugins. If you do not see the plugin you need, refer to the Installing FogLAMP section of this manual. In addition, this screen allows you to specify a display name for the data destination. In addition, you can specify how frequently data will be forwarded to the destination in days, hours, minutes and seconds. Enter the number of days in the interval in the left box and the number of hours, minutes and seconds in format HH:MM:SS in the right box.
2. The second screen allows you to configure the plugin and the data assets it will send. See the section below for specifics of configuring a PI, EDS or OCS destination.
3. The final screen loads the plugin. You can specify whether it will be enabled immediately for data sending or to await enabling in the future.

2.7.2 Configuring Data Destinations

To modify the configuration of a data destination, click on its name in the North Services screen. This will display a list of all parameters available for that data source. If you make any changes, click on the “save” button in the top panel to save the new configuration. Click on the “x” button in the upper right corner to return to the North Services screen.

2.7.3 Enabling and Disabling Data Destinations

To enable or disable a data source, click on its name in the North Services screen. Under the list of data source parameters, there is a check box to enable or disable the service. If you make any changes, click on the “save” button in the bottom panel near the check box to save the new configuration.

2.7.4 Using the OMF plugin

OSISoft data historians are one of the most common destinations for FogLAMP data. FogLAMP supports the full range of OSISoft historians; the PI System, Edge Data Store (EDS) and OSISoft Cloud Services (OCS). To send data to a PI server you may use either the older PI Connector Relay or the newer PI Web API OMF endpoint. It is recommended that new users use the PI Web API OMF endpoint rather than the Connector Relay which is no longer supported by OSISoft.

2.8 PI Web API OMF Endpoint

To use the PI Web API OMF endpoint first ensure the OMF option was included in your PI Server when it was installed.

Now go to the FogLAMP user interface, create a new North instance and select the “OMF” plugin on the first screen. The second screen will request the following information:

?

Endpoint

PI Web API

Send full structure

☒

Naming Scheme

Concise

Server hostname

localhost

Server port, 0=use the default

0

Producer Token

omf_north_0001

Data Source

readings

Static Data

Location: Palo Alto, Company: Dianomic

Sleep Time Retry

1

Maximum Retry

3

HTTP Timeout

10

Integer Format

int64

Number Format

float64

Compression

☒

Default Asset Framework Location

/fledge/data_piwebapi/default

Asset Framework hierarchy rules

1

{}

PI Web API Authentication Method

anonymous

PI Web API User Id

user_id

PI Web API Password

.....

PI Web API Kerberos keytab file

piwebapi_kerberos_https.keytab

OCS Namespace

name_space

OCS Tenant ID

ocs_tenant_id

OCS Client ID

ocs_client_id

OCS Client Secret

.....

Select PI Web API from the Endpoint options.

- **Basic Information**

- **Endpoint:** This is the type of OMF endpoint. In this case, choose PI Web API.
- **Send full structure:** Used to control if Asset Framework structure messages are sent to the PI Server. If this is turned off then the data will not be placed in the Asset Framework.
- **Naming scheme:** Defines the naming scheme to be used when creating the PI points in the PI Data Archive. See [Naming Scheme](#).
- **Server hostname:** The hostname or address of the PI Web API server. This is normally the same address as the PI Server.
- **Server port:** The port the PI Web API OMF endpoint is listening on. Leave as 0 if you are using the default port.
- **Data Source:** Defines which data is sent to the PI Server. Choices are: readings or statistics (that is, FogLAMP's internal statistics).
- **Static Data:** Data to include in every reading sent to PI. For example, you can use this to specify the location of the devices being monitored by the FogLAMP server.

- **Asset Framework**

- **Default Asset Framework Location:** The location in the Asset Framework hierarchy into which the data will be inserted. All data will be inserted at this point in the Asset Framework hierarchy unless a later rule overrides this. Note this field does not include the name of the target Asset Framework Database; the target database is defined on the PI Web API server by the PI Web API Admin Utility.
- **Asset Framework Hierarchies Rules:** A set of rules that allow specific readings to be placed elsewhere in the Asset Framework. These rules can be based on the name of the asset itself or some metadata associated with the asset. See [Asset Framework Hierarchy Rules](#).

- **PI Web API authentication**

- **PI Web API Authentication Method:** The authentication method to be used: anonymous, basic or kerberos. Anonymous equates to no authentication, basic authentication requires a user name and password, and Kerberos allows integration with your single signon environment.
- **PI Web API User Id:** For Basic authentication, the user name to authenticate with the PI Web API.
- **PI Web API Password:** For Basic authentication, the password of the user we are using to authenticate.
- **PI Web API Kerberos keytab file:** The Kerberos keytab file used to authenticate.

- **Connection management (These should only be changed with guidance from support)**

- **Sleep Time Retry:** Number of seconds to wait before retrying the HTTP connection (FogLAMP doubles this time after each failed attempt).
- **Maximum Retry:** Maximum number of times to retry connecting to the PI Server.
- **HTTP Timeout:** Number of seconds to wait before FogLAMP will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match FogLAMP data types to the data type configured in PI. This defaults to int64 but may be set to any OMF data type compatible with integer data, e.g. int32.
- **Number Format:** Used to match FogLAMP data types to the data type configured in PI. The default is float64 but may be set to any OMF datatype that supports floating point values.

- **Compression:** Compress the readings data before sending them to the PI Web API OMF endpoint. This setting is not related to data compression in the PI Data Archive.

2.9 Edge Data Store OMF Endpoint

To use the OSIsoft Edge Data Store first install Edge Data Store on the same machine as your FogLAMP instance. It is a limitation of Edge Data Store that it must reside on the same host as any system that connects to it with OMF.

Now go to the FogLAMP user interface, create a new North instance and select the “OMF” plugin on the first screen. The second screen will request the following information:

?

Endpoint

Edge Data Store

Send full structure

☒

Naming Scheme

Concise

Server hostname

localhost

Server port, 0=use the default

0

Producer Token

omf_north_0001

Data Source

readings

Static Data

Location: Palo Alto, Company: Dianomic

Sleep Time Retry

1

Maximum Retry

3

HTTP Timeout

10

Integer Format

int64

Number Format

float64

Compression

☒

Default Asset Framework Location

/fledge/data_piwebapi/default

Asset Framework hierarchy rules

1

{}

PI Web API Authentication Method

anonymous

PI Web API User Id

user_id

PI Web API Password

.....

PI Web API Kerberos keytab file

piwebapi_kerberos_https.keytab

OCS Namespace

name_space

OCS Tenant ID

ocs_tenant_id

OCS Client ID

ocs_client_id

OCS Client Secret

.....

Select Edge Data Store from the Endpoint options.

- **Basic Information**

- **Endpoint:** This is the type of OMF endpoint. In this case, choose Edge Data Store.
- **Naming scheme:** Defines the naming scheme to be used when creating the PI points within the PI Server. See [Naming Scheme](#).
- **Server hostname:** Normally the hostname or address of the OMF endpoint. For Edge Data Store, this must be *localhost*.
- **Server port:** The port the Edge Data Store is listening on. Leave as 0 if you are using the default port.
- **Data Source:** Defines which data is sent to the Edge Data Store. Choices are: readings or statistics (that is, FogLAMP's internal statistics).
- **Static Data:** Data to include in every reading sent to PI. For example, you can use this to specify the location of the devices being monitored by the FogLAMP server.

- **Connection management (These should only be changed with guidance from support)**

- **Sleep Time Retry:** Number of seconds to wait before retrying the HTTP connection (FogLAMP doubles this time after each failed attempt).
- **Maximum Retry:** Maximum number of times to retry connecting to the PI server.
- **HTTP Timeout:** Number of seconds to wait before FogLAMP will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match FogLAMP data types to the data type configured in PI. This defaults to int64 but may be set to any OMF data type compatible with integer data, e.g. int32.
- **Number Format:** Used to match FogLAMP data types to the data type configured in PI. The default is float64 but may be set to any OMF datatype that supports floating point values.
- **Compression:** Compress the readings data before sending them to the Edge Data Store.

2.10 OSIssoft Cloud Services OMF Endpoint

Go to the FogLAMP user interface, create a new North instance and select the “OMF” plugin on the first screen. The second screen will request the following information:

Endpoint

OSIsoft Cloud Services

Send full structure

☒

Naming Scheme

Concise

Server hostname

localhost

Server port, 0=use the default

0

Producer Token

omf_north_0001

Data Source

readings

Static Data

Location: Palo Alto, Company: Dianomic

Sleep Time Retry

1

Maximum Retry

3

HTTP Timeout

10

Integer Format

int64

Number Format

float64

Compression

☒

Default Asset Framework Location

/fledge/data_piwebapi/default

Asset Framework hierarchy rules

1

{}

PI Web API Authentication Method

anonymous

PI Web API User Id

user_id

PI Web API Password

.....

PI Web API Kerberos keytab file

piwebapi_kerberos_https.keytab

OCS Namespace

name_space

OCS Tenant ID

ocs_tenant_id

OCS Client ID

ocs_client_id

OCS Client Secret

.....

Select OSIsoft Cloud Services from the Endpoint options.

- **Basic Information**

- **Endpoint:** This is the type of OMF endpoint. In this case, choose OSIsoft Cloud Services.
- **Naming scheme:** Defines the naming scheme to be used when creating the PI points within the PI Server. See [Naming Scheme](#).
- **Data Source:** Defines which data is sent to OSIsoft Cloud Services. Choices are: readings or statistics (that is, FogLAMP's internal statistics).
- **Static Data:** Data to include in every reading sent to OSIsoft Cloud Services. For example, you can use this to specify the location of the devices being monitored by the FogLAMP server.

- **Authentication**

- **OCS Namespace:** Your namespace within OSIsoft Cloud Services.
- **OCS Tenant ID:** Your OSIsoft Cloud Services Tenant ID for your account.
- **OCS Client ID:** Your OSIsoft Cloud Services Client ID for your account.
- **OCS Client Secret:** Your OSIsoft Cloud Services Client Secret.

- **Connection management (These should only be changed with guidance from support)**

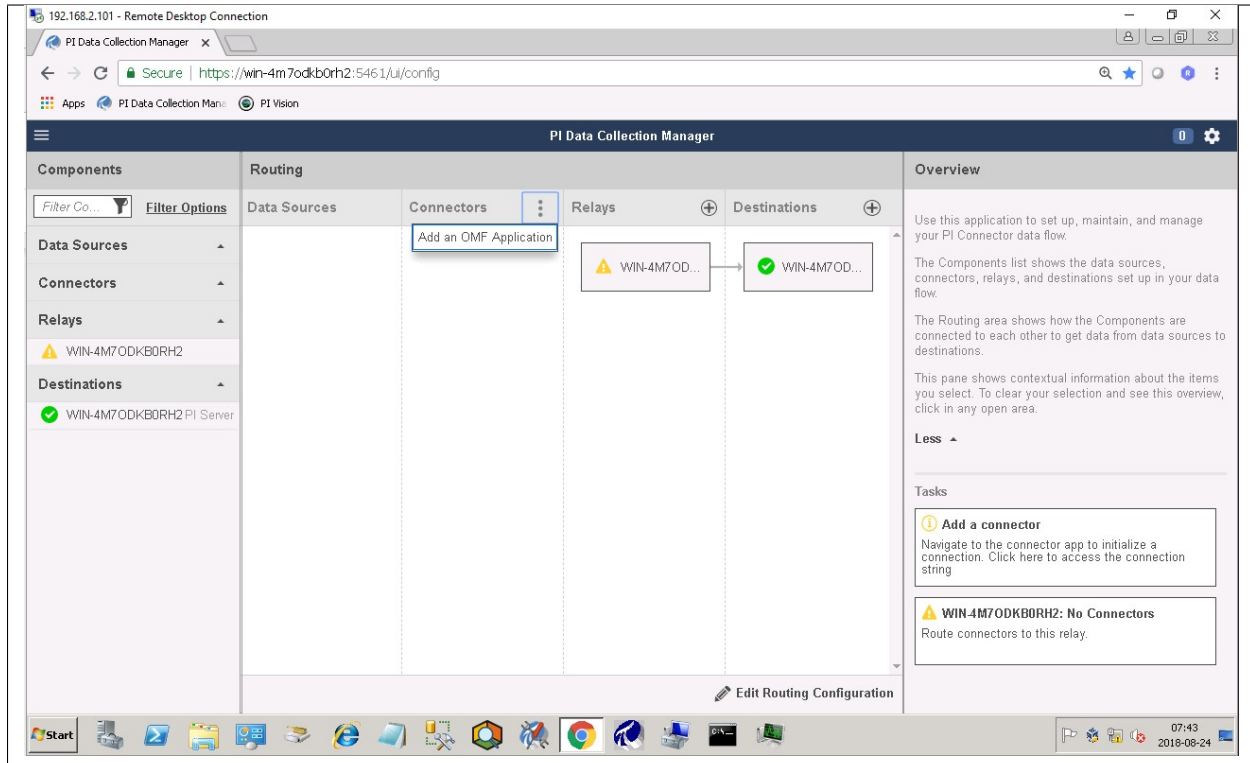
- **Sleep Time Retry:** Number of seconds to wait before retrying the HTTP connection (FogLAMP doubles this time after each failed attempt).
- **Maximum Retry:** Maximum number of times to retry connecting to the PI server.
- **HTTP Timeout:** Number of seconds to wait before FogLAMP will time out an HTTP connection attempt.

- **Other (Rarely changed)**


- **Integer Format:** Used to match FogLAMP data types to the data type configured in PI. This defaults to int64 but may be set to any OMF data type compatible with integer data, e.g. int32.
- **Number Format:** Used to match FogLAMP data types to the data type configured in PI. The default is float64 but may be set to any OMF datatype that supports floating point values.
- **Compression:** Compress the readings data before sending them to OSIsoft Cloud Services.

2.11 PI Connector Relay


The PI Connector Relay has been discontinued by OSIsoft. All new deployments should use the PI Web API endpoint. Existing installations will still be supported. The PI Connector Relay was the original mechanism by which OMF data could be ingesting into a PI Server. To use the PI Connector Relay, open and sign into the PI Relay Data Connection Manager.



To add a new connector for the FogLAMP system, click on the drop down menu to the right of “Connectors” and select “Add an OMF application”. Add and save the requested configuration information.



Connect the new application to the PI Connector Relay by selecting the new FogLAMP application, clicking the check box for the PI Connector Relay and then clicking “Save Configuration”.



Finally, select the new FogLAMP application. Click “More” at the bottom of the Configuration panel. Make note of the Producer Token and Relay Ingress URL.

Now go to the FogLAMP user interface, create a new North instance and select the “OMF” plugin on the first screen. The second screen will request the following information:

?

Endpoint

Connector Relay

Send full structure

☒

Naming Scheme

Concise

Server hostname

localhost

Server port, 0=use the default

0

Producer Token

omf_north_0001

Data Source

readings

Static Data

Location: Palo Alto, Company: Dianomic

Sleep Time Retry

1

Maximum Retry

3

HTTP Timeout

10

Integer Format

int64

Number Format

float64

Compression

☒

Default Asset Framework Location

/fledge/data_piwebapi/default

Asset Framework hierarchy rules

1

{}

PI Web API Authentication Method

anonymous

PI Web API User Id

user_id

PI Web API Password

.....

PI Web API Kerberos keytab file

piwebapi_kerberos_https.keytab

OCS Namespace

name_space

OCS Tenant ID

ocs_tenant_id

OCS Client ID

ocs_client_id

OCS Client Secret

.....

- **Basic Information**

- **Endpoint:** This is the type of OMF endpoint. In this case, choose Connector Relay.
- **Server hostname:** The hostname or address of the PI Connector Relay.
- **Server port:** The port the PI Connector Relay is listening on. Leave as 0 if you are using the default port.
- **Producer Token:** The Producer Token provided by the PI Relay Data Connection Manager.
- **Data Source:** Defines which data is sent to the PI Connector Relay. Choices are: readings or statistics (that is, FogLAMP's internal statistics).
- **Static Data:** Data to include in every reading sent to PI. For example, you can use this to specify the location of the devices being monitored by the FogLAMP server.

- **Connection management (These should only be changed with guidance from support)**

- **Sleep Time Retry:** Number of seconds to wait before retrying the HTTP connection (FogLAMP doubles this time after each failed attempt).
- **Maximum Retry:** Maximum number of times to retry connecting to the PI server.
- **HTTP Timeout:** Number of seconds to wait before FogLAMP will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match FogLAMP data types to the data type configured in PI. This defaults to int64 but may be set to any OMF data type compatible with integer data, e.g. int32.
- **Number Format:** Used to match FogLAMP data types to the data type configured in PI. The default is float64 but may be set to any OMF datatype that supports floating point values.
- **Compression:** Compress the readings data before sending it to the PI System.

2.11.1 Naming Scheme

The naming of objects in the Asset Framework and of the attributes of those objects has a number of constraints that need to be understood when storing data into a PI Server using OMF. An important factor in this is the stability of your data structures. If you have objects in your environment that are likely to change, you may wish to take a different naming approach. Examples of changes are a difference in the number of attributes between readings, and a change in the data types of attributes.

This occurs because of a limitation of the OMF interface to the PI Server. Data is sent to OMF in a number of stages. One of these is the definition of the Types used to create AF Element Templates. OMF uses a Type to define an AF Element Template but once defined it cannot be changed. If an updated Type definition is sent to OMF, it will be used to create a new AF Element Template rather than changing the existing one. This means a new AF Element Template is created each time a Type changes.

The OMF plugin names objects in the Asset Framework based upon the asset name in the reading within FogLAMP. Asset names are typically added to the readings in the south plugins, however they may be altered by filters between the south ingest and the north egress points in the data pipeline. Asset names can be overridden using the *OMF Hints* mechanism described below.

The attribute names used within the objects in the PI System are based on the names of the datapoints within each Reading within FogLAMP. Again *OMF Hints* can be used to override this mechanism.

The naming used within the objects in the Asset Framework is controlled by the *Naming Scheme* option:

Concise No suffix or prefix is added to the asset name and property name when creating objects in the Asset Framework and PI Points in the PI Data Archive. However, if the structure of an asset changes a new AF Element Template will be created which will have the suffix `-type*x*` appended to it.

Use Type Suffix The AF Element names will be created from the asset names by appending the suffix `-type*x*` to the asset name. If the structure of an asset changes a new AF Element name will be created with an updated suffix.

Use Attribute Hash AF Attribute names will be created using a numerical hash as a prefix.

Backward Compatibility The naming reverts to the rules that were used by version 1.9.1 and earlier of FogLAMP: both type suffixes and attribute hashes will be applied to the name.

2.11.2 Asset Framework Hierarchy Rules

The Asset Framework rules allow the location of specific assets within the Asset Framework to be controlled. There are two basic types of hint:

- Asset name placement: the name of the asset determines where in the Asset Framework the asset is placed,
- Meta data placement: metadata within the reading determines where the asset is placed in the Asset Framework.

The rules are encoded within a JSON document. This document contains two properties in the root of the document: one for name-based rules and the other for metadata based rules.

```
{
  "names" :
  {
    "asset1" : "/Building1/EastWing/GroundFloor/Room4",
    "asset2" : "Room14"
  },
  "metadata" :
  {
    "exist" :
    {
      "temperature" : "temperatures",
      "power" : "/Electrical/Power"
    },
    "nonexist" :
    {
      "unit" : "Uncalibrated"
    },
    "equal" :
    {
      "room" :
      {
        "4" : "ElecticalLab",
        "6" : "FluidLab"
      }
    },
    "notequal" :
    {
      "building" :
      {
        "plant" : "/Office/Environment"
      }
    }
  }
}
```

(continues on next page)

(continued from previous page)

```

    }
}

```

The name type rules are simply a set of asset name and Asset Framework location pairs. The asset names must be complete names; there is no pattern matching within the names.

The metadata rules are more complex. Four different tests can be applied:

- **exists:** This test looks for the existence of the named datapoint within the asset.
- **nonexist:** This test looks for the lack of a named datapoint within the asset.
- **equal:** This test looks for a named datapoint having a given value.
- **notequal:** This test looks for a name datapoint having a value different from that specified.

The *exist* and *nonexist* tests take a set of name/value pairs that are tested. The name is the datapoint name to examine and the value is the Asset Framework location to use. For example

```

"exist" :
{
    "temperature" : "temperatures",
    "power"       : "/Electrical/Power"
}

```

If an asset has a datapoint called *temperature* it will be stored in the AF hierarchy *temperatures*, if the asset had a datapoint called *power* the asset will be placed in the AF hierarchy */Electrical/Power*.

The *equal* and *notequal* tests take an object as a child, the name of the object is datapoint to examine, the child nodes are sets of values and locations. For example

```

"equal" :
{
    "room" :
    {
        "4" : "ElectricalLab",
        "6" : "FluidLab"
    }
}

```

In this case if the asset has a datapoint called *room* with a value of *4* then the asset will be placed in the AF location *ElectricalLab*, if it has a value of *6* then it is placed in the AF location *FluidLab*.

If an asset matches multiple rules in the ruleset it will appear in multiple locations in the hierarchy, the data is shared between each of the locations.

If an OMF Hint exists within a particular reading this will take precedence over generic rules.

The AF location may be a simple string or it may also include substitutions from other datapoints within the reading. For example if the reading has a datapoint called *room* that contains the room in which the readings were taken, an AF location of */BuildingA/\${room}* would put the reading in the Asset Framework using the value of the room datapoint. The reading

```

"reading" : {
    "temperature" : 23.4,
    "room"        : "B114"
}

```

would be put in the AF at */BuildingA/B114* whereas a reading of the form


```
"reading" : {
  "temperature" : 24.6,
  "room"       : "2016"
}
```

would be put at the location */BuildingA/2016*.

It is also possible to define defaults if the referenced datapoint is missing. In our example above if we used the location */BuildingA/\${room:unknown}* a reading without a *room* datapoint would be placed in */BuildingA/unknown*. If no default is given and the data point is missing then the level in the hierarchy is ignored. E.g. if we use our original location */BuildingA/\${room}* and we have the reading

```
"reading" : {
  "temperature" : 22.8,
}
```

this reading would be stored in */BuildingA*.

2.11.3 OMF Hints

The OMF plugin also supports the concept of hints in the actual data that determine how the data should be treated by the plugin. Hints are encoded in a specially named datapoint within the asset, *OMFHint*. The hints themselves are encoded as JSON within a string.

2.12 Number Format Hints

A number format hint tells the plugin what number format to use when inserting data into the PI Server. The following will cause all numeric data within the asset to be written using the format *float32*.

```
"OMFHint" : { "number" : "float32" }
```

The value of the *number* hint may be any numeric format that is supported by the PI Server.

2.13 Integer Format Hints

An integer format hint tells the plugin what integer format to use when inserting data into the PI Server. The following will cause all integer data within the asset to be written using the format *integer32*.

```
"OMFHint" : { "number" : "integer32" }
```

The value of the *number* hint may be any numeric format that is supported by the PI Server.

2.14 Type Name Hints

A type name hint specifies that a particular name should be used when defining the name of the type that will be created to store the object in the Asset Framework. This will override the *Naming Scheme* currently configured.

```
"OMFHint" : { "typeName" : "substation" }
```

2.15 Type Hint

A type hint is similar to a type name hint, but instead of defining the name of a type to create it defines the name of an existing type to use. The structure of the asset *must* match the structure of the existing type with the PI Server, it is the responsibility of the person that adds this hint to ensure this is the case.

```
"OMFHint" : { "type" : "pump" }
```

2.16 Tag Name Hint

Specifies that a specific tag name should be used when storing data in the PI Server.

```
"OMFHint" : { "tagName" : "AC1246" }
```

2.17 Datapoint Specific Hint

Hints may also be targeted to specific data points within an asset by using the datapoint hint. A *datapoint* hint takes a JSON object as its value; the object defines the name of the datapoint and the hint to apply.

```
"OMFHint" : { "datapoint" : { "name" : "voltage:", "number" : "float32" } }
```

The above hint applies to the datapoint *voltage* in the asset and applies a *number format* hint to that datapoint.

2.18 Asset Framework Location Hint

An Asset Framework location hint can be added to a reading to control the placement of the asset within the Asset Framework. An Asset Framework hint would be as follows:

```
"OMFHint" : { "AFLocation" : "/UK/London/TowerHill/Floor4" }
```

Note the following when defining an *AFLocation* hint:


- An asset in a FogLAMP Reading is used to create a [Container in the OSIsoft Asset Framework](#). A *Container* is an AF Element with one or more AF Attributes that are mapped to PI Points using the OSIsoft PI Point Data Reference. The name of the AF Element comes from the FogLAMP Reading asset name. The names of the AF Attributes come from the FogLAMP Reading datapoint names.
- If you edit the AF Location hint, the Container will be moved to the new location in the AF hierarchy.
- If you disable the OMF Hint filter, the Container will not move.

- If you wish to move a Container, you can do this with the PI System Explorer. Right-click on the AF Element that represents the Container. Choose Copy. Select the AF Element that will serve as the new parent of the Container. Right-click and choose *Paste*. You can then return to the original Container and delete it. *Note that PI System Explorer does not have the traditional Cut function for AF Elements.*
- If you move a Container, OMF North will not recreate it. If you then edit the AF Location hint, the Container will appear in the new location.

2.19 Adding OMF Hints


An OMF Hint is implemented as a string data point on a reading with the data point name of *OMFHint*. It can be added at any point in the processing of the data, however a specific plugin is available for adding the hints, the .

2.20 Backing up and Restoring FogLAMP



You can make a complete backup of all FogLAMP data and configuration. To do this, click on “Backup & Restore” in the left menu bar. This screen will show a list of all backups on the system and the time they were created. To make a new backup, click the “Backup” button in the upper right corner of the screen. You will briefly see a “Running” indicator in the lower left of the screen. After a period of time, the new backup will appear in the list. You may need to click the refresh button in the upper left of the screen to refresh the list. You can restore, delete or download any backup simply by clicking the appropriate button next to the backup in the list.

2.21 Troubleshooting and Support Information



FogLAMP keep detailed logs of system events for both auditing and troubleshooting use. To access them, click “Logs” in the left menu bar. There are five logs in the system:

- **Audit:** Tracks all configuration changes and data uploads performed on the FogLAMP system.
- **Notifications:** If you are using the FogLAMP notification service this log will give details of notifications that have been triggered
- **Packages:** This log will give you information about the installation and upgrade of FogLAMP packages for services and plugins.
- **System:** All events and scheduled tasks and their status.
- **Tasks:** The most recent scheduled tasks that have run and their status

If you have a service contract for your FogLAMP system, your support technician may ask you to send system data to facilitate troubleshooting an issue. To do this, click on “Support” in the left menu and then “Request New” in the upper right of the screen. This will create an archive of information. Click download to retrieve this archive to your system so you can email it to the technician.

2.22 Package Uninstallation

2.22.1 Debian Platform

Use the `apt` or the `apt-get` command to uninstall FogLAMP:

```
sudo apt -y purge foglamp
```

Note: You may notice the warning in the last row of the package removal output:

`dpkg: warning: while removing foglamp, directory '/usr/local/foglamp' not empty so not removed`

This is due to the fact that the data directory (`/usr/local/foglamp/data` by default) has not been removed, in case we might want to analyze or reuse the data further. So, if you want to remove foglamp completely from your system, then do `rm -rf /usr/local/foglamp` directory.

PROCESSING DATA

We have already seen that FogLAMP can collect data from a variety of sources, buffer it locally and send it on to one or more destination systems. It is also possible to process the data within FogLAMP to edit, augment or remove data as it traverses the FogLAMP system. In the same way FogLAMP makes extensive use of plugin components to add new sources of data and new destinations for that data, FogLAMP also uses plugins to add processing filters to the FogLAMP system.

3.1 Why Use Filters?

The concept behind filters is to create a set of small, useful pieces of functionality that can be inserted into the data flow from the south data ingress side to the north data egress side. By making these elements small and dedicated to a single task it increases the re-usability of the filters and greatly improves the chances when a new requirement is encountered that it can be satisfied by creating a filter pipeline from existing components or by augmenting existing components with the addition of any incremental processing required. The ultimate aim being to be able to create new applications within FogLAMP by merely configuring filters from the existing pool of available filters into a suitable pipeline without the need to write any new code.

3.2 What Can Be Done?

Data processing is done via plugins that are known as *filters* in FogLAMP, therefore it is not possible to give a definitive list of all the different processing that can occur, the design intent is that it is expandable by the user. The general types of things that can be done are;

- **Modify a value in a reading.** This could be as simple as applying a scale factor to convert from one measurement scale to another or more complex mathematical operation.
- **Modify asset or datapoint names.** Perform a simple textual substitution in order to change the name of an asset or a data point within that asset.
- **Add a new calculated value.** A new value can be calculated from a set of values, either based over a time period or based on a combination of different values, e.g. calculate power from voltage and current.
- **Add metadata to an asset.** This allows data such as units of measurement or information about the data source to be added to the data.
- **Compress data.** Only send data forward when the data itself shows significant change from previous values. This can be a useful technique to save bandwidth in low bandwidth or high cost network connections.
- **Conditionally forward data.** Only send data when a condition is satisfied or send low rate data unless some *interesting* condition is met.

- **Data conditioning.** Remove data from the data stream if the values are suspect or outside of reasonable conditions.

3.3 Where Can it Be Done?

Filters can be applied in two locations in the FogLAMP system;

- In the south service as data arrives in FogLAMP and before it is added to the storage subsystem for buffering.
- In the north tasks as the data is sent out to the upstream systems that receive data from the FogLAMP system.

More than one filter can be added to a single south or north within a FogLAMP instance. Filters are placed in an ordered pipeline of filters that are applied to the data in the order of the pipeline. The output of the first filter becomes the input to the second. Filters can thus be combined to perform complex sets of operations on a particular data stream into FogLAMP or out of FogLAMP.

The same filter plugin can appear in multiple places within a filter pipeline, a different instance is created for each and each one has its own configuration.

3.3.1 Adding a South Filter

In the following example we will add a filter to a south service. The filter we will use is the *expression* filter and we will convert the incoming value to a logarithmic scale. The south plugin used in this simple example is the *sinusoid* plugin that creates a simulated sine wave.

The process starts by selecting the *South* services in the FogLAMP GUI from the left-hand menu bar. Then click on the south service of interest. This will display a dialog that allows the south service to be edited.

The screenshot shows a configuration window titled "Sine South Service". It contains the following elements:

- Asset name:** A text input field containing "sinusoid".
- Enabled:** A checkbox that is checked.
- Show Advanced Config:** A link to expand the configuration options.
- Applications:** A section header with a plus icon, indicating where to add or manage applications for this service.
- Buttons:** "Cancel" and "Save" buttons are located at the bottom right.
- Service Info:** A section at the bottom with "Export Readings" and "Delete Service" buttons.

Towards the bottom of this dialog is a section labeled *Applications* with a + icon to the right, select the + icon to add a filter to the south service. A filter wizard is now shown that allows you to select the filter you wish to add and give that filter a name.

Sine South Service

1 Plugin Name 2 Review Configuration

Plugin

- asset
- change
- delta
- expression

Apply an expression to the data stream

[Install from available plugins](#)

Name name

Back Next

Select the *expression* filter and enter a name in the dialog. Now click on the *Next* button. A new page in the wizard appears that allows the configuration of the filter.

Sine South Service

1 Plugin Name 2 Review Configuration

Datapoint Name LogSine

Expression to apply log(sinusoid)


Enabled ☒

Previous Done

In the case of our expression filter we should add the expression we wish to execute *log(sinusoid)* and the name of the datapoint we wish to put the result in, *LogSine*. We can also choose to enable or disable the execution of this filter. We will enable it and click on *Done* to complete adding the filter.

Click on *Save* in the south edit dialog and our filter is now installed and running.

If we select the *Assets & Readings* option from the menu bar we can examine the sinusoid asset and view a graph of that asset. We will now see a second datapoint has been added, *LogSine* which is the result of executing our expression in the filter.



A second filter can be added in the same way, for example a *metadata* filter to create a pipeline. Now when we go back and view the south service we see two applications in the dialog.

Sine South Service

Asset name

sinusoid

Enabled

☒

[Show Advanced Config](#)

Applications +

≡ MyExpression

▼

≡ Location

▼

Cancel

Save

Service Info

http://localhost:37799

Export Readings

Delete Service

Reordering Filters

The order in which the filters are applied can be changed in the south service dialog by clicking and dragging one filter above another in the *Applications* section of dialog.

Sine South Service

Asset name

sinusoid

Enabled

☒

[Show Advanced Config](#)

Applications +

≡ Location

≡ MyExpression

Cancel

Save

Service Info

http://localhost:37799


Export Readings

Delete Service

Filters are executed in a top to bottom order always. It may not matter in some cases what order a filter is executed in, in others it can have significant effect on the result.

Editing Filter Configuration

A filters configuration can be altered from the south service dialog by selecting the down arrow to the right of the filter name. This will open the edit area for that filter and show the configuration that can be altered.



You can also remove a filter from the pipeline of filters by select the trash can icon at the bottom right of the edit area for the filter.

3.3.2 Adding Filters To The North

Filters can also be added to the north in the same way as the south. The same set of filters can be applied, however some may be less useful in the north than in the south as they apply to all assets that are sent north.

In this example we will use the metadata filter to label all the data that goes north as coming via a particular FogLAMP instance. As with the *South* service we start by selecting our north task from the *North* menu item in the left-hand menu bar.

PI Server

PI Web API Password

....

PI Web API Kerberos keytab file

piwebapi_kerberos_https.keytab

OCS Namespace

name_space

OCS Tenant ID

ocs_tenant_id

OCS Client ID

ocs_client_id

OCS Client Secret

....

Enabled

☒

[Show Advanced Config](#)

Exclusive

☒

Interval

0

00:00:30

Applications +

Cancel

Save

Delete Instance

At the bottom of the dialog there is a *Applications* area, you may have to scroll the dialog to find it, click on the + icon. A selection dialog appears that allows you to select the filter to use. Select the *metadata* filter.

The screenshot shows a web interface titled "PI Server" with a close button in the top right corner. A progress bar at the top indicates two steps: "1 Plugin Name" (highlighted with a green circle) and "2 Review Configuration". The main content area contains a "Plugin" dropdown menu with options: "fft", "FlirValidity" (with a sub-label "Metadata filter plugin"), "metadata" (highlighted), and "python27". Below the dropdown is a blue link that says "Install from available plugins". Underneath is a "Name" label followed by a text input field containing the word "Floor". At the bottom of the form are two buttons: "Back" and "Next".

After clicking *Next* you will be shown the configuration page for the particular filter you have chosen. We will edit the JSON that defines the metadata tags to add and set a name of *floor* and a value of *1*.

The screenshot shows a window titled "PI Server" with a close button in the top right corner. A progress bar at the top has two steps: "1 Plugin Name" and "2 Review Configuration", with the second step being the active one. Below the progress bar, there is a section titled "Metadata to add" which contains a text area with a JSON snippet:

```
1 {  
2   "floor": "1"  
3 }
```

. Below the text area, there is a checkbox labeled "Enabled" which is checked. At the bottom of the window, there are two buttons: "Previous" and "Done".

After enabling and clicking on *Done* we save the north changes. All assets sent to this PI Server connection will now be tagged with the tag “floor” and value “1”.

Although this is a simple example of labeling data other things can be done here, such as limiting the rate we send data to the PI Server until an *interesting* condition becomes true, perhaps to save costs on an expensive link or prevent a network becoming loaded until normal operating conditions. Another option might be to block particular assets from being sent on this link, this could be useful if you have two destinations and you wish to send a subset of assets to each.

This example used a PI Server as the destination, however the same mechanism and filters may be used for any north destination.

3.4 Some Useful Filters

A number of simple filters are worthy of mention here, a complete list of the currently available filters in FogLAMP can be found in the section .

3.4.1 Scale

The filter *foglamp-filter-scale* applies a scale factor and offset to the numeric values within an asset. This is useful for operations such as changing the unit of measurement of a value. An example might be to convert a temperature reading from Centigrade to Fahrenheit.

3.4.2 Metadata

The filter *foglamp-filter-metadata* will add metadata to an asset. This could be used to add information such as unit of measurement, machine data (make, model, serial no) or the location of the asset to the data.

3.4.3 Delta

The filter *foglamp-filter-delta* allows duplicate data to be removed, only forwarding data that changes by more than a configurable percentage. This can be useful if a value does not change often and there is a desire not to forward all the *similar* values in order to save network bandwidth or reduce storage requirements.

3.4.4 Rate

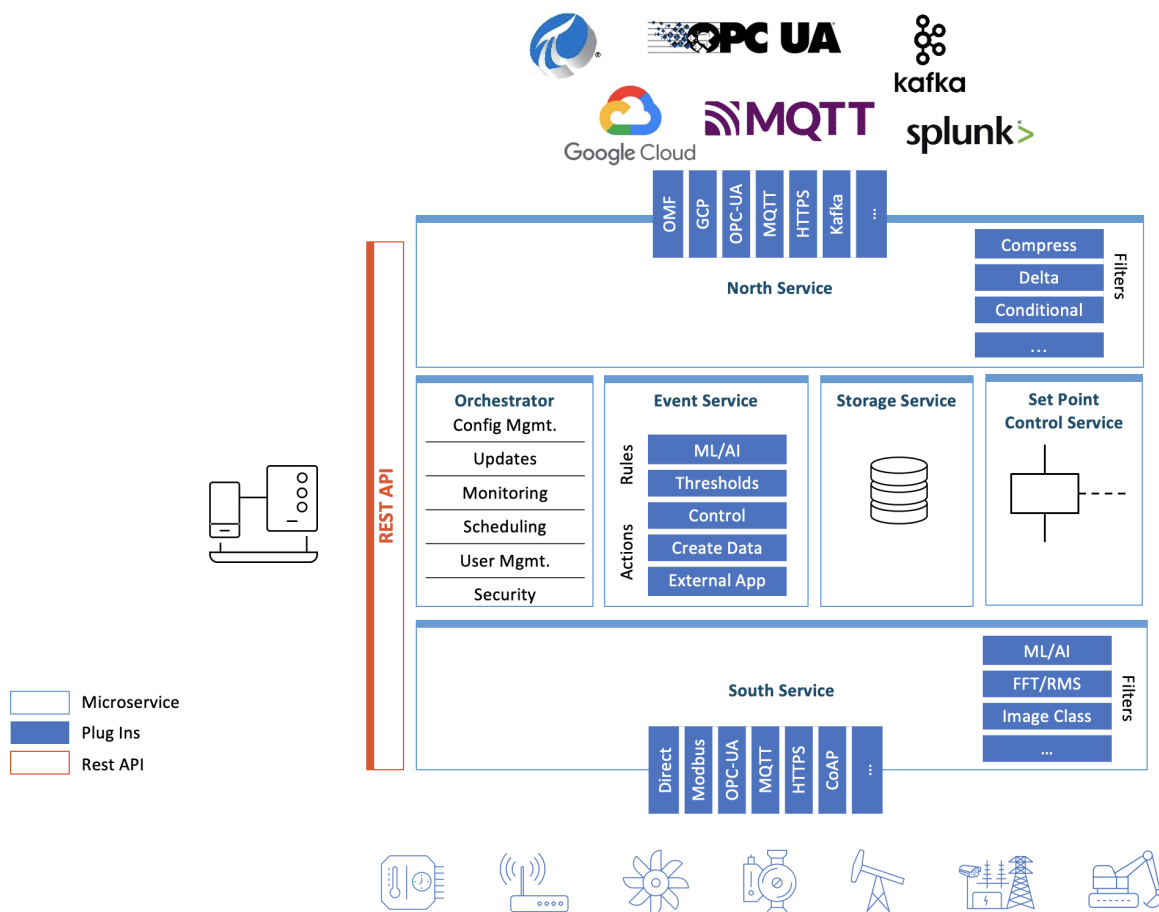
The filter *foglamp-filter-rate* is similar to the delta filter above, however it forwards data at a fixed rate that is lower the rate of the oncoming data but can send full rate data should an *interesting* condition be detected. The filter is configured with a rate to send data, the values sent at that rate are an average of the values seen since the last value was sent.

A rate of one reading per minute for example would average all the values for 1 minute and then send that average as the reading at the end of that minute. A condition can be added, when that condition is triggered all data is forwarded at full rate of the incoming data until a further condition is triggered that causes the reduced rate to be resumed.

FOGLAMP ARCHITECTURE

The following diagram shows the architecture of FogLAMP:

- Components in blue are **plugins**. Plugins are light-weight modules that enable FogLAMP to be extended. There are a variety of types of plugins: south-facing, north-facing, storage engine, filters, event rules and event delivery mechanisms. Plugins can be written in python (for fast development) or C++ (for high performance).
- Components with a blue line at the top of the box are **microservices**. They can co-exist in the same operating environment or they can be distributed across multiple environments.



4.1 FogLAMP Core

The Core microservice coordinates all of the FogLAMP operations. Only one Core service can be active at any time.

Core functionality includes:

Scheduler: Flexible scheduler to bring up processes.

Configuration Management: maintain configuration of all FogLAMP components. Enable software updates across all FogLAMP components.

Monitoring: monitor all FogLAMP components, and if a problem is discovered (such as an unresponsive microservice), attempt to self-heal.

REST API: expose external management and data APIs for functionality across all components.

Backup: FogLAMP system backup and restore functionality.

Audit Logging: maintain logs of system changes for auditing purposes.

Certificate Storage: maintain security certificates for different components, including south services, north services, and API security.

User Management: maintain authentication and permission info on FogLAMP administrators.

Asset Browsing: enable querying of stored asset data.

4.2 Storage Layer

The Storage microservice provides two principal functions: a) maintenance of FogLAMP configuration and run-time state, and b) storage/buffering of asset data. The type of storage engine is pluggable, so in installations with a small footprint, a plugin for SQLite may be chosen, or in installations with a high number of concurrent requests and larger footprint Postgresql may be suitable. In micro installations, for example on Edge devices, or when high bandwidth is required, an in-memory temporary storage may be the best option.

4.3 South Microservices

South microservices offer bi-directional communication of data and metadata between Edge devices, such as sensors, actuators or PLCs and FogLAMP. Smaller systems may have this service installed onboard Edge devices. South components are typically deployed as always-running services, which continuously wait for new data.

4.4 North Microservices

North microservices offer bi-directional communication of data and metadata between the FogLAMP platform and larger systems located locally or in the cloud. Larger systems may be private and public Cloud data services, proprietary solutions or FogLAMP instances with larger footprints. North components are typically deployed as one-shot tasks, which periodically spin up and send data which has been batched, then spin down. However, they can also be deployed as continually-running services.

4.5 Filters

Filters are plugins which modify streams of data that flow through FogLAMP. They can be deployed at ingress (in a South service), or at egress (in a North service). Typically, ingress filters are used to transform or enrich data, and egress filters are used to reduce flow to northbound pipes and infrastructure, i.e. by compressing or reducing data that flows out. Multiple filters can be applied in “pipelines”, and once configured, pipelines can be applied to multiple south or north services.

A sample of existing Filters:

Expression: apply an arbitrary mathematical equation across one or more assets.

Python35: run user-specified python code across one or more assets.

Metadata: apply tags to data, to note the device/location it came from, or to attribute data to a manufactured part.

RMS/Peak: summarize vibration data by generating a Root Mean Squared (RMS) across n samples.

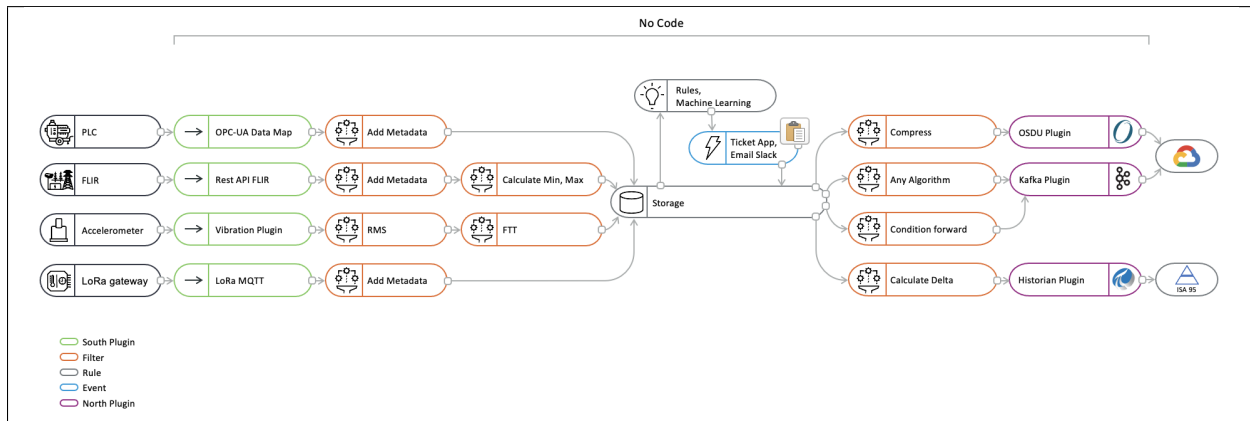
FFT: generate a Fast Fourier Transform (FFT) of vibration data to discover component waveforms.

Delta: Only send data that has changed by a specified amount.

Rate: buffer data but don't send it, then if an error condition occurs, send the previous data.

Contrast: Enhance the contrast of image type data

Filters may be concatenated together to form a data pipeline from the data source to the storage layer, in the south microservice. Or from the storage layer to the destination in the north.



This allows for data processing to be built up via the graphical interface of FogLAMP with little or no coding required. Filters that are applied in a south service will affect all out going streams whilst those applied in the north only affect the data that is sent on that particular connection to an external system.

4.6 Event Service

The event engine maintains zero or more rule/action pairs. Each rule subscribes to desired asset data, and evaluates it. If the rule triggers, its associated action is executed.

Data Subscriptions: Rules can evaluate every data point for a specified asset, or they can evaluate the minimum, maximum or average of a specified window of data points.

Rules: the most basic rule evaluates if values are over/under a specified threshold. The Expression plugin will evaluate an arbitrary math equation across one or more assets. The Python35 plugin will execute user-specified python code to one or more assets.

Actions: A variety of delivery mechanisms exist to execute a python application, create arbitrary data, alter the configuration of FogLAMP, send a control message, raise a ticket in a problem ticking system or email/slack/hangout/communicate a message.

4.7 Set Point Control Service

FogLAMP is not designed to replace real time control systems, it does however allow for non-time-critical control using the control microservice. Control messages may originate from a number of sources; the north microservice, the event service, the REST API or from scheduled events. It is the job of the control service to route these control messages to the correct destination. It also provides a simple form of scripting to allow control messages to generate chains of writes and operations on the south service and also modify the configuration of the FogLAMP itself.

4.8 REST API

The FogLAMP API provides methods to administer FogLAMP, and to interact with the data inside it.

4.9 Graphical User Interface

A GUI enables administration of FogLAMP. All GUI capability is through the REST API, so FogLAMP can also be administered through scripts or other management tools. The GUI contains pages to:

Health: See if services are responsive. See data that's flowed in and out of FogLAMP

Assets & Readings: analytics of data in FogLAMP

South: manage south services

North: manage north services

Notifications: manage event engine rules and delivery mechanisms

Configuration Management: manage configuration of all components

Schedules: flexible scheduler management for processes and tasks

Certificate Store: manage certificates

Backup & Restore: backup/restore FogLAMP

Logs: see system, notification, audit, packages and tasks logging information

Support: support bundle contents with system diagnostic reports

Settings: set/reset connection and GUI related settings

BUFFERING & STORAGE

One of the micro-services that makes up the core of a FogLAMP implementation is the storage micro-service. This is responsible for

- storing the configuration of FogLAMP
- buffering the data read from the south
- maintaining the FogLAMP audit log
- persisting the state of the system

The storage service is configurable, like other services within FogLAMP and uses plugins to extend the functionality of the storage system. These storage plugins provide the underlying mechanism by which data is stored within FogLAMP. FogLAMP can make use of either one or two of these plugins at any one time. If a single plugin is used then this plugin provides the storage for all data. If two plugins are used, one will be for the buffering of readings and the other for the storage of the configuration.

As standard FogLAMP comes with 3 storage plugins

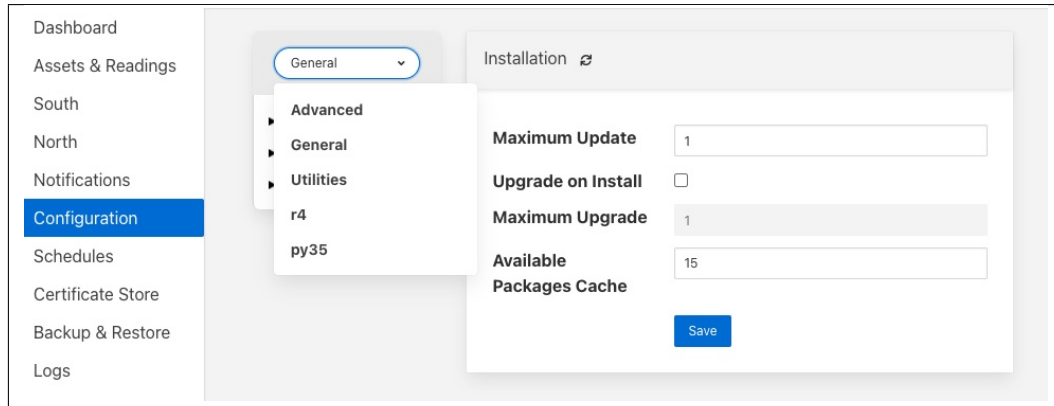
- **SQLite**: A plugin that can store both configuration data and the readings data using SQLite files as the backing store. The plugin uses multiple SQLite database to store different assets, allowing for high bandwidth data at the expense of limiting the number of assets that a single instance can ingest.,
- **SQLiteLB**: A plugin that can store both configuration data and the readings data using SQLite files as the backing store. This version of the SQLite plugin uses a single readings database and is better suited for environments that do not have very high bandwidth data. It does not limit the number of distinct assets that can be ingested.
- **PostgreSQL**: A plugin that can store both configuration and readings data which uses the PostgreSQL SQL server as a storage medium.
- **SQLiteMemory**: A plugin that can only be used to store reading data. It uses SQLite's in memory storage engine to store the reading data. This provides a high performance reading store however capacity is limited by available memory and if FogLAMP is stopped or there is a power failure the buffered data will be lost.

The default configuration uses the SQLite disk based storage engine for both configuration and reading data

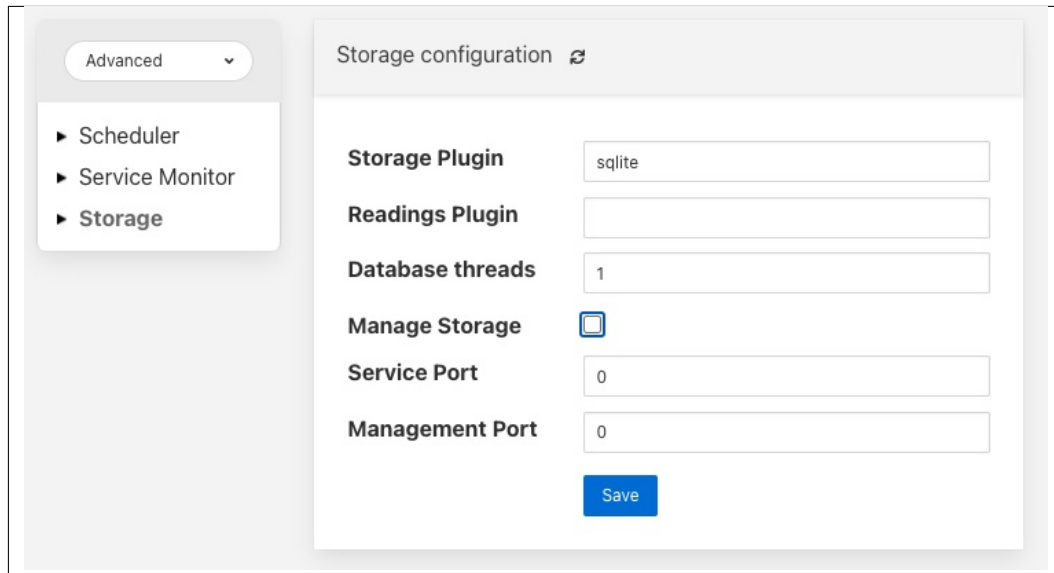
5.1 Configuring The Storage Plugin

Once installed the storage plugin can be reconfigured in much the same way as any FogLAMP configuration, either using the API or the graphical user interface to set the storage engine and its options.

- Using the user interface to configuration the storage, select the *Configuration* item in the left hand menu bar.



- In the category pull down menu select *Advanced*.



- To change the storage plugin to use for both configuration and readings enter the name of the new plugin in the *Storage Plugin* entry field. If *Readings Plugin* is left empty then the storage plugin will also be used to store reading data. The default set of plugins installed with FogLAMP that can be used as *Storage Plugin* values are:
 - *sqlite* - the SQLite file based storage engine.
 - *postgres* - the PostgreSQL server. Note the Postgres server is not installed by default when FogLAMP is installed and must be installed before it can be used.
- The *Readings Plugin* may be set to any of the above and may also be set to use the SQLite In Memory plugin by entering the value *sqlitememory* into the configuration field.

- The *Database threads* field allows for the number of threads used for database housekeeping to be controlled. In normal circumstances 1 is sufficient. If performance issues are seen this can be increased however it is rarely required to be greater than 1 and can have counter productive effects on heavily loaded systems.
- The *Manage Storage* option is only used when the database storage uses an external database server, such as PostgreSQL. Toggling this option on causes FogLAMP to start as stop the database server when FogLAMP is started and stopped. If it is left off then FogLAMP will assume the database server is running when it starts.
- The *Management Port* and *Service Port* options allow fixed ports to be assigned to the storage service. These settings are for debugging purposes only and the values should be set to 0 in normal operation.

Note: Additional storage engines may be installed to extend the set that is delivered with the standard FogLAMP installation. These will be documented in the packages that provide the storage plugin.

Storage plugin configurations are not dynamic and FogLAMP *must* be restarted after changing these values. Changing the plugin used to store readings will *not* cause the data in the previous storage system to be migrated to the new storage system and this data may be lost if it has not been sent onward from FogLAMP.

5.1.1 SQLite Plugin Configuration

The SQLite plugin has a more complex set of configuration options that can be used to configure how and when it creates more database to accommodate more distinct assets. This plugin is designed to allow greater ingest rates for readings by separating the readings for each asset into a database table for that asset. It does however result in limiting the number of distinct assets that can be handled due to the requirement to handle large number of database files.

- **Purge Exclusions:** This option allows the user to specify that the purge process should not be applied to particular assets. The user can give a comma separated list of asset names that should be excluded from the purge process. Note, it is recommended that this option is only used for extremely low bandwidth, lookup data that would otherwise be completely purged from the system when the purge process runs.
- **Pool Size:** The number of connections to create in the database connection pool.
- **No. Readings per database:** This option control how many assets can be stored in a single database. Each asset will be stored in a distinct table within the database. Once all tables within a database are allocated the plugin will use more databases to store further assets.
- **No. databases allocate in advance:** This option defines how many databases are create initially by the SQLite plugin.

- **Database allocation threshold:** The number of unused databases that must exist within the system. Once the number of available databases falls below this value the system will begin the process of creating extra databases.
- **Database allocation size:** The number of databases to create when the above threshold is crossed. Database creation is a slow process and hence the tuning of these parameters can impact performance when an instance receives a large number of new asset names for which it has previously not allocated readings tables.

5.2 Installing A PostgreSQL server

The precise commands needed to install a PostgreSQL server vary for system to system, in general a packaged version of PostgreSQL is best used and these are often available within the standard package repositories for your system.

5.2.1 Ubuntu Install

On Ubuntu or other apt based distributions the command to install postgres:

```
sudo apt install -y postgresql postgresql-client
```

Now, make sure that PostgreSQL is installed and running correctly:

```
sudo systemctl status postgresql
```

Before you proceed, you must create a PostgreSQL user that matches your Linux user. Supposing that user is *<foglamp_user>*, type:

```
sudo -u postgres createuser -d <foglamp_user>
```

The *-d* argument is important because the user will need to create the FogLAMP database.

A more generic command is:

```
sudo -u postgres createuser -d $(whoami)
```

5.3 SQLite Plugin Configuration

The SQLite storage engine has further options that may be used to configure its behavior. To access these configuration parameters click on the *sqlite* option under the *Storage* category in the configuration page.

The screenshot displays the FogLAMP configuration interface for the Storage Plugin. On the left, a sidebar shows a navigation menu with 'Advanced' at the top, followed by 'Scheduler', 'Service Monitor', 'Storage' (which is expanded to show 'sqlite'), and 'sqlite'. The main content area is titled 'Storage Plugin' and contains five configuration parameters, each with a text input field and a 'Save' button at the bottom.

Parameter	Value
Pool Size	5
No. Readings per database	15
No. databases to allocate in advance	3
Database allocation threshold	1
Database allocation size	2

Many of these configuration options control the performance of SQLite and it is important to have some background on how readings are stored within SQLite. The storage plugin stores readings for each distinct asset in a table for that asset. These tables are stored within a database. In order to improve concurrency multiple databases are used within the storage plugin. A set of parameters are used to define how these tables and databases are used.

- **Pool Size:** The number of connections to maintain to the database server.
- **No. Readings per database:** This controls the number of different assets that will be stored in each database file within SQLite.
- **No. databases to allocate in advance:** The number of SQLite databases that will be created at startup.
- **Database allocation threshold:** The point at which new databases are created. If the number of empty databases falls below this value then another set of databases will be created.
- **Database allocation size:** The number of database to allocate each time a new set of databases is required.

The setting of these parameters also imposes an upper limit on the number of assets that can be stored within a FogLAMP instance as SQLite has a maximum limit of 61 databases that can be in use at any time. Therefore the maximum number of readings is 60 times the number of readings per database. One database is reserved for the configuration data.

5.4 Storage Management

FogLAMP manages the amount of storage used by means of purge processes that run periodically to remove older data and thus limit the growth of storage use. The purging operations are implemented as FogLAMP tasks that can be scheduled to run periodically. There are two distinct tasks that are run

- **purge:** This task is responsible for limiting the readings that are maintained within the FogLAMP buffer.
- **system purge:** This task limit the amount of system data in the form of logs, audit trail and task history that is maintained.

5.4.1 Purge Task

The purge task is run via a scheduled called *purge*, the default for this schedule is to run the purge task every hour. This can be modified via the user interface in the *Schedules* menu entry or via the REST API by updating the schedule.

The purge task has two metrics it takes into consideration, the age of the readings within the system and the number of readings in the system. These can be configured to control how much data is retained within the system. Note however that this does not mean that there will never be data older than specified or more rows than specified as purge runs periodically and between executions of the purge task the readings buffered will continue to grow.

The configuration of the purge task can be found in the *Configuration* menu item under the *Utilities* section.

The screenshot shows the FogLAMP configuration interface for the Purge task. On the left, there is a sidebar with a 'Utilities' dropdown menu and two sub-items: 'Purge' and 'Purge System'. The main content area is titled 'Purge the readings, log, statistics history table' and contains the following configuration fields:

- Age Of Data To Be Retained (In Hours):** A text input field with the value '1'.
- Max rows of data to retain:** A text input field with the value '1000000'.
- Retain Unsent Data:** A dropdown menu with the selected option 'purge unsent'.
- Retain Stats History Data (In Days):** A text input field with the value '30'.
- Retain Audit Trail Data (In Days):** A text input field with the value '60'.

A blue 'Save' button is located at the bottom right of the configuration panel.

- **Age Of Data To Be Retained:** This configuration option sets the limit on how old data has to be before it is considered for purging from the system. It defines a value in hours, and only data older than this is considered for purging from the system.
- **Max rows of data to retain:** This defines how many readings should be retained in the buffer. This can override the age of data to retain and defines the maximum allowed number of readings that should be in the buffer after the purge process has completed.
- **Retain Unsent Data:** This defines how to treat data that has been read by FogLAMP but not yet sent onward to one or more of the north destinations for data. It supports a number of options

Retain Unsent Data

Retain Stats History Data (In Days)

☒ **purge unsent**
☐ retain unsent to any destination
☐ retain unsent to all destinations

- **purge unsent:** Data will be purged regardless if it has been sent onward from FogLAMP or not.
- **retain unsent to any destination:** Data will not be purged, i.e. it will be retained, if it has not been sent to any of the north destinations. If it has been sent to at least one of the north destinations then it will be purged.
- **retain unsent to all destinations:** Data will be retained until it has been sent to all north destinations that are enabled at the time the purge process runs. Disabled north destinations are not included in order to prevent them from stopping all data from being purged.

Note: This configuration category will not appear until after the purge process has run for the first time. By default this will be 1 hour after the FogLAMP instance is started for the first time.

5.4.2 System Purge Task

The system purge task is run via a scheduled called *system_purge*, the default for this schedule is to run the system purge task every 23 hours and 50 minutes. This can be modified via the user interface in the *Schedules* menu entry or via the REST API by updating the schedule.

The configuration category for the system purge can be found in the *Configuration* menu item under the *Utilities* section.

Utilities

Purge

Purge System

Configuration of the Purge System

Statistics Retention

Audit Retention

Task Retention

Save

- **Statistics Retention:** This defines the number of days for which full statistics are held within FogLAMP. Statistics older than this number of days are removed and only a summary of the statistics is held.
- **Audit Retention:** This defines the number of day for which the audit log entries will be retained. Once the entries reach this age they will be removed from the system.
- **Task Retention:** This defines the number of days for which history of task execution within FogLAMP is maintained.

Note: This configuration category will not appear until after the system purge process has run for the first time.

ADDITIONAL SERVICES

The following additional services are currently available to extend the functionality of FogLAMP. These are optional services not installed as part of the base FogLAMP installation.

6.1 Bucket Storage Service

The FogLAMP Bucket Storage Service is an optional service that can be installed as part of a FogLAMP instance which provides storage services to other components of the FogLAMP instance. These storage services take the form of a single file of contents that are indexed by a set of key/value pair attributes. The contents of the file can be anything required by the client of the bucket storage service and may be of any size.

The motivation behind the bucket storage service is to provide a mechanism for other components of the FogLAMP system to persist arbitrary objects, which may be any size from relatively small items to large binary objects. The overhead of the bucket storage service is such that storing very small objects may be wasteful on resources.

6.1.1 Limitations

The storage of the content of the buckets is in files on the file system of the machine on which the bucket storage service is running and is thus dependent upon the amount of storage available on that machine.

There are no limitations as to what the contents of the buckets may be, they can be textual objects, binary data or archives of multiple files.

6.1.2 Indexing

When a storage bucket is created a number of key/values pairs are supplied with the content of the bucket, it is these that are used to later locate the bucket for retrieval. These key values pairs must consist of string values and each bucket must have a unique set of attributes. Each bucket must have at least one attribute given when it is created.

The creation of a bucket will return a unique ID for a bucket that can be used to retrieve the contents of the bucket.

Searches of the bucket storage service operate against the attributes of the bucket. A subset of the key/value pairs are given in order to find the set of matching storage buckets. The system then returns the unique identifiers for all those buckets that match this subset of attributes.

6.1.3 API

The bucket service server a REST API to all of the other services within FogLAMP, as well as offering a modified version of that API via the public REST API of FogLAMP.

A service within FogLAMP wishing to use the bucket service should first contact the FogLAMP core with which it is registered and request the registration record for the bucket storage. Note that there is nothing stopping one FogLAMP service having multiple bucket storage service, although normally each instance would have a single bucket storage service. Access to the bucket service API would then be via the service port of the bucket service API.

Components outside of the FogLAMP instance may also access the bucket service via the FogLAMP public API, in this case a modified form of the service API would be used with the */foglamp/extension* prefix added to the URL path of the REST call.

The examples below assume that the FogLAMP instance is running on a host called *foglamp.local*, this name should be substituted with the real address or hostname of the FogLAMP machine. The examples also assume all the FogLAMP service are running on the same host, hence the localhost address of 127.0.0.1 is used when accessing the service port. Again this should be substituted with the address obtained from the service registry.

Adding a bucket

A bucket is added into the bucket store using the *POST* method of the API. The URL for adding a bucket is */bucket* if using the service API from within FogLAMP or */foglamp/extension/bucket* if using the FogLAMP public API. The payload associated with this call should be a multipart message. The data itself is sent in a part named *bucket*, the content of which is an octet stream. The attributes, used for indexing, should be carried in a part named *attributes* and are expressed as a single JSON object with the set of key value pairs.

If using the curl command to add a new bucket, the command would be of the form

```
curl -v -X POST -F "bucket=@data.csv" -F 'attributes=@attributes.json' http://foglamp.  
↪local:8081/foglamp/extension/bucket
```

The contents of the *attributes.json* file would be as follows

```
{  
    "name"    : "Tag Mapping",  
    "type"    : "TagSheet",  
    "device"  : "Acme Press"  
}
```

This will create a bucket whose contents are the file given in the *bucket=* section with three attributes, *name*, *type* and *device*. These attributes names are arbitrary and have no internal significance within the bucket service, they should however be carefully chosen as they must be unique, as a set, for each bucket and they will be used to match particular buckets.

The return payload, upon successfully adding a bucket is a JSON document that contains the unique ID of the bucket.

```
{  
    "bucket" : 3293  
}
```

The above example shows the URL for adding the bucket via the public REST API of FogLAMP, to add it via the service API of the bucket service, which is the option that would be used only by other services within the same FogLAMP instance, the example would be

```
curl -v -X POST -F "bucket=@data.csv" -F 'attributes=@attributes.json' http://127.0.0.  
↪1:<service_port>/bucket
```


Where `<service_port>` is the service port of the bucket storage service that has previously be obtain from the service registry of the FogLAMP instance using the management REST API.

Error Responses

A *bad request* response will be generated in the following conditions

- The request is missing the attributes set
- The request is missing the bucket contents
- The attribute JSON document could not be parsed
- An empty set of attributes was given
- One of the attributes had a value that was not a string
- A bucket exists that has the same set of attributes as the request

The response will also contain an explanatory message, within a JSON document that describes the nature of the issue.

```
{
  "message" : "A bucket must have at least one attribute"
}
```

Retrieving Bucket Contents

The API to retrieve the contents of a bucket is a simple *GET* HTTP request to either the public REST API of FogLAMP or the Bucket Storage Service's service API from other FogLAMP services. The call requires that you give the unique identifier of the bucket that was returned as the result of the add request or by using the attribute matching API entry point.

```
curl http://foglamp.local:8081/foglamp/extension/bucket/3293 -o b3293
```

The response is an octet stream and may be binary if the bucket contains a binary file. The above example shows using the public REST API, the call from within a FogLAMP instance would use the service port of the Bucket Storage Service and a slightly modified URL.

```
curl http://127.0.0.1:<service_port>/bucket/3293 -o b3293
```

Where `<service_port>` is the service port of the bucket storage service that has previously be obtain from the service registry of the FogLAMP instance using the management REST API.

Error Responses

A *bad request* response will be generated in the following conditions

- The identifier given does not related to any bucket stored within the system
- The bucket contents could not be read

The response will also contain an explanatory message, within a JSON document that describes the nature of the issue.

```
{
  "message" : "Invalid bucket id"
}
```

Matching Bucket Attributes

Normally it is expected that clients of the Bucket Storage Server would not know the unique identifiers for the buckets of interest and instead would search for them using the attributes of the buckets. This is done using the *match* API entry point. The entry point is given a set of attributes to match against and will return all those buckets that have the attributes given in the set with identical values. This set may be a subset of the attributes a bucket actually has associated with it.

Matching is done using a PUT method on the REST API with a payload of those attributes to match. The payload is a JSON document similar to the one given in the API call to create the bucket, however it may be a subset of the attributes that the bucket has. The following payload would return all those buckets that have a *type* attribute whose value is *TagSheet*.

```
{
  "attributes" : {
    "type" : "TagSheet"
  }
}
```

The example curl command that would be used to perform the match, via the public FogLAMP REST API would be

```
curl -X PUT http://foglamp.local:8081/foglamp/extension/bucket/match -d@match.json
```

Assuming the file *match.json* has the contents should above. In the same way as shown in the other example, other services within the same FogLAMP instance would use the service API of the Bucket Storage Server itself to get the matching set of buckets.

```
curl -X PUT http://127.0.0.1:<service_port>/bucket/match -d@match.json
```

The return of this call is a JSON document that lists all the matching buckets stored within the Bucket Storage Service.

```
{
  "matches" : [
    {
      "bucket" : 3293,
      "file" : "/usr/local/foglamp/data/buckets/3/3293",
      "attributes" : {
        "name" : "Tag Mapping",
        "type" : "TagSheet",
        "device" : "Acme Press"
      }
    },
    {
      "bucket" : 2712,
      "file" : "/usr/local/foglamp/data/buckets/2/2712",
      "attributes" : {
        "name" : "Drier 002",
        "type" : "TagSheet",
        "device" : "Drier"
      }
    }
  ]
}
```

The result of this can then be used to retrieve the particular bucket of interest. Note this return payload includes the internal filename in which the bucket is stored, callers should not normally use this information as the Bucket Storage Service may not be running on the same host and the filename will then not be valid locally. All access should be via the REST API.

Error Responses

A *bad request* response will be generated in the following conditions

- The attributes are missing from the payload
- One of the attribute values is not a string
- The attributes JSON document could not be parsed

The response will also contain an explanatory message, within a JSON document that describes the nature of the issue.

```
{
  "message" : "Badly formed payload, missing attributes to match"
}
```

Updating Bucket Attributes

It is possible to update bucket attributes by overwriting the values of existing attributes or adding new attributes, it is not possible to remove attributes or to change the content of the bucket itself. The API entry point must be given the unique identifier for the bucket that should be updated and the set of new attributes as a JSON document.

The API entry point for the public FogLAMP API is

```
curl -X PUT http://foglamp.local:8081/foglamp/extension/bucket/3293 -d '{ "attributes
↪" : { "status" : "current" } }'
```

Where the 3293 is the unique identifier for the bucket to be updated. The above example will add a new attribute *status* to the bucket with the value of *current*. If the bucket already has an attribute names *status* then the value of that attribute will be updated.

Again a service that is part of a FogLAMP instance should use the service API port of the Bucket Storage Service itself rather than the public REST API.

```
curl -X PUT http://127.0.0.1:<service_port>/bucket/3293 -d '{ "attributes" : { "status
↪" : "current" } }'
```

Error Responses

A *bad request* response will be generated in the following conditions

- The identifier is not a valid identifier of a bucket stored in the Bucket Storage Server
- The attributes are missing from the update
- The attributes JSON document could not be parsed
- One of the attributes has a value that is not a string

The response will also contain an explanatory message, within a JSON document that describes the nature of the issue.

```
{
  "message" : "Badly formed payload, missing attributes to update"
}
```

Deleting a Bucket

To delete a bucket the DELETE method should be used, giving the unique identifier of the bucket to be deleted in the URL

```
curl -X DELETE http://localhost:8081/foglamp/extension/bucket/3292
```

The above example shows the use of the public FogLAMP API to delete bucket 3292 from the system. As with all entry points internal services would use the Bucket Storage Service API via the service port of the service itself.

```
curl -X DELETE http://localhost:<service_port>/bucket/3292
```

Error Responses

A *bad request* response will be generated in the following conditions

- The identifier given does not related to any bucket stored within the system

The response will also contain an explanatory message, within a JSON document that describes the nature of the issue.

```
{
  "message" : "Invalid bucket id"
}
```

6.2 Notifications Service

FogLAMP supports an optional service, known as the notification service that adds an event engine to the FogLAMP installation. The notification services observed data as it flows into the FogLAMP storage service buffer and processes that data against a set of rules that are configurable by the user to determine if an event has occurred. Events may be either when a condition that was previously not met being is, or a condition that was previously met becoming no longer true. The notification service can then send a notification when an event occurs or, in the case of a condition that is met, it can send notifications as long as that condition is met.

The notification services operates on data that is in the storage layer, and is independent of the individual south services. This means that the notification rules can use data from several south services to evaluate if a condition has occurred. Also the data that is observed by the notification is after any filtering rules have been applied in the south services but before any filtering that occurs in the north tasks. The mechanism used to allow the notification service to observe data is that the notifications register with the storage service to be given the values for particular assets as they arrive at the storage service. A notification may register for several assets and is free to buffer that data internally within the notification service. This registration does not impact how the data that is requested is treated in the rest of the system; it will still for example follow the normal processing rules to be sent onward to the north systems.

6.2.1 Notifications

The notification services manages *Notifications*, these are a set of parameters that it uses to determine if an event has occurred and a notification delivery should be made on the basis of that event.

A notification within the notification service consists of;

- A notification rule plugin that contains the logic to evaluate if a rule has been triggered, thus creating an event.
- A set of assets that are required to execute a notification rule.
- Information that defines how the data for each asset should be delivered to the notification rule.

- Configuration for the rule plugin that customizes that logic to this notification instance.
- A delivery plugin that provides the mechanism to delivery an event to destination for the notification.
- Configuration that may be required for the delivery plugin to operate.

Notification Rules

Notification rules are the logic that is used by the notification to determine if an event has occurred or not. An event is basically based on the values of a number of attributes, either at a single point in time or over a period of time. The notification services is delivered with one built in rule, this is a very simple rule called the *threshold rule* it simply looks at a single asset to determine if the value of a datapoint within the asset goes above or below a set value.

A notification rule has associated with it a set of configuration options, these define how the plugin behaves but also what data the plugin requires to execute the evaluation logic within the plugin. These configuration parameters can be divided into two sets; those items that define the data the rule requires from the notification service itself and those that relate directly to the logic of the rule.

A rule may work across one or more assets, the assets it requires are configured in the rule configuration and passed the the notification service to enable the service to subscribe to those assets and be sent that data by the storage service. A rule plugin may ask for every value of the asset as it changes or it may ask for a window of data. A window is defined as the values of an asset within a given time frame. An example might be the last 10 minutes of values. In the case of the window the rule may be passed the average value, minimum, maximum or all values in that window. The requirements about how data is delivered to a rule may be hard coded within the logic of a rule or may be part of the configuration a user of the rule should provide.

The second type of configuration parameter a rule might include are those that control the logic itself, in the example of the *threshold rule* this would be the threshold value itself and the control if the event is considered to have triggered if the value is above or below the threshold.

The section contains a full list of currently available rule plugins for FogLAMP. As with other plugin types they are designed to be easily written by end users and developers, a guide is available for anyone wishing to write a notification rule plugin of their own.

Notification Types

Notifications can be delivered under a number of different conditions based on the state returned from a notification rule and how it related to the previous state returned by the notification rule, this is known as the notification type. A notification may be one of three types, these types are used to define when and how often notification are delivered.

One shot

A one shot notification is sent once when the notification triggers but will not be resent again if the notification triggers on successive evaluations. Once the evaluation does not trigger, the notification is cleared and will be sent again the next time the notification rule triggers.

One shot notifications may be further tailored with a maximum repeat frequency, e.g. no more than once in any 15 minute period.

Toggle

A toggle notification is sent when the notification rule triggers and will not be resent again until the rule fails to trigger, in exactly the same way as a one shot trigger. However in this case when the notification rule first stops triggering a cleared notification is sent.

Again this may be modified by the addition of a maximum repeat frequency.

Retriggered

A retriggered notification will continue to be sent when a notification rule triggers. The rate at which the notification is sent can be controlled by a maximum repeat frequency, e.g. send a notification every 5 minutes until the condition fails to trigger.

Notification Delivery

The notification service does not natively support any form of notification delivery, it relies upon a notification delivery plugin in order to delivery a notification of an event to a user or external system that should be alerted to the event that has occurred. Typical notification deliveries might be to alert a user via some form of paging or messaging system, push an event to an external application by sending some machine level message, execute an external program or code segment to make an action occur, switching on an indication light or in extreme cases maybe shutting down a machine for which a critical fault has been detected. The section contains a full list of currently available notification delivery plugins, however like other plugins these are easily extended and a guide is available for writing notification plugins to extend the available set of plugins.

6.2.2 Installing the Notification Service

The notification service is not part of the base FogLAMP installation and is not a plugin, it is a separate microservice dedicated to the detection of events and the sending of notifications.

Installing Notification Service Package

If you are using the packaged binaries for you system then you can use the package manager to install the *foglamp-service-notification* package. The exact command depends on your package manager and how you obtained your packages.

If you downloaded you packages then you should navigate to the directory that contains your package files and run the package manager. If you have deb package files run the command

```
$ sudo apt -y install ./foglamp-service-notification-1.7.0-armhf.deb
```

Note: The version number, 1.7.0 may be different on your system, this will depend which version you have downloaded. Also the armhf may be different for your machine architecture. Verify the precise name of your package before running the above command.

If you are using a RedHat or CentOS distribution and have rpm package files then run the command

```
$ sudo yum -y localinstall ./foglamp-service-notification-1.7.0-x86_64.deb
```

Note: The version number, 1.7.0 may be different on your system, this will depend which version you have downloaded. Verify the precise name of your package before running the above command.

If you have configured your system to search a package repository that contains the FogLAMP packages then you can simply run the command

```
$ sudo apt-get -y install foglamp-service-notification
```


On a Debian/Ubuntu system, or

```
$ sudo yum -y install foglamp-service-notification
```

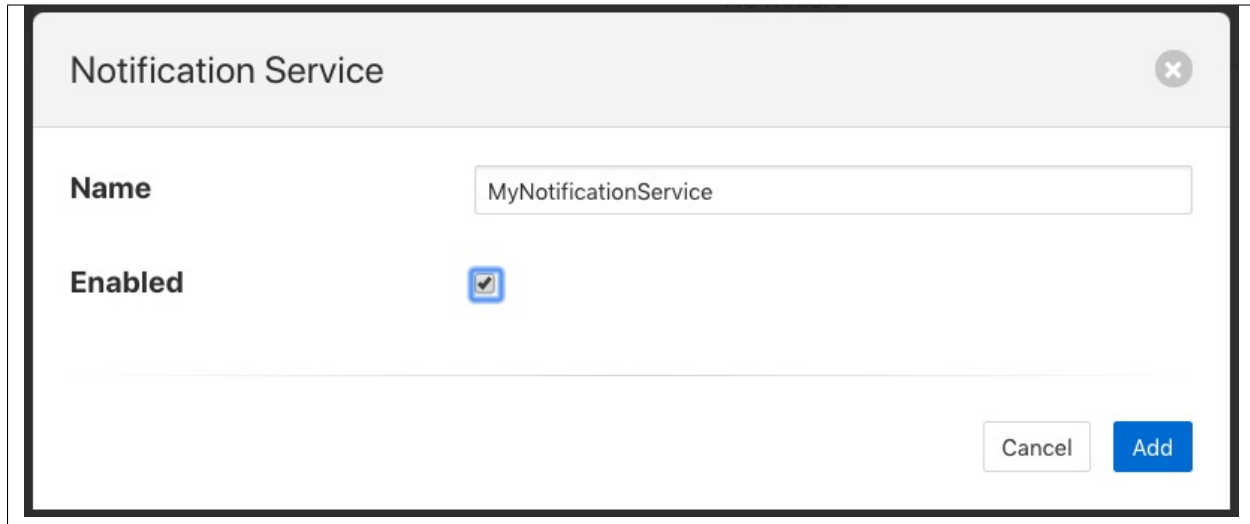
On a RedHat/CentOS system. This will install the latest version of the notification service on your machine.

6.2.3 Starting The Notification Service

Once installed you must configure FogLAMP to start the notification service. This is simply done from the GUI by selecting the *Notifications* option from the left-hand menu. In the page that is then shown you will see a panel at the top that allows you to *add & enable now* the notification service. This only appears if one has not already be added.




Select this link to *add & enable now* the notification service, a new dialog will appear that allows you to name and enable your service.

A screenshot of a 'Notification Service' configuration dialog box. The dialog has a title bar with the text 'Notification Service' and a close button (an 'x' in a circle) on the right. Below the title bar, there are two fields: 'Name' with a text input containing 'MyNotificationService', and 'Enabled' with a checked checkbox. At the bottom right, there are two buttons: 'Cancel' and 'Add'.

6.2.4 Configuring The Notification Service

Once the notification service has been added and enabled a new icon will appear in the *Notifications* page that allows you to configure the notification service. The icon appears in the top right and is in the shape of a gear wheel.



Clicking on this icon will display the notification service configuration dialog.

MyNotificationService Notification Service

Minimum Log Level warning

Maximun number of delivery threads 2

Enabled ☒

Cancel Save

Delete Service

You can use this dialog to control the level of logging that is done from the service by setting the *Minimum Log Level* to the least severity log level you wish to see. All log entries at the select level and of greater severity will be logged.

It is also possible to set the number of threads that will be used for delivering notifications. This defines how many notifications can be delivered in parallel. This only needs to be increased if the delivery process of any of the in use delivery plugins are long running.

The final setting allows you to disable the notification service.

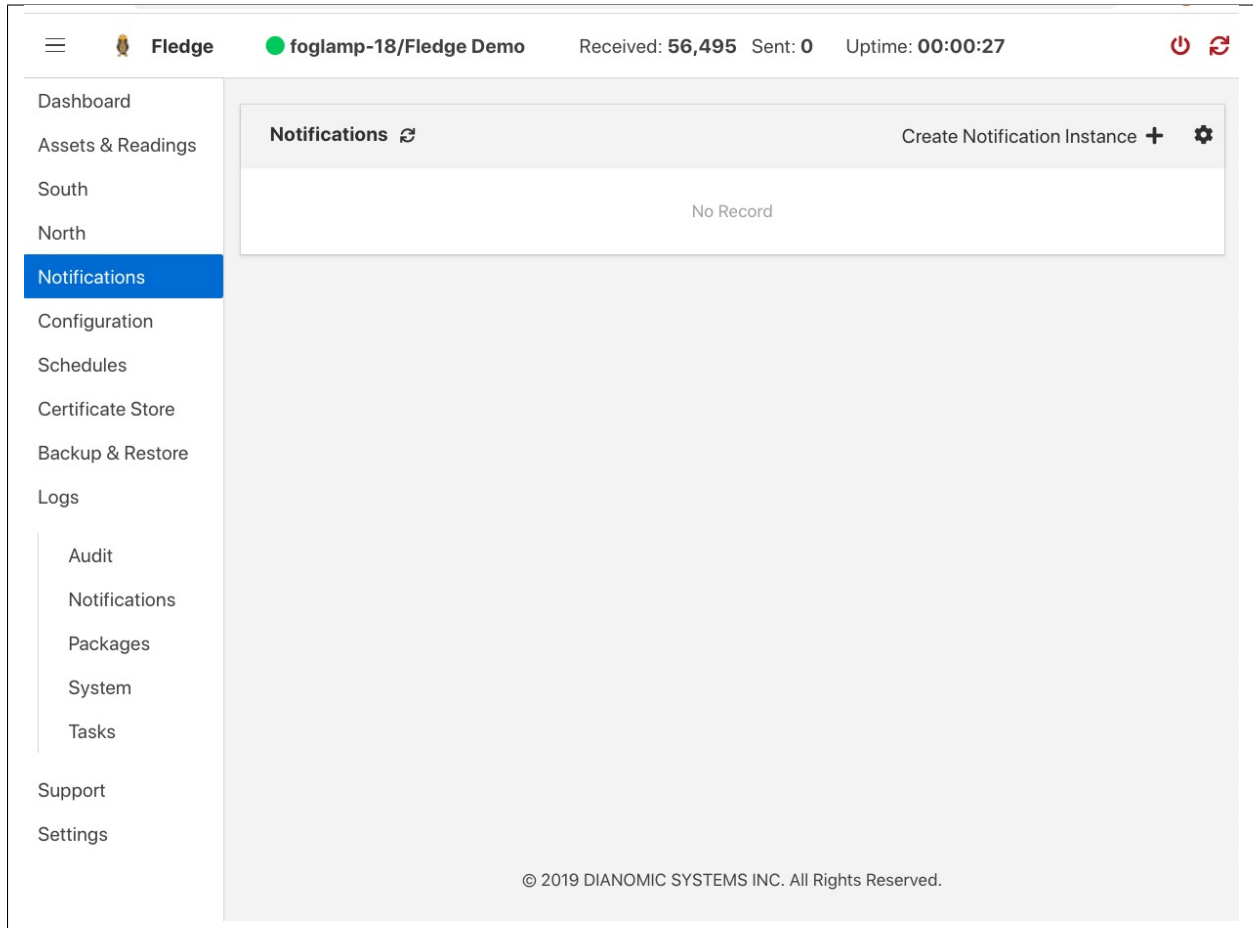
Once you have updated the configuration of the service click on *Save*.

It is also possible to delete the notification service using the *Delete Service* button at the bottom of this dialog.

6.2.5 Using The Notification Service

Add A Notification

In order to add s notification, select the Notifications page in the left-hand menu, an empty set of notifications will appear.



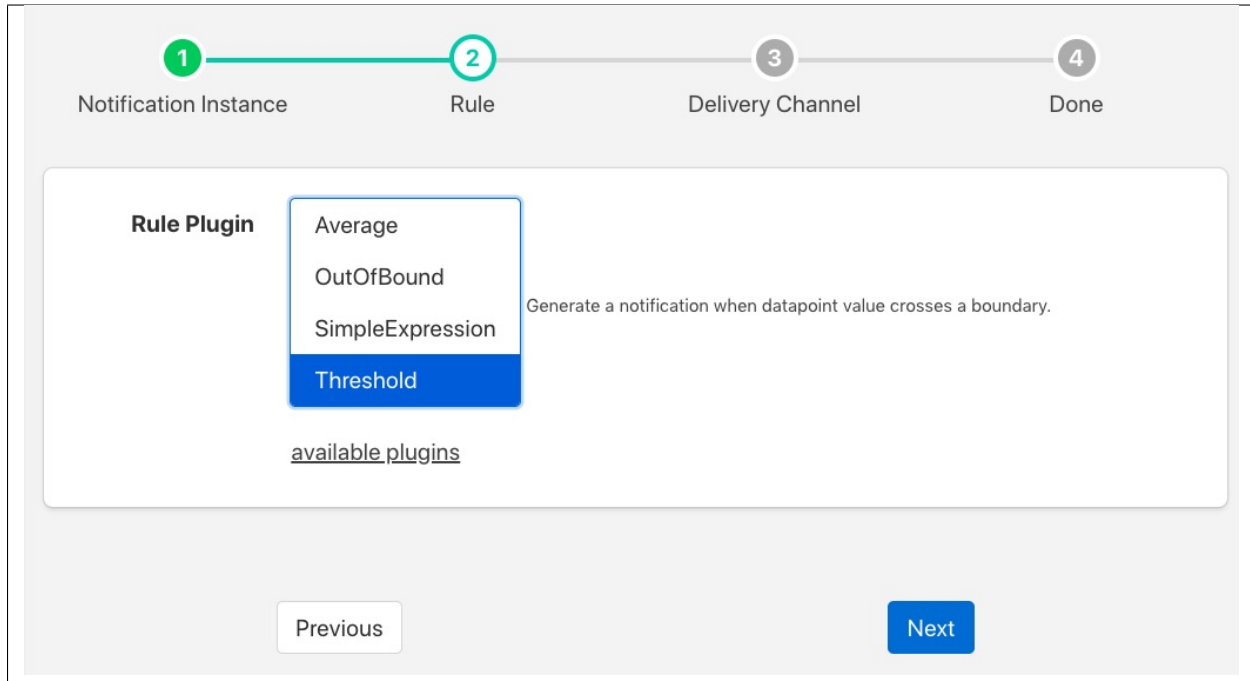
Click on the + icon to add a new notification.

The screenshot shows the FogLAMP Fledge interface. At the top, there's a header with the Fledge logo, a green status indicator, the text "foglamp-18/Fledge Demo", and system metrics: "Received: 92,998 Sent: 0 Uptime: 18:22:45". On the left is a sidebar menu with options: Dashboard, Assets & Readings, South, North, Notifications, Configuration, Schedules, Certificate Store, Backup & Restore, Logs, Audit, Notifications, Packages, System, Tasks, Support, and Settings. The main area displays a progress bar with four steps: 1 (Notification Instance, highlighted in green), 2 (Rule), 3 (Delivery Channel), and 4 (Done). Below the progress bar is a form with two fields: "Name" with a placeholder "name" and "Description" with a placeholder "description". At the bottom of the form are "Back" and "Next" buttons. A copyright notice "© 2019 DIANOMIC SYSTEMS INC. All Rights Reserved." is at the very bottom.

You will be presented with a dialog to enter a name and description for your notification.

This is a close-up of the configuration dialog. The progress bar at the top shows step 1 (Notification Instance) as the active step. The form contains two input fields. The "Name" field has the text "Above0.5" entered. The "Description" field has the text "Above0.5 notification instance" entered. "Back" and "Next" buttons are positioned at the bottom of the dialog.

Enter text for the name you require, a suggested description will be automatically added, however you can modify this to any string you desire. When complete click on the *Next* button to move forwards in the definition process. You can always click on *Previous* to go back a screen and modify what has been entered.



You are presented with the set of installed rules on the system. If the rule you wish to use is not installed and you wish to install it then use the link *available plugins* to be presented with the list of plugins that are available to be installed.

Note: The *available plugins* link will only work if you have added the FogLAMP package repository to the package manager of your system.

When you select a rule plugin a short description of what the rules does will be displayed to the right of the list. In this example we will use the threshold rule that is built into the notification service. Click on *Next* once you have selected the rule you wish to use.

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Asset name FastSine

Datapoint name sinusoid

Condition >

Trigger value 0.5

Evaluation data Single Item

Window evaluation Average

Time window 30

Previous Next

You will be presented with the configuration parameters applicable to the rule you have chosen. Enter the name of the asset and the datapoint within that asset that you wish the rule to operate on. In the case of the *threshold rule* you can also define if you want the rule to trigger if the value is greater than, greater than or equal, less than or less than or equal to a *Trigger value*.

You can also choose to look at *Single Item* or *Window* data. If you choose the later you can then choose to define if the minimum, maximum or average within the window that must cross the threshold value.

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Asset name FastSine

Datapoint name sinusoid

Condition >

Trigger value 0.5

Evaluation data Maximum
Minimum
✓ Average

Window evaluation

Time window 30

Previous Next

Once you have set the parameters for the rule click on the *Next* button to select the delivery plugin to use to delivery the notification data.

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Delivery Plugin

- alexa
- asset
- Blynk
- email

[available plugins](#)

Previous Next

A list of available delivery plugins will be presented, along with a similar link that allows you to install new delivery

plugins if desired. As you select a plugin a short text description will be displayed to the right of the plugin list. In this example we will select the *Slack* messaging platform for the delivery of the notification.

Once you have selected the plugin you wish to use click on the *Next* button.

The screenshot shows a four-step progress bar at the top: 1 Notification Instance, 2 Rule, 3 Delivery Channel (active), and 4 Done. Below the progress bar is a configuration form for a Slack Webhook. The form contains three fields: 'Slack Webhook URL' with the value 'https://hooks.slack.com/services/T2GBZ52AF/BGFNTP7NG/YJxQwiJda5ZHMirFZqUjUci', 'Message Text' with the value 'The value of the sinusoid is greater than 0.5', and 'Enabled' with a checked checkbox. At the bottom of the form are two buttons: 'Previous' and 'Next'.

You will then be presented with the configuration parameters the delivery plugin requires to deliver the notification. In the case of the *Slack* plugin this consists of the webhook that you should obtain from the *Slack* application and a message text that will be sent when the event triggers.

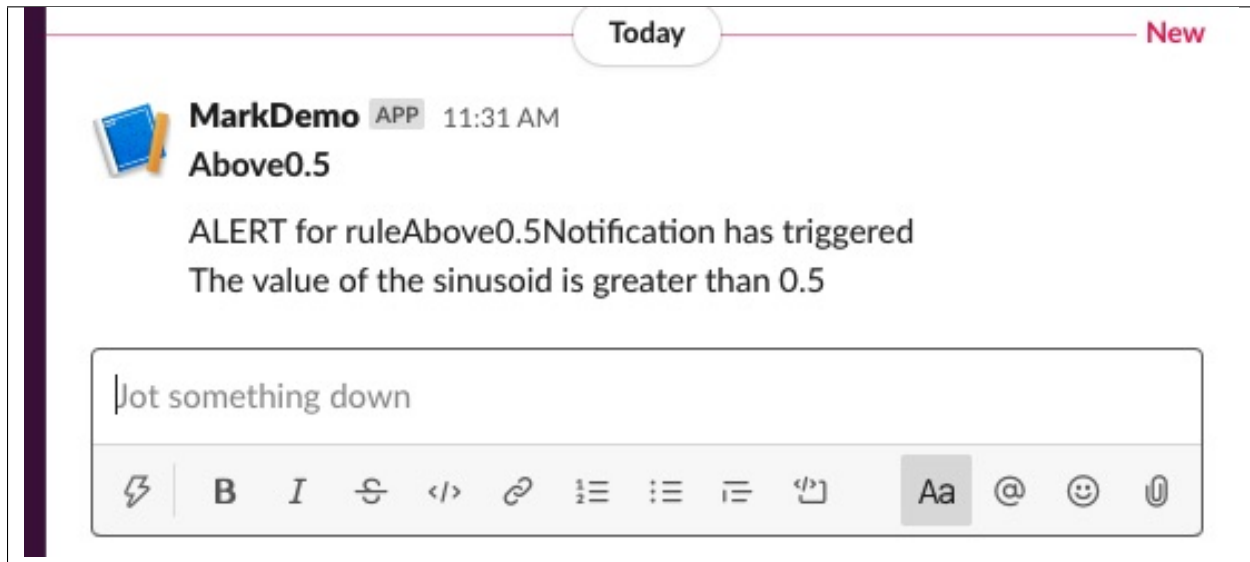
Note: You may disable the delivery of a notification separately to enabling or disabling the notification. This allows you to test the logic of a notification without delivering the notification. Entries will still be made in the notification log when delivery is disabled.

Once you have completed the configuration of the delivery plugin click on *Next* to move to the final stage in setting up your notification.

The screenshot shows the same four-step progress bar, but now step 4 'Done' is active. The configuration form below shows the 'Trigger' dropdown menu open, displaying three options: 'one shot', 'retriggered', and 'toggled'. The 'Enabled' checkbox is checked. A 'Done' button is visible at the bottom right of the form.

The final stage of setting up your configuration is to set the notification type and the retrigger time for the notification. Enable the notification and click on *Done* to complete setting up your notification.

After a period of time, when a *sinusoid* value greater than 0.5 is received, a message will appear in your *Slack* window.



This will repeat at a maximum rate defined by the *Retrigger Time* whenever a value of greater than 0,5 is received.

Notification Log

You can see activity related to the notification service by selecting the *Notifications* option under *Logs* in the left-hand menu.

6.2. Notifications Service

Notification Logs

Count: 13

ALL

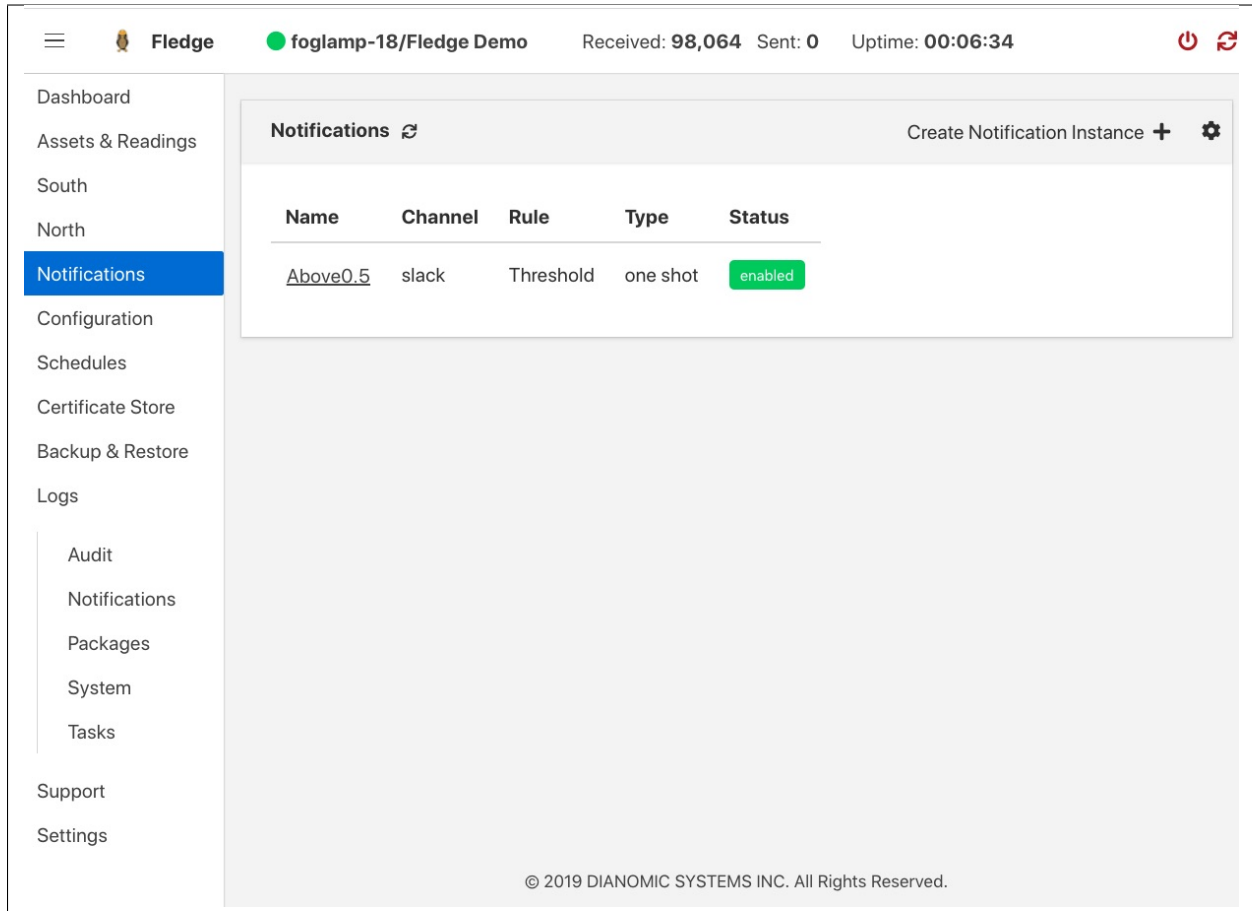
INFORMATION

Name

		Severity	Source
NTFDL - Notification Deleted		INFORMATION	NTFSN
NTFAD - Notification Added		INFORMATION	NTFSN
NTFSN - Notification Sent		INFORMATION	NTFSN
NTFCL - Notification Cleared		INFORMATION	NTFSN
NTFST - Notification Server Startup		INFORMATION	NTFST
NTFSD - Notification Server Shutdown		INFORMATION	NTFST
2020-04-21 09:20:15	MyNotificationService	INFORMATION	NTFST
2020-04-21 09:15:32	Above0.5	INFORMATION	NTFAD
2020-04-20 14:47:38	MyNotificationService	INFORMATION	NTFST
2020-04-20 14:28:17	MyNotificationService	INFORMATION	NTFST
2020-04-20 13:56:17	MyNotificationService	INFORMATION	NTFST
2020-04-20 13:54:27	MyNotificationService	INFORMATION	NTFST
2020-04-20 11:22:38	MyNotificationService	INFORMATION	NTFST

Editing Notifications

It is possible to update existing notifications or remove them using the *Notifications* option from the left-hand menu. Clicking on *Notifications* will bring up a list of the currently defined notifications within the system.



Dashboard

Assets & Readings

South

North

Notifications

Configuration

Schedules

Certificate Store

Backup & Restore

Logs

Audit

Notifications

Packages

System

Tasks

Support

Settings

Notifications

Create Notification Instance +

Name	Channel	Rule	Type	Status
Above0.5	slack	Threshold	one shot	enabled

© 2019 DIANOMIC SYSTEMS INC. All Rights Reserved.

Click on the name of the notification of interest to display the details of that notification and allow it to be edited.

Above0.5

Asset name

sinusoid

Datapoint name

sinusoid

Condition

>

Trigger value

0.5

Evaluation data

Single Item

Window evaluation

Maximum

Time window

30

Delivery Channel - slack

Slack Webhook URL

https://hooks.slack.com/services/T2GBZ52AF/BLH4E9VPX/SpTueiK9t73KSaNSe3

Message Text

The value of the sinusoid is greater than 0.5

Enabled

☒

Cancel

Save

A single page dialog appears that allows you to change any of the parameters of your notification.

Note: You can not change the rule plugin or delivery plugin you are using. If you wish to change either of these then you must delete this notification and create a new one with the desired plugins.

Once you have updated your notification click *Save* to action the changes.

If you wish to delete your notification this may be done by clicking the *Delete* button at the base of the dialog.

FOGLAMP CONTROL FEATURES

FogLAMP supports facilities that allows control of devices via the south service and plugins. This control is known as *set point control* as it is not intended for real time critical control of devices but rather to modify the behavior of a device based on one of many different information flows. The latency involved in these control operations is highly dependent on the control path itself and also the scheduling limitations of the underlying operating system. Hence the caveat that the control functions are not real time or guaranteed to be actioned within a specified time window.

7.1 Control Functions

There are two types of control function supported

- Modify the value in a device via the south service and plugin.
- Request the device to perform an action.

7.1.1 Set Point

Setting the value within the device is known as a set point action in FogLAMP. This can be as simple as setting a speed variable within a controller for a fan or it may be more complete. Typically a south plugin would provide a set of values that can be manipulated, giving each a symbolic name that would be available for a set point command. The exact nature of these is defined by the south plugin.

7.1.2 Operation

Operations, as the name implies provides a means for the south service to request a device to perform an operation, such as reset or re-calibrate. The names of these operations and any arguments that can be given are defined within the south plugin and are specific to that south plugin.

7.2 Control Paths

Set point control may be invoked via a number of paths with FogLAMP

- As the result of a notification within FogLAMP itself.
- As a result of a request via the FogLAMP public REST API.
- As a result of a control message flowing from a north side system into a north plugin and being routed onward to the south service.

Currently only the notification method is fully implemented within FogLAMP.

The use of a notification in the FogLAMP instance itself provides the fastest response for an edge notification. All the processing for this is done on the edge by FogLAMP itself.

7.2.1 Edge Based Control

Edge based control is the name we use for a class of control applications that take place solely within the FogLAMP instance at the edge. The data that is required for the control decision to be made is gathered in the FogLAMP instance, the logic to trigger the control action runs in the FogLAMP instance and the control action is taken within the FogLAMP instance. Typically this will involve one or more south plugins to gather the data required to make the control decision, possibly some filters to process that data, the notification engine to make the decision and one or more south services to deliver the control messages.

As an example of how edge based control might work lets consider the following case.

We have a machine tool that is being monitored by FogLAMP using the OPC/UA south plugin to read data from the machine tools controlling PLC. As part of that data we receive an asset which contains the temperature of the motor which is running the tool. We can assume this asset is called *MotorTemperature* and it contains a single data point called *temperature*.

We also have a fan unit that is able to cool that motor which is controlled via a Modbus interface. The modbus contains one a coil that toggles the fan on and off and a register that controls the speed of the fan. We configure the *foglamp-south-modbus* as a service called *MotorFan* with a control map that will map the coil and register to a pair of set points.

```
{
  "values" : [
    {
      "name" : "run",
      "coil" : 1
    },
    {
      "name" : "speed",
      "register" : 1
    }
  ]
}
```

Control

Control Map

Use Control Map

```

1 {
2   "values": [
3     {
4       "name": "run",
5       "coil": 1
6     },
7     {
8       "name": "speed",
9       "register": 1
10    }
11  ]
12 }
```

If the measured temperature of the motor going above 35 degrees centigrade we want to turn the fan on at 1200 RPM. We create a new notification to do this. The notification uses the *threshold* rule and triggers if the asset *MotorTemperature*, data point *temperature* is greater than 35.

The screenshot displays the 'Rule' configuration step in a four-part sequence: 1. Notification Instance, 2. Rule (active), 3. Delivery Channel, and 4. Done. The configuration form includes the following fields:

- Asset name:** MotorTemperature
- Datapoint name:** temperature
- Condition:** >
- Trigger value:** 0.0
- Evaluation data:** Single Item
- Window evaluation:** Average
- Time window:** 30

Navigation buttons 'Previous' and 'Next' are located at the bottom of the form.

We select the *setpoint* delivery plugin from the list and configure it.

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Service MotorFan

Trigger Value

```

1 {
2   "values": {
3     "run": "1",
4     "speed": 1200
5   }
6 }

```

Cleared Value

```

1 {
2   "values": {
3     "run": "0"
4   }
5 }

```

Enabled ☒

Previous Next

- In *Service* we set the name of the service we are going to use to control the fan, in this case *MotorFan*
- In *Trigger Value* we set the control message we are going to send to the service. This will turn the fan on and set the speed to 1200RPM
- In *Cleared Value* we set the control message we are going to send to turn off the fan when the value falls below 35 degrees.

The plugin is enabled and we go on to set the notification type to toggled, since we want to turn off the fan if the motor cools down, and set a retrigger time to prevent the fan switching on and off too quickly. The notification type and the retrigger time are important parameters for tuning the behavior of the control system and are discussed in more detail below.

If we required the fan to speed up at a higher temperature then this could be achieved with a second notification. In this case it would have a higher threshold value and would set the speed to a higher value in the trigger condition and set it back to 1200 in the cleared condition. Since the notification type is *toggled* the notification service will ensure that these are called in the correct order.

Data Substitution

There is another option that can be considered in our example above that would allow the fan speed to be dependent on the temperature, the use of data substitution in the *setpoint* notification delivery.

Data substitution allows the values of a data point in the asset that caused the notification rule to trigger to be substituted into the values passed in the set point operation. The data that is available in the substitution is the same data that is given to the notification rule that caused the alert to be triggered. This may be a single asset with all of its data points for simple rules or may be multiple assets for more complex rules. If the notification rule is given averaged data then it is these averages that will be available rather than the individual values.

Parameters are substituted using a simple macro mechanism, the name of an asset and data point with in the asset is inserted into the value surrounded by the \$ character. For example to substitute the value of the *temperature* data point of the *MotorTemperature* asset into the *speed* set point parameter we would define the following in the *Trigger Value*

```
{
  "values" : {
    "speed" : "$MotorTemperature.temperature$"
  }
}
```

Note that we separate the asset name from the data point name using a period character.

This would have the effect of setting the fan speed to the temperature of the motor. Whilst allowing us to vary the speed based on temperature it would probably not be what we want as the fan speed is too low. We need a way to map a temperature to a higher speed.

A simple option is to use the macro mechanism to append a couple of 0s to the temperature, a temperature of 21 degrees would result in a fan speed of 2100 RPM.

```
{
  "values" : {
    "speed" : "$MotorTemperature.temperature$00"
  }
}
```

This works, but is a little primitive and limiting. Another option is to add data to the asset that triggers the notification. In this case we could add an expression filter to create a new data point with a desired fan speed. If we were to add an expression filter and give it the expression *desiredSpeed = temperature > 20 ? temperature * 50 + 1200 : 0* then we would create a new data point in the asset called *desiredSpeed*. The value of *desiredSpeed* would be 0 if the temperature was 20 degrees or below, however for temperatures above it would be 1200 plus 50 times the temperature.

This new desired speed can then be used to set the temperature in the *setpoint* notification plugin.

```
{
  "values" : {
    "speed" : "$MotorTemperature.desiredSpeed$"
  }
}
```

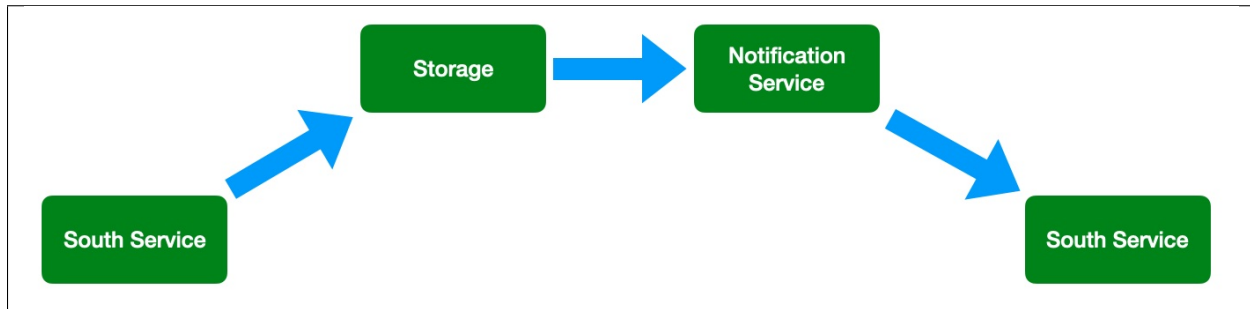
The user then has the choice of adding the desired speed item to the data stored in the north, or adding an asset filter in the north to remove this data point from the data that is sent onward to the north.

Tuning edge control systems

The set point control features of FogLAMP are not intended to replace real time control applications such as would be seen in PLCs that are typically implemented in ladder logic, however FogLAMP does allow for high performance control to be implemented within the edge device. The precise latency in control decisions is dependent on a large number of factors and there are various tuning parameters that can be used to reduce the latency in the control path.

In order to understand the latency inherent in the control path we should first start by examining that path to discover where latency can occur. To do this we will choose a simple case of a single south plugin that is gathering data required by a control decision within FogLAMP. The control decision will be taken in a notification rule and delivered via the *foglamp-notify-setpoint* plugin to another south service.

A total of four services within FogLAMP will be involved in the control path



- the south service that is gathering the data required for the decision
- the storage service that will dispatch the data to the notification service
- the notification service that will run the decision rule and trigger the delivery of the control message
- the south service that will send the control input to the device that is being controlled

Each of these services can add to that latency in the control path, however the way in which these are configured can significantly reduce that latency.

The south service that is gathering the data will typically be either polling a device or obtaining data asynchronously from the device. This will be sent to the ingest thread of the south service where it will be buffered before sending the data to the storage service.

The advanced settings for the south service can be used to trigger how often that data is sent to the storage service. Since it is the storage service that is responsible for routing the data onward to the notification service this impacts the latency of the delivery of the control messages.

[Hide Advanced Config](#)

Maximum Reading Latency (mS)	<input style="width: 90%;" type="text" value="5000"/>
Maximum buffered Readings	<input style="width: 90%;" type="text" value="100"/>
Reading Rate	<input style="width: 90%;" type="text" value="1"/>
Throttle	<input type="checkbox"/>

The above shows the default configuration of a south service. In this case data will not be sent to the storage service until there are either 100 readings buffered in the south service, or the oldest reading in the south service buffer has been in the buffer for 5000 milliseconds. In this example we are reading 1 new readings every second, therefore will send data to the storage service every 5 seconds, when the oldest reading in the buffer has been there for 5000mS.

When it sends data it will send all the data it has buffered, in this case 5 readings as one block. If the oldest reading is the one that triggers the notification we have therefore introduced a 5 second latency into the control path.

The control path latency can be reduced by reducing the *Maximum Reading Latency* of this south plugin. This will of course put greater load on the system as a whole and should be done with caution as it increases the message traffic between the south service and the storage service.

The storage service has little impact on the latency, it is designed such that it will forward data it receives for buffering to the notification service in parallel to buffering it. The storage service will only forward data the notification service has subscribed to receive and will forward that data in the blocks it arrives at the storage service in. If a block of 5 readings arrives at the the storage service then all 5 will be sent to the notification service as a single block.

The next service in the edge control path is the notification service, this is perhaps the most complex step in the journey. The behavior of the notification service is very dependent upon how each individual notification instance has been configured, factors that are important are the notification type, the retrigger interval and the evaluation data options.

The notification type is used to determine when notifications are delivered to the delivery channel, in the case of edge control this might be the *setpoint* plugin or the *operation* plugin. FogLAMP implements three options for the notification type

- **One shot:** A one shot notification is sent once when the notification triggers but will not be resent again if the notification triggers on successive evaluations. Once the evaluation does not trigger, the notification is cleared and will be sent again the next time the notification rule triggers. One shot notifications may be further tailored with a maximum repeat frequency, e.g. no more than once in any 15 minute period.
- **Toggle:** A toggle notification is sent when the notification rule triggers and will not be resent again until the rule fails to trigger, in exactly the same way as a one shot trigger. However in this case when the notification rule first stops triggering a cleared notification is sent. Again this may be modified by the addition of a maximum repeat frequency.
- **Retriggered:** A retriggered notification will continue to be sent when a notification rule triggers. The rate at which the notification is sent can be controlled by a maximum repeat frequency, e.g. send a notification every 5 minutes until the condition fails to trigger.

It is very important to choose the right type of notification in order to ensure the data delivered in your set point control path is what you require. The other factor that comes into play is the *Retrigger Time*, this defines a dead period during which notifications will not be sent regardless of the notification type.

Setting a retrigger time that is too high will mean that data that you expect to be sent will not be sent. For example if you a new value you wish to be updated once every 5 seconds then you should use a retrigger type notification and set the retrigger time to less than 5 seconds.


It is very important to understand however that the retrigger time defines when notifications can be delivered, it does not related to the interval between readings. As an example, assume we have a retrigger time of 1 second and a reading that arrives every 2 seconds that causes a notification to be sent.

- If the south service is left with the default buffering configuration it will send the readings in a block to the storage service every 5 seconds, each block containing 2 readings.
- These are sent to the notification service in a single block of two readings.
- The notification will evaluate the rule against the first reading in the block.
- If the rule triggers the notification service will send the notification via the set point plugin.
- The notification service will now evaluate the rule against the second readings.
- If the rule triggers the notification service will note that it has been less than 1 second since it sent the last notification and it will not deliver another notification.


Therefore, in this case you appear to see only half of the data points you expect being delivered to you set point notification. In order to rectify this you must alter the tuning parameters of the south service to send data more frequently to the storage service.

The final hop in the edge control path is the call from the notification service to the south service and the delivery via the plugin in the south service. This is done using the south service interface and is run on a separate thread in the south service. The result would normally be expected to be very low latency, however it should be noted that plugins commonly protect against simultaneous ingress and egress, therefore if the south service being used to deliver the data to the end device is also reading data from that device, there may be a requirement for the current read to complete before the write operation can commence.

To illustrate how the buffering in the south service might impact the data sent to the set point control service we will use a simple example of sine wave data being created by a south plugin and have every reading sent to a modbus device and then read back from the modbus device. The input data as read at the south service gathering the data is a smooth sine wave,




The data observed that is written to the modbus device is not however a clean sine wave as readings have been missed due to the retrigger time eliminating data that arrived in the same buffer.



Some jitter caused by occasional differences in the readings that arrive in a single block can be seen in the data as well.

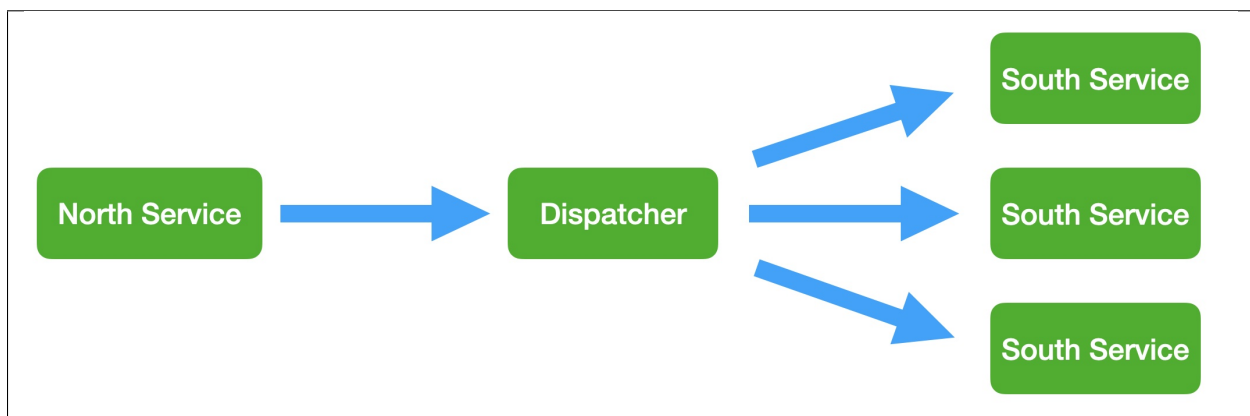
Changing the buffering on the south service to only buffer a single reading results in a much smooth sine wave as can be seen below as the data is seen to transition from one buffering policy to the next.



At the left end of the graph the south service is buffering 5 readings before sending data onward, on the right end it is only buffering one reading.

7.2.2 End to End Control

The end to end control path in FogLAMP is a path that allows control messages to enter the FogLAMP system from the north and flow through to the south. Both the north and south plugins involved in the path must support control operations, a dedicated service, the control dispatcher, is used to route the control messages from the source of the control input, the north service to the objects of the control operations, via the south service and plugins. Multiple south services may receive control inputs as a result of a single north control input.



It is the job of the north plugin to define how the control input is received, as this is specific to the protocol of device to the north of FogLAMP. The plugin then takes this input and maps it to a control message that can be routed by the dispatcher. The way this mappings is defined is specific to each of the north plugins that provide control input.

The control messages that the dispatcher is able to route are defined by the following set

- Write one or more values to a specified south service
- Write one or more values to the south service that ingests a specified asset
- Write one or more values to all south services supporting control
- Run an automation script within the dispatcher service
- Execution an operation on a specified south service
- Execute an operation on the south service that ingests a specified asset
- Execute an operation on all the south services that support control

An example of how a north plugin might define this mapping is shown below

Control Map

```

1 {
2   "nodes": [
3     {
4       "name": "test",
5       "type": "integer",
6       "destination": "service",
7       "asset": "fan0213"
8     }
9   ]
10  }

```

In this case we have an OPCUA north plugin that offers a writable node called *test*, we have defined this as accepting integer values and also set a destination of *service* and a name of *fan0213*. When the OPCUA node *test* is written the plugin will send a control message to the dispatcher to ask it to perform a write operation on the named service.

Alternately the dispatcher can send the request based on the assets that the south service is ingesting. In the following example, again taken from the OPCUA north plugin, we send a value of *EngineSpeed* which is an integer within the OPCUA server that FogLAMP presents to the service that is ingesting the asset *pump0014*.

Control Map

```

1 {
2   "nodes": [
3     {
4       "name": "EngineSpeed",
5       "type": "integer",
6       "asset": "pump0014"
7     }
8   ]
9   }

```

If browsing the OPCUA server which FogLAMP is offering via the north service you will see a node with the browse name *EngineSpeed* which when written will cause the north plugin to send a message to the dispatcher service and ultimately cause the south service ingesting *pump0014* to have that value written to its *EngineSpeed* item. That south service need not be an OPCUA service, it could be any south service that supports control.

The screenshot shows the FogLAMP OPCUA browser interface. The top bar indicates the connection is 'Running' to 'opc.tcp://foglamp-18.local:4840/fledge/server -- urn://fledge.dianomic.com'. The left sidebar shows a tree view with 'Objects' expanded, containing 'Booth1', 'Control', and 'EngineSpeed' (selected). The right pane shows the 'Attributes and References' for the selected node.

Attribute	Value
NodeId	ns=2;i=2001
NodeClass	Variable
BrowseName	2:EngineSpeed
DisplayName	EngineSpeed
Description	EngineSpeed
WriteMask	1
UserWriteMask	1
Value	7
StatusCode	GOOD (0x00000000) ""
ServerTimestamp	02/09/22 15:01:47.48336...
SourceTimestamp	null
ServerPicoSeconds	0
SourcePicoSeconds	0
Value	7
DataType	Int32
ValueRank	Scalar
ArrayDimensions	null
AccessLevel	CurrentRead
UserAccessLevel	CurrentRead
MinimumSamplingInterval	1.0
Historizing	false

It is also possible to get the dispatcher to send the control request to all services that support control. In the case of the OPCUA north plugin this is specified by omitting the other types of destination.

Control Map

```

1 {
2   "nodes": [
3     {
4       "name": "EngineSpeed",
5       "type": "integer"
6     }
7   ]
8 }

```

All south services that support control will be sent the request, these may be of many different types and are free to ignore the request if it can not be mapped locally to a resource to update. The semantics of how the request is treated is determined by the south plugin, each plugin receiving the request may take different actions.

The dispatcher can also be instructed to run a local automation script, these are discussed in more detail below, when a write occurs on the OPCUA node via this north plugin. In this case the control map is passed a script key and name to execute. The script will receive the value *EngineSpeed* as a parameter of the script.

Control Map

```

1 {
2   "nodes": [
3     {
4       "name": "EngineSpeed",
5       "type": "integer",
6       "script" : "SetPumpSpeed"
7     }
8   ]
9 }

```

Note, this is an example and does not mean that all or any plugins will use the exact syntax for mapping described above, the documentation for your particular plugin should be consulted to confirm the mapping implemented by the plugin.

7.3 Control Dispatcher Service

The *control dispatcher* service is a service responsible for receiving control messages from other components of the FogLAMP system and taking the necessary actions against the south services in order to achieve the request result. This may be as simple as forwarding the write or operation request to one to more south services or it may require the execution of an automation script by the *dispatcher service*.

7.3.1 Forwarding Requests

The *service dispatcher* supports three forwarding regimes which may be used to either forward write requests or operation requests, these are;

- Forward to a single service using the name of the service. The caller of the dispatcher must provide the name of the service to which the request will be sent.
- Forward to a single service that is responsible for ingesting a named asset into the FogLAMP system. The caller of the dispatcher must provide the name of an asset, the *service dispatcher* will then look this asset up in the asset tracker database to determine which service ingested the named asset. The request will then be forwarded to that service.
- Forward the request to all south services that are currently running and that support control operations. Note that if a service is not running then the request will not be buffered for later sending.

7.3.2 Automation Scripts

The control dispatcher service supports a limited scripting designed to allow users to easily create sequences of operations that can be executed in response to a single control write operation. Scripts are created within FogLAMP and named externally to any control operations and may be executed by more than one control input. These scripts consist of a linear set of steps, each of which results in one of a number of actions, the actions supported are


- Perform a write request. A new write operation is defined in the step and it may take the form of any of the three styles of forwarding supported by the dispatcher; write to a named service, write to a service providing an asset or write to all south services.
- Perform an operation request on one or all south services. As with the write request above the three forwards of defining the target south service are defined.
- Delay the execution of a script. Add a delay between execution of the script steps.
- Update the FogLAMP configuration. Change the value of a configuration item within the system.
- Execute another script. A mechanism for calling another named script, the named script is executed and then the calling script will continue.

The same data substitution rules described above can also be used within the steps of an automation script. This allows data that is sent to the write or operation request in the dispatcher to be substituted in the steps themselves, for example a request to run a script with the values *param1* set to *value1* and *param2* set to *value2* would result in a step that wrote the value *\$param1\$* to a south service actually writing the value *value1*, i.e the value of *param1*.

Each step may also have associated with it a condition, if specified that condition must evaluate to true for the step to be executed. If it evaluates to false then the step is not executed and execution moves to the next step in the script.

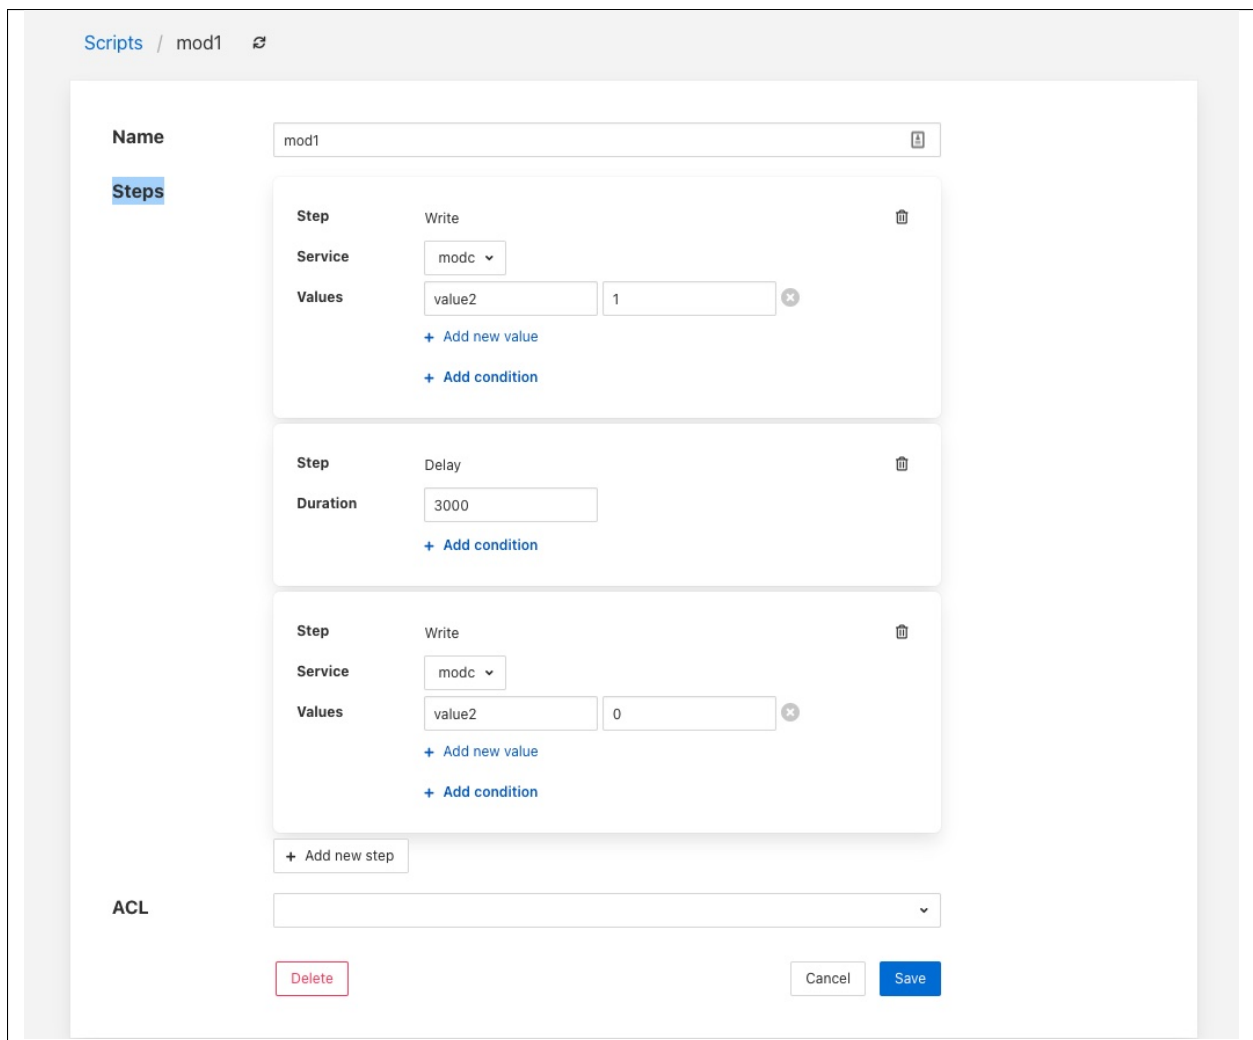
Graphical Interface

The automation scripts are available via the *Control Service* menu item in the left hand menu panel. Selecting this will give you access to the user interface associated with the control functions of FogLAMP. Click on the *Scripts* tab to select the scripts, this will display a list of scripts currently defined within the system and also show an add button icon at the top right corner.



Viewing & Editing Existing Scripts

Simply click on the name of a script to view the script



The steps within the script are each displayed within a panel for that step. The user is then able to edit the script provided they have permission on the script.

There are then a number of options that allow you to modify the script, note however it is not possible to change the type of a step in the script. The user must add a new step and remove the old step they wish to replace.

- To add a new step to a script click on the *Add new step* button
 - The new step be created in a new panel and will prompt for the user to select the step type

- The next step in to process will depend on the type of automation step chosen.
 - * A *Configure* step will request the configuration category to update to be chosen. This is displayed in a drop down menu.

The configuration categories are shown as a tree structure, allowing the user to navigate to the configuration category they wish to change.

Once chosen the user is presented with the items in that configuration category from which to choose.

The screenshot shows a configuration window with the following elements:

- Step:** A dropdown menu set to 'Configure'.
- Category:** A dropdown menu set to 'Kafka'.
- Config Item:** A dropdown menu with the following options:
 - Simple plugin to send data to a Kafka topic (selected)
 - Simple plugin to send data to ...
 - Bootstrap Brokers
 - Kafka Topic
 - Send JSON
 - Data Source
 - A switch that can be used to e...
 - Identifies the specific stream...
 - Filter pipeline
- Value:** A text box next to the 'Config Item' dropdown containing the text 'Kafka'.
- Buttons:**
 - '+ Add new step' (bottom left)
 - 'Delete' (bottom left, red border)
 - 'Cancel' (bottom right)
 - 'Save' (bottom right, blue)

Selecting an item will give you a text box with the current value of that item. Simply type the new value that should be assigned to that item when this step of the script runs into that text box.

- * A *Delay* step will request the duration of the delay. The *Duration* is merely typed into the text box and is expressed in milliseconds.
- * An *Operation* step will request you to enter the name of the operation to perform and then select the service to which the operation request should be sent

The screenshot shows a 'Step' configuration form. The 'Operation' dropdown is open, showing a list of services: modc, Sine, modbus2, Coolant, HTTPIn, Lathe, MQTTTest, Open62451, S2OPCUA, and Spinnaker. The 'Name' field is empty, and the 'Service' dropdown is set to 'Select service'. The 'Parameters' section is empty. The 'Delete' button is highlighted in red. The 'Cancel' and 'Save' buttons are at the bottom right.

Operations can be passed zero or more parameters, to add parameters to an operation click on the *Add parameter* option. A pair of text boxes will appear allowing you to enter the key and value for the parameter.

The screenshot shows the 'Add new parameter' dialog. The 'Operation' dropdown is set to 'Operation'. The 'Name' field is set to 'shutdown'. The 'Service' dropdown is set to 'Lathe'. The 'Parameters' section shows a 'key' field with 'key' and a 'value' field with 'value'. The 'Add new parameter' button is highlighted in blue. The 'Add condition' button is also visible.

To add another parameter simply press the *Add parameter* option again.

- * A *Script* step will request you to choose the name of the script to run from a list of all the currently defined scripts.

The screenshot shows a configuration panel for a 'Script' step. The 'Step' dropdown is set to 'Script'. The 'Name' dropdown is open, showing 'Select script'. The 'Execution' dropdown is open, showing 'Unlock Lathe'. The 'Parameters' section has two links: '+ Add parameter' and '+ Add condition'.

Note that the script that you are currently editing is not included in this list of scripts. You can then choose if you want the execution of this script to block the execution of the current script or to run in parallel with the execution of the current script.

The screenshot shows the same configuration panel for a 'Script' step. The 'Name' dropdown is now set to 'Unlock Lathe'. The 'Execution' dropdown is open, showing 'Select execution type'. The 'Parameters' section has two options: 'Blocking' and 'Non Blocking'.

Scripts may also have parameters added by choosing the *Add parameter* option.

- * A *Write* step will request you to choose the service to which you wish to send the write request. The list of available services is given in a drop down selection.

The screenshot shows a configuration panel for a 'Write' step. The 'Step' dropdown is set to 'Write'. The 'Service' dropdown is open, showing a list of services: modc, Sine, modbus2, Coolant, HTTPIn, Lathe, MQTTTest, Open62451, S2OPCUA, and Spinnaker. The 'Values' section is empty. There are buttons for '+ Add new step', 'Delete', 'Cancel', and 'Save'.

Values are added to the write request by clicking on the *Add new value* option. This will present a pair of text boxes in which the key and value of the write request value can be typed.

The screenshot shows the 'Add new value' dialog. The 'Step' is 'Write' and the 'Service' is 'Coolant'. The 'Values' section shows a 'key' and 'value' input field. There are buttons for '+ Add new value' and '+ Add condition'.

Multiple values can be sent in a single write request, to add another value simply click on the *Add new value* option again.

- Any step type may have a condition added to it. If a step has a condition associated with it, then that condition must evaluate to true if the step is to be run executed. If it does not evaluate to true the step is skipped and the next step is executed. To add a condition to a step click on the *Add condition* option within the step's panel.

Step Write

Service Coolant

Values rate 25

+ Add new value

Condition key == value

A key and a value text box appears, type the key to test, this is usually a script parameter and the value to test. Script parameters are referenced using the \$ character to enclose the name of the script parameter.

Step Write

Service Coolant

Values rate 25

+ Add new value

Condition \$flow\$ == 0

A selection list is provided that allows the test that you wish to perform to be chosen.

The screenshot shows a configuration panel for a script step. The 'Step' field is set to 'Write'. The 'Service' field is set to 'Coolant'. The 'Values' section has a field 'rate' with the value '25'. Below it is a link '+ Add new value'. The 'Condition' section has a field '\$flow\$' followed by a dropdown menu showing '==' and a field '0'. Below the condition is a link '+ Add new condition'. At the bottom left is a 'Delete' button. At the bottom right are 'Cancel' and 'Save' buttons. A dropdown menu is open for the condition operator, showing options: '==', '!=', '<', '>', '<=', and '>='.

- To remove a step from a script click on the bin icon on the right of the step panel

The screenshot shows a configuration panel for a script step. The 'Step' field is set to 'Write'. The 'Service' field is set to 'modc'. The 'Values' section has a field 'value2' with the value '1'. Below it are links '+ Add new value' and '+ Add condition'. A red circle highlights the bin icon in the top right corner of the panel.

- To reorder the steps in a script it is a simple case of clicking on one of the panels that contains a step and dragging and dropping the step into the new position within the script in which it should run.

The screenshot shows the 'Scripts / mod1' configuration page. At the top, the 'Name' field is set to 'mod1'. Below it, the 'Steps' section contains three steps:

- Step 1: Write**
 - Service: modc
 - Values: value2, 1
 - + Add new value
 - + Add condition
- Step 2: Delay**
 - Duration: 3000
 - + Add condition
- Step 3: (Service: modc)**
 - Values: value2, 0
 - + Add new value
 - + Add condition

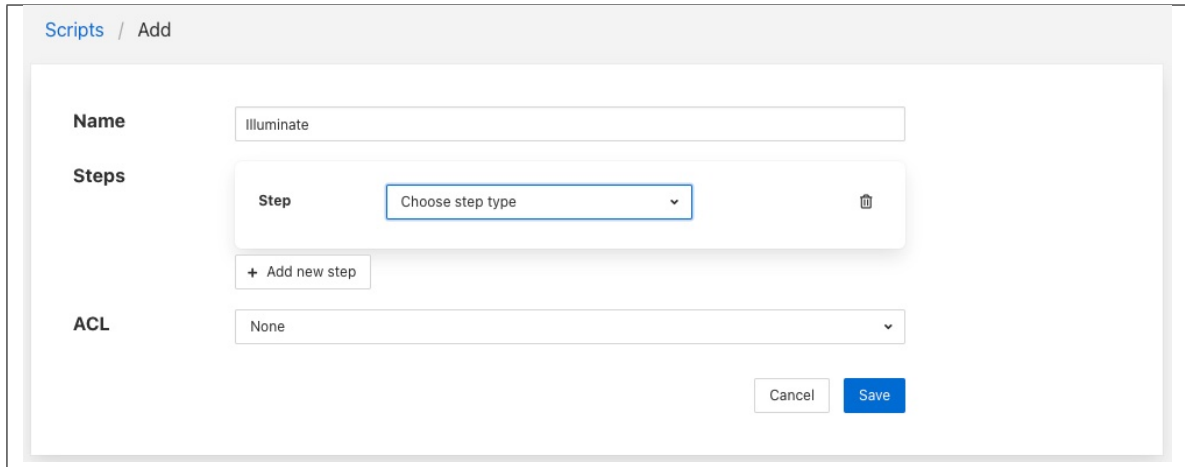
Below the steps, there is a '+ Add new step' button. At the bottom, the 'ACL' field is empty, and there are 'Delete', 'Cancel', and 'Save' buttons.

- A script may have an access control list associated to it. This controls how a script can be access, it allows the script to limit access to certain services, notifications or APIs. The creation of ACLs is covered elsewhere, to associate an ACL to a script simply select the name of the ACL from the ACL drop down at foot of the screen. If not ACL is assigned access to the script will not be limited.

Adding a Script

The process for adding new scripts is similar to editing an existing script.

- To add a new script click on the *Add* option in the top right corner.
- Enter a name for the script in the text box that appears



- Now start to add the steps to your script in the same way as above when editing an existing script.
- Once you have added all your steps you may also add optional access control list
- Finally click on *Save* to save your script

Step Conditions

The conditions that can be applied to a step allow for the checking of the values in the original request sent to the dispatcher. For example attaching a condition of the form

```
speed != 0
```

to a step, would result in the step being executed if the value in the parameter called *speed* that was in the original request to the dispatcher, had a value other than 0.


Conditions may be defined using the equals and not equals operators or for numeric values also greater than and less than.

7.3.3 Access Control Lists

Control features within FogLAMP have the ability to add access control to every stage of the control process. Access control lists are used to limit which services have access to specific south services or scripts within the FogLAMP system.

Graphical Interface

A graphical interface is available to allow the creation and management of the access control lists used within the FogLAMP system. This is available within the *Control Dispatcher* menu item of the FogLAMP graphical interface.



Adding An ACL

Click on the *Add* button in the top right corner of the ACL screen, the following screen will then be displayed.

You can enter a name for your access control list in the *Name* item on the screen. You should give each of your ACLs a unique name, this may be anything you like, but should ideally be descriptive or memorable as this is the name you will use when associating the ACL with services and scripts.

The *Services* section is use to define a set of services that this ACL is allowing access for. You may select services either by name or by service type. Multiple services may be granted access by a single ACL.

The screenshot shows the 'Services' section of the FogLAMP interface. On the left, there are labels for 'Names', 'Types', 'URLs', and 'ACLs'. A dropdown menu is open, displaying a list of service names: 'Core', 'Fledge Core', 'Southbound', 'Sine', 'Press001', and 'Press002'. The 'Sine' option is currently selected and highlighted in blue. To the right of the dropdown is a trash icon. Below the dropdown is a button labeled '+ Add new URL'.

To add a named service to the ACL select the names drop down list and select the service name from those displayed. The display will change to show the service that you added to the ACL.

The screenshot shows the 'Services' section with the 'Names' field now populated with 'HTTP'. The 'Types' field is still empty, showing the placeholder text 'Select service type'. The trash icon remains to the right of the 'Names' field.

More names may be added to the ACL by selecting the drop down again.

The screenshot shows the 'Services' section with the 'Names' field now containing two entries: 'HTTP' and 'Press001'. The 'Types' field is still empty. The dropdown menu is open again, showing a list of service names: 'Sine', 'Press001', 'Press002', 'Press003', 'Northbound', and 'HTTP'. The 'Press001' option is currently selected and highlighted in blue. The trash icon remains to the right of the 'Names' field.

If you wish to remove a named service from the list of services simply click on the small *x* to the left of the service name you wish to remove.

It is also possible to add a service type to an ACL. In this case all services of this type in the local FogLAMP instance will be given access via this ACL.

The screenshot shows the 'Services' configuration page. On the left, there are sections for 'Names', 'Types', 'URLs', and 'ACLs'. The 'Names' field has two entries: 'HTTP' and 'Press001'. The 'Types' dropdown menu is open, displaying a list of service types: 'Core', 'Storage', 'Southbound', 'Northbound', 'Notifications' (which is highlighted in blue), and 'Management'. Below the 'Types' dropdown is a '+ Add new URL' button. The 'URLs' section on the left has a 'URL' field and an 'ACLs' field, with a trash icon to the right.


For example to create an ACL that allows all north services to have be granted access you would select *Northbound* in the *Services Types* drop down list.

The screenshot shows the 'ACLs / Add' configuration page. At the top, the 'Name' field is set to 'All North'. Below it, the 'Services' section contains 'Names' and 'Types' dropdowns. The 'Types' dropdown is open, showing a list of service types: 'Core', 'Storage', 'Southbound', 'Northbound' (highlighted in blue), 'Notifications', and 'Management'. Below the 'Types' dropdown is a '+ Add new URL' button. The 'URLs' section on the left has a 'URL' field and an 'ACLs' field, with a trash icon to the right. At the bottom of the form are 'Cancel' and 'Save' buttons.

The *URLs* section of the ACL is used to grant access to specific URLs accessing the system.

Note: This is intended to allow control access via the REST API of the FogLAMP instance and is currently not implemented in FogLAMP.

Once you are satisfied with the content of your access control list click on the *Save* button at the bottom of the page. You will be taken back to a display of the list of ACLs defined in your system.



Updating an ACL

In the page that displays the set of ACLs in your system, click on the name of the ACL you wish to update, a page will then be displayed showing the current contents of the ACL.

To completely remove the ACL from the system click on the *Delete* button at the bottom of the page.

You may add and remove service names and types using the same procedure you used when adding the ACL.

Once you are happy with your updated ACL click on the *Save* button.

7.3.4 Configuration

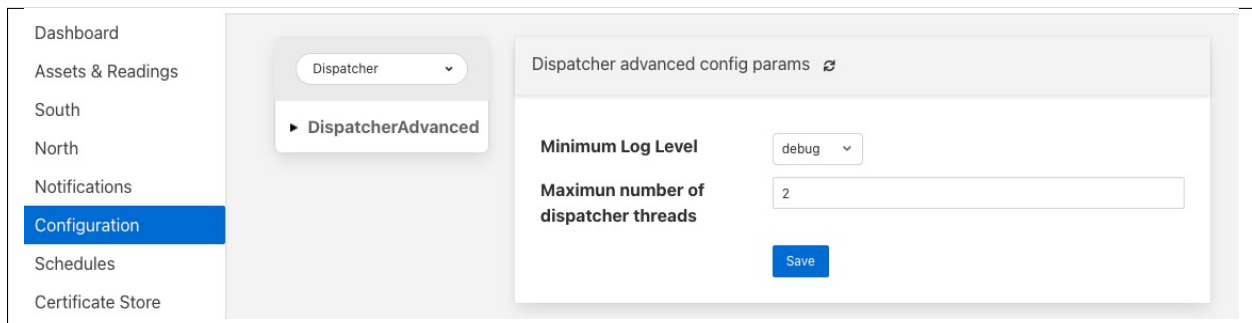
The *control dispatcher service* has a small number of configuration items that are available in the *Dispatcher* configuration category within the general Configuration menu item on the user interface.

Two subcategories exist, Server and Advanced.

Server Configuration

The server section contains a single option which can be used to either turn on or off the forwarding of control messages to the various services within FogLAMP. Clicking this option off will turn off all control message routing within FogLAMP.

Advanced Configuration



The screenshot displays the FogLAMP user interface for configuration. On the left is a sidebar menu with options: Dashboard, Assets & Readings, South, North, Notifications, Configuration (highlighted in blue), Schedules, and Certificate Store. The main content area shows a 'Dispatcher' dropdown menu with a 'DispatcherAdvanced' subcategory expanded. To the right, a 'Dispatcher advanced config params' panel contains two settings: 'Minimum Log Level' set to 'debug' and 'Maximum number of dispatcher threads' set to '2'. A blue 'Save' button is located at the bottom right of the configuration panel.

- **Minimum Log Level:** Allows the minimum level at which logs will get written to the system log to be defined.
- **Maximum number of dispatcher threads:** Dispatcher threads are used to execute automation scripts. Each script utilizes a single thread for the duration of the execution of the script. Therefore this setting determines how many scripts can be executed in parallel.

PLUGIN DOCUMENTATION

The following external plugins are currently available to extend the functionality of FogLAMP.

8.1 FogLAMP South Plugins

8.1.1 ABB Ability Smart Cloud Service

The *foglamp-south-abb* plugin is designed to pull data from the ABB Ability™ Smart Sensor Cloud into FogLAMP. It pulls data for a list of ABB assets into the local FogLAMP system at a rate defined for the service.

To create a south service with the ABB plugin

- Click on *South* in the left hand menu bar
- Select *ABB* from the plugin list
- Name your service and click *Next*

1 Plugin & Service Name 2 Review Configuration 3 Done

ABB Assets

```

1 {
2   "assets": [
3     ""
4   ]
5 }

```

ABB Service api.smartsensor.abb.com

Username FogLAMP

Auth. Key ---

Asset Structure Single Asset

Previous Next

- Configure the plugin
 - **ABB Assets:** A list of the assets in the ABB cloud service that should be read. This is a JSON document with an array called assets which contains the assets name as strings.
 - **ABB Service:** The hostname of the ABB service to which to connect. Usually this is the default api.smartsensor.abb.com.
 - **Username:** The ABB cloud user name.
 - **Auth. Key:** The authentication key that has been created in the ABB cloud for the given username.
 - **Asset Structure:** This defines how the FogLAMP assets that will be created should be organized.

Asset Structure

✓ Single Asset

Group Assets

Individual Asset

- * **Single Asset:** A single asset in the ABB cloud will be stored as a single asset in FogLAMP with the same name as the ABB asset. Within each FogLAMP asset a data point will be

created for each data value within the asset using the ABB measurement type name.

- * **Group Assets:** An asset will be created for each group of sensors for each asset within the ABB cloud. The asset will be named `<ABB asset>_<group name>`. Within each FogLAMP asset a data point will be created for each data value within the group using the ABB measurement type name.
- * **Individual Assets:** An asset will be created for each data item for each ABB cloud asset. The asset will be named `<ABB asset>_<item name>`.

8.1.2 AM2315 Temperature & Humidity Sensor



The `foglamp-south-am2315` is a south plugin for a temperature and humidity sensor. The sensor connects via the I2C bus and can provide temperature data in the range -40°C to $+125^{\circ}\text{C}$ with an accuracy of 0.1°C .

The plugin will produce a single asset that has two data points; temperature and humidity.

Note: The AM2315 is only available on the Raspberry Pi as it requires an I2C bus connection

To create a south service with the AM2315 plugin

- Click on *South* in the left hand menu bar
- Select `am2315` from the plugin list
- Name your service and click *Next*

- Configure the plugin

- **Asset Name:** The name of the asset that will be created. To help when multiple AM2315 sensors are used a %M may be added to the asset name. This will be replaced with the I2C address of the sensor.
- **I2C Address:** The I2C address of the sensor, this allows multiple sensors to be added to the same I2C bus.
- Click *Next*
- Enable the service and click on *Done*

Wiring The Sensor

The following table details the four connections that must be made from the sensor to the Raspberry Pi GPIO connector.

Colour	Name	GPIO Pin	Description
Red	VDD	Pin 2 (5V)	Power (3.3V - 5V)
Yellow	SDA	Pin 3 (SDA)	Serial Data
Black	GND	Pin 6 (GND)	Ground
White	SCL	Pin 5 (SCL)	Serial Clock

8.1.3 Beckhoff TwinCAT

The *foglamp-south-beckhoff* plugin is a plugin that allows collection of data from Beckhoff PLC's using the TwinCAT 2 or TwinCAT 3 protocols. It utilises the ADS library to allow updates the values held within the PLC to be captured in FogLAMP and sent onward as with any other data in FogLAMP.

The plugin uses a subscription model to register for changes to variables within the PLC and each of these becomes a data point in the asset that is created within FogLAMP.

To create a south service with the Beckhoff TwinCAT plugin

- Click on *South* in the left hand menu bar
- Select *Beckhoff* from the plugin list
- Name your service and click *Next*

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name:

EtherCAT Server:

Remote NetId:

Protocol:

Source NetId:

TwinCAT Map:

```

1 {
2   "items": [
3     {
4       "datapoint": "engine",
5       "name": "MAIN.engine"
6     }
7   ]
8 }

```

- Configure the plugin
 - **Asset Name:** The default asset name that is used for the data that is extracted from the PLC if the map does not define an explicit asset name.
 - **EtherCAT Server:** The hostname or IP address of the ADS master, this is the IP address of the Beckhoff PLC.
 - **Remote NetId:** The Beckhoff netId of the PLC. This is normally the IP address of the PLC with .1.1 appended to it.
 - **Protocol:** Define if the Automatic, TwinCAT 2 or TwinCAT 3 protocol is to be used. If *Automatic* is chosen the plugin will attempt to determine if the PLC supports TwinCAT 2 or TwinCAT 3.
 - **Source NetId:** The Beckhoff AMS NetId to assign to this plugin. This may be left blank, in which case an NetId will be generated from the IP address of the machine. However in some circumstances this is not acceptable or does not work correctly. A source NetId must always be provided when running within a container.
 - **TwinCAT Map:** A JSON document that is the data mapping for the PLC. This defines what variables are to be extracted from the PLC. See below for details of the map format.
- You must also authorise the FogLAMP plugin by adding an AMS route on your PLC

Adding AMS Route

Sample AMS route:

Name:	MyAdsClient
AMS Net Id:	192.168.0.1.1.1 # Derived from the IP address of your FogLAMP
Address:	192.168.0.1 # The IP address of your FogLAMP
Transport Type:	TCP/IP

Routes can be configured using one of several different methods;

TwinCAT Engineering: Go to the tree item SYSTEM/Routes and add a static route.

TwinCAT Systray: Open the context menu by right click the TwinCAT systray icon. (not available on Windows CE devices)

TC2: Go to Properties/AMS Router/Remote Computers. This requires a restart of TwinCAT on your PLC

TC3: Go to Router/Edit routes.

TcAmsRemoteMgr: Windows CE devices can be configured locally (TC2/TC3). Tool location: /Hard Disk/System/TcAmsRemoteMgr.exe. If uses TwinCAT 2 then a restart will be required after adding the AMS Route.

IPC Diagnose: Beckhoff IPC's provide a web interface for diagnose and configuration. Further information: [Beckhoff Device Manager](<http://infosys.beckhoff.de/content/1033/devicemanager/index.html?id=286>)

Map Format

The map is a JSON document that describes the variables to be extracted from the PLC. The variables may be defined either by name or by group and index id. Each variable will become a datapoint with the asset added to FogLAMP. The map itself is a single JSON array called "items", with each element in the array being an object that define the variable and what to do with it.

These objects have the following members within them

Key	Description
as-set	An optional element that defines an asset code that should be used to store the variable. If this is not given then the default asset code for the plugin is used.
dat-a-point	The name of the datapoint into which the variable is stored within the asset code. The datapoint name must be given for each object in the map.
name	The variable name within the PLC that is extracted. This may be obtained either by examining the PLC code that is running or by extracted from the .TPY file for the PLC. Either name or group and index must be given for each item in the map.
group	The numeric group within the PLC from which data is extracted. This allow data to be extracted without the use of variable names. It is not recommended for production use as it is very dependent on the layout of the PLC code, using variable names is more robust than group and index.
in-dex	The numeric index within the group from which to extract data, see above.

Example

An example TwinCAT map is shown below

```
{
  "items": [
    {
      "datapoint": "engine",
      "name": "MAIN.engine"
    }
  ]
}
```

This is based on the simulation that is available from Beckhoff and creates a single data point within the default asset called engine. It is populated with the value of the internal PLC variable *MAIN.engine*. A new asset will be created and added to the FogLAMP buffer every time this variable changes.

Multiple items may be read from the PLC by adding an element for each to the *items* array. For example to extract the two variables *MAIN.oilPressure* and *Main.engineSpeed* from the PLC a map as shown below could be used.

```
{
  "items": [
    {
      "datapoint": "oilPressure",
      "name": "MAIN.oilPressure"
    },
    {
      "datapoint": "rpm",
      "name": "MAIN.engineSpeed"
    }
  ]
}
```

Testing

The easiest way to test the Beckhoff plugin is to setup a simulation on a windows machine and run the Beckhoff PLC in simulator mode. The Beckhoff PLC can be freely downloaded from the Beckhoff site.

```
https://beckhoff.co.uk/english/download/tc3-downloads.htm?id=1905053019883865
```

This is designed to be run on a Windows 7 machine.

You can then create some sample variables to try to link to.

Downloading the code from Beckhoff includes a simple example that can be run that defines an engine variable, this is the example for which the default configuration is setup for.

Note: You will need to setup a static route in the Beckhoff PLC with the AMSNetId and IP address for the plugin and the type as TCP/IP.

8.1.4 CC2650 SensorTag



The *foglamp-south-cc2650* is a plugin that connects using Bluetooth to a Texas Instruments . The SensorTag offers 10 sensors within a small, low powered package which may be read by this plugin and ingested into FogLAMP. These sensors include;

- ambient light
- magnetometer
- humidity
- pressure
- accelerometer
- gyroscope
- object temperature
- digital microphone

Note: The sensor requires that you have a Bluetooth low energy adapter available that supports at least BLE 4.0.

To create a south service with the

- Click on *South* in the left hand menu bar
- Select *cc2650* from the plugin list
- Name your service and click *Next*

1 Plugin & Service Name 2 Review Configuration 3 Done

Bluetooth Address	B0:91:22:EA:79:04
Asset Name Prefix	CC2650/%M/
Shutdown Threshold	10
Connection Timeout	3
Temperature Sensor	<input checked="" type="checkbox"/>
Temperature Sensor Name	temperature
Luminance Sensor	<input type="checkbox"/>
Luminance Sensor Name	luminance
Humidity Sensor	<input type="checkbox"/>
Humidity Sensor Name	humidity
Pressure Sensor	<input type="checkbox"/>
Pressure Sensor Name	pressure
Movement Sensor	<input type="checkbox"/>
Gyroscope Sensor Name	gyroscope
Accelerometer Sensor Name	accelerometer
Magnetometer Sensor Name	magnetometer
Battery Data	<input type="checkbox"/>
Battery Sensor Name	battery

Previous Next

- Configure the plugin
 - **Bluetooth Address:** The Bluetooth MAC address of the device
 - **Asset Name Prefix:** A prefix to add to the asset name
 - **Shutdown Threshold:** The time in seconds allowed for a shutdown operation to complete
 - **Connection Timeout:** The Bluetooth connection timeout to use when attempting to connect to the device
 - **Temperature Sensor:** A toggle to include the temperature data in the data ingested
 - **Temperature Sensor Name:** The data point name to assign the temperature data
 - **Luminance Sensor:** Toggle to control the inclusion of the ambient light data
 - **Luminance Sensor Name:** The data point name to use for the luminance data
 - **Humidity Sensor:** A toggle to include the humidity data
 - **Humidity Sensor Name:** The data point name to use for the humidity data
 - **Pressure Sensor:** A toggle to control the inclusion of pressure data
 - **Pressure Sensor Name:** The name to be used for the data point that will contain the atmospheric pressure data

- **Movement Sensor:** A toggle that controls the inclusion of movement data gathered from the gyroscope, accelerometer and magnetometer
- **Gyroscope Sensor Name:** The data point name to use for the gyroscope data
- **Accelerometer Sensor Name:** The name of the data point that will record the accelerometer data
- **Magnetometer Sensor Name:** The name to use for the magnetometer data
- **Battery Data:** A toggle to control inclusion of the state of charge of the battery
- **Battery Sensor Name:** The data point name for the battery charge percentage
- Click *Next*
- Enable the service and click on *Done*

8.1.5 CoAP

The *foglamp-south-coap* plugin implements a passive listener that will accept data from sensors implementing the CoAP protocol. CoAP is an Internet application protocol for constrained devices to send data over the internet, it is similar to HTTP but may be run over UDP or TCP and is considerably simplified to allow implementation in small footprint devices. CoAP stands for Constrained Application Protocol.

The plugin listens for POST requests to the URI defined in the configuration. It expects the content of this PUT request to be a CBOR payload which it will expand and create assets for the items read from the CBOR payload.

To create a south service with the CoAP plugin

- Click on *South* in the left hand menu bar
- Select *coap* from the plugin list
- Name your service and click *Next*

The screenshot displays a configuration window for the CoAP plugin. At the top, a progress bar indicates three steps: 1. Plugin & Service Name, 2. Review Configuration (currently active), and 3. Done. Below the progress bar, there are two input fields: 'Port' with the value '5683' and 'URI' with the value 'sensor-values'. At the bottom, there are 'Previous' and 'Next' buttons.

- Configure the plugin
 - **Port:** The port on which the CoAP plugin will listen
 - **URI:** The URI the plugin expects to receive POST requests
- Click *Next*
- Enable the service and click on *Done*

8.1.6 Simple CSV Plugin

The *foglamp-south-csv* plugin is a simple plugin for reading comma separated variable files and injecting them as if there were sensor data. There are a number of variants of plugin that support this functionality with varying degrees of sophistication. These may also be considered as simple examples of how to write plugin code.

This particular CSV reader supports single or multi-column CSV files, without timestamps in the file. It assumes every value is a data value. If the multi-column option is not set then it will read data from the file up until a newline or a comma character and make that as single data point in an asset and return that.

If the multi-column option is selected then each column in the CSV file becomes a data point within a single asset. It is assumed that every row of the CSV file will have the same number of values.

Upon reaching the end of the file the plugin will restart sending data from the beginning of the file.

To create a south service with the csv plugin

- Click on *South* in the left hand menu bar
- Select *Csv* from the plugin list
- Name your service and click *Next*

- Configure the plugin
 - **Asset Name:** The name of the asset that will be created
 - **Datapoint:** The name of the data point to insert. If multi-column is selected this becomes the prefix of the name, with the column number appended to create the full name
 - **Multi-Column:** If selected then each row of the CSV file is treated as a single asset with each column becoming a data point within that asset.
 - **Path Of File:** The file that should be read by the CSV plugin, this may be any location within the host operating system. The FogLAMP process should have permission to read this file.
- Click *Next*
- Enable the service and click on *Done*

8.1.7 CSV Playback

The plugin plays a csv file inside some given directory in file system (The default being FOGlamp_ROOT/data). It converts the columns of csv file into readings which are datapoints of an output asset. The plugin plays readings at some configured rate.

We can also convert the columns of csv file into some other data type. For example from float to integer. The converted data will be part of reading not the CSV file.

The plugin has the ability to play the readings in either burst or continuous mode. In burst mode all readings are ingested into database at once and there is no adjustment of timestamp of a single reading. Whereas in continuous mode readings are ingested one by one and the timestamp of each reading is adjusted according to sampling rate. (For example if sampling rate is 8000 then the user_ts of every reading differs by 125 micro seconds.)

We can also copy the timestamp if present in the CSV file. This time stamp becomes the user_ts of a reading.

The plugin can also play the file in a loop which means it can start again if end of the file has reached.

The plugin can also play a file that has variable columns in every line.

- **‘assetName’:** type: string default: **‘vibration’**: The output asset that contains the readings.
- **‘csvDirName’:** type: string default: **‘FOGLAMP_DATA’**: The directory where CSV file exists. Default is FOGlamp_DATA or FOGlamp_ROOT/data
- **‘csvFileName’:** type: string default: **‘’**: CSV file name or pattern to search inside directory. Not necessarily an exact file name. If there are multiple files matching with the pattern, then the plugin will pick the first file in alphabetical order. If postProcessMethod is rename or delete then it will rename or delete the played file and pick the next one and so on.
- **‘headerMethod’:** type: enumeration default: **‘do_not_skip’**: The method for processing the header of csv file.
 1. skip_rows : If this is selected then the plugin will skip a given number of rows. The number of rows should be given in noOfRows config parameter given below.

2. `pass_in_datapoint` : If this is selected then the given number of rows will be combined into a string. This string will be present inside some given datapoint. Useful in cases where we want to ingest meta data along with readings from the csv file.
 3. `do_not_skip`: This option will not take any action on the header.
- **‘datapointForCombine’**: **type: string default: ‘metadata’**: If header method is `pass_in_datapoint` then it is the datapoint name where the given number of rows will get combined.
 - **‘noOfRows’**: **type: integer default: ‘1’**: No. of rows to skip or combine to single value. Used when header-Method is either `skip_rows` or `pass_in_datapoint`.
 - **‘variableCols’**: **type: boolean default: ‘false’**: It should be set true, when the columns in every row of CSV are not fixed. For example If you have a file like this

```
a,b,c
```

```
2,3,,23
```

```
4
```

Then you should set it true.

Note: Only one reading will be ingested at a time in this case. If you want to increase the rate then increase `readingPerSec` parameter in advanced plugin configuration.

- **‘columnMethod’**: **type: enumeration default: ‘pick_from_file’**: If variable Columns is false then it indicates how columns are considered.
 1. `pick_from_file` : The columns will be picked using a row index given.
 2. `explicit` : Specify the columns inside `useColumns` parameter.
- **‘autoGeneratePrefix’**: **type: string default: ‘column’**: If variable Columns is set true then data points will be generated using the prefix. For example if there is row like this 1,,2 and we chose `autoGeneratePrefix` to be `column`, then we will get data points like this `column_1: 1, column_3: 2`. Empty values will be ignored.
- **‘useColumns’**: **type: string default: ‘’**: Format **column1:type,column2:type**

The data types supported are: int, float, str, datetime, bool

We can perform three tasks with this config parameter.

1. The column name will get renamed in the reading if different name is used other than present in CSV file.
2. We can select a subset of columns from total columns.
3. We can convert the data type of each column.

Example if the file is like the following

```
id,value,status
```

```
1,2.5,'OK'
```

```
2,2.7,'OK'
```

Then we can give

1. `id:int,temperature:float,status:str`

The column value will be renamed to temperature.

2. `id:int,value:float`

Only two columns will be selected here.

3. id:int,temperature:int,status:str

The data type will be converted to integer. Also column will be renamed.

Row index for column names	<input type="text" value="0"/>
Ingest mode	<input type="button" value="burst"/>
Sample rate	<input type="text" value="8000"/>
Burst interval (ms)	<input type="text" value="1000"/>
Timestamp processing mode	<input type="button" value="current time"/>
Timestamp column name	<input type="text" value=""/>
Timestamp format	<input type="text" value="%Y-%m-%d %H:%M:%S.%f%z"/>
Ignore or report for NaN	<input type="button" value="ignore"/>
Post process method	<input type="button" value="continue_playing"/>
Suffix name	<input type="text" value=".tmp"/>

- **‘rowIndexForColumnNames’**: type: integer default: **‘0’**: If column method is pick_from_file then it is the index from where column names are taken.
- **‘ingestMode’**: type: enumeration default: **‘burst’**: Burst or continuous mode for ingestion.
- **‘sampleRate’**: type: integer default: **‘8000’**: No of readings per second to ingest.
- **‘burstInterval’**: type: integer default: **‘1000’**: Used for burst mode. Time interval between consecutive bursts in milliseconds.
- **‘timestampStyle’**: type: enumeration default: **‘current time’**: Controls how to give timestamps to reading. Works in four ways:
 1. current time: The timestamp in the readings is whatever the local time in the machine.
 2. copy csv value: Copy the timestamp present in the CSV file.
 3. move csv value: Used when we do not want to include timestamps from files in actual readings.
 4. use csv sample delta: Pick the delta between two readings in the file and construct the timestamp of reading using this delta. Assuming the delta remains constant through out the file.)
- **‘timestampCol’**: type: string default: **‘’**: The timestamp column to pick from the file. Used only when timestampStyle is not ‘current time’.
- **‘timestampFormat’**: type: string default: **‘%Y-%m-%d %H:%M:%S.%f%z’**: The timestamp format that will be used to parse the time stamps present in the file. Used only when timestampStyle is not ‘current time’.
- **‘ignoreNaN’**: type: enumeration default: **ignore**: Pandas takes the white spaces and missing values as NaN’s. These NaN’s cause problem while ingesting into database. It is left to the user to ensure there

are no missing values in CSV file. However if the option selected is report. Then plugin will check for NaN's and report error to user. This can serve as a way to check the CSV file for missing values. However the user has to take action on what to do with NaN values. The default action is to ignore them. When error is reported the user must delete the south service and try again with clean CSV file.

- **'postProcessMethod': type: enumeration default: 'continue_playing'**: It is the method to process the CSV file once all rows are ingested. It could be:
 1. continue_playing
Play the file again if finished.
 2. delete
Delete the played file once finished.
 3. rename
Rename the file with suffix after playing.
- **'suffixName': type: string default: '.tmp'**: The suffix name for renaming the file if postProcess method is rename.

Execution

Assuming you have a csv file named vibration.csv inside FOGlamp_ROOT/data/csv_data (Can give a pattern like vib. The plugin will search for all the files starting with vib and therefore find out the file named vibration.csv). The csv file has fixed number of columns per row. Also assuming the column names are present in the first line. The plugin will rename the file with suffix .tmp after playing. Here is the cURL command for that.

```
res=$(curl -sX POST http://localhost:8081/foglamp/service -d @- << EOF | jq
↪ '.')
{
  "name": "csv_player",
  "type": "south",
  "plugin": "csvplayback",
  "enabled": false,
  "config": {
    "assetName": {"value": "My_csv_asset"},
    "csvDirName": {"value": "FOGLAMP_DATA/csv_data"},
    "csvFileName": {"value": "vib"},
    "headerMethod": {"value": "do_not_skip"},
    "variableCols": {"value": "false"},
    "columnMethod": {"value": "pick_from_file"},
    "rowIndexForColumnNames": {"value": "0"},
    "ingestMode": {"value": "burst"},
    "sampleRate": {"value": "8000"},
    "postProcessMethod": {"value": "rename"},
    "suffixName": {"value": ".tmp"}
  }
}
EOF
)

echo $res
```

Poll Vs Async

The plugin also works in async mode. Though the default mode is poll. The async mode is faster but suffers with memory growth when sample rate is too high for the machine configuration.

Use the following sed operation for async and start the plugin again. The second sed operation, in similar way, can be used if you want to revert back to poll mode. Restart for the plugin service is required.

```
plugin_path=$FOGLAMP_ROOT/python/foglamp/plugins/south/csvplayback/csvplayback.py
value='s/POLL_MODE=True/POLL_MODE=False/'
sudo sed -i $value $plugin_path

# for reverting back to poll the commands will be
plugin_path=$FOGLAMP_ROOT/python/foglamp/plugins/south/csvplayback/csvplayback.py
value='s/POLL_MODE=False/POLL_MODE=True/'
sudo sed -i $value $plugin_path
```

Behaviour under various modes

Table 1: Behaviour of CSV playback plugin


Plugin mode	Ingest mode	Behaviour
poll	burst	No memory growth. Resembles the way sensors give data in real life. However the timestamps of readings won't differ by a fixed delta.
poll	continuous	No memory growth. Readings differ by a constant delta. However it is slow in performance.
async	continuous	Similar to poll continuous but faster. However memory growth is observed over time.
async	burst	Similar to poll burst. Not used generally.

For using poll mode in continuous setting increase the readingPerSec category to the sample rate.

```
sampling_rate=8000
curl -sX PUT http://localhost:8081/foglamp/category/csv_playerAdvanced -d '{
  ↪ "bufferThreshold": "'$sampling_rate'", "readingsPerSec": "'$sampling_rate'" }' | jq
```

It is advisable to increase the buffer threshold to atleast half the sample rate for good performance. (As done in above command)

8.1.8 DHT11 (C version)



The *foglamp-south-dht* plugin implements a temperature and humidity sensor using the DHT11 sensor module. Two versions of plugins for the DHT11 are available and are used as the example for . The other DHT11 plugin is *foglamp-south-dht11* and is a .

The DHT11 and the associated DHT22 sensors may be used, however they have slightly different characteristics;

	DHT11	DHT22
Voltage	3 to 5 Volts	3 to 5 Volts
Current	2.5mA	2.5mA
Humidity Range	0-50 % humidity 5% accuracy	0-100% humidity 2.5% accuracy
Temperature Range	0-50 +/- 2 degrees C	-40 to 80 +/- 0.5 degrees C
Sampling Frequency	1Hz	0.5Hz


Note: Due to the requirement for attaching to GPIO pins this plugin is only available for the Raspberry Pi platform.

To create a south service with the DHT11 plugin

- Click on *South* in the left hand menu bar
- Select *dht11_V2* from the plugin list
- Name your service and click *Next*

- Configure the plugin
 - **Asset Name:** The asset name which will be used for all data read.
 - **Rpi Pin:** The GPIO pin on the Raspberry Pi to which the DHT11 serial pin is connected.
- Click *Next*
- Enable the service and click on *Done*

8.1.9 DHT11 (Python version)



The *foglamp-south-dht11* plugin implements a temperature and humidity sensor using the DHT11 sensor module. Two versions of plugins for the DHT11 are available and are used as the example for . The other DHT11 plugin is *foglamp-south-dht* and is a .

The DHT11 and the associated DHT22 sensors may be used, however they have slightly different characteristics;

	DHT11	DHT22
Voltage	3 to 5 Volts	3 to 5 Volts
Current	2.5mA	2.5mA
Humidity Range	0-50 % humidity 5% accuracy	0-100% humidity 2.5% accuracy
Temperature Range	0-50 +/- 2 degrees C	-40 to 80 +/- 0.5 degrees C
Sampling Frequency	1Hz	0.5Hz

Note: Due to the requirement for attaching to GPIO pins this plugin is only available for the Raspberry Pi platform.

To create a south service with the DHT11 plugin

- Click on *South* in the left hand menu bar
- Select *dht11* from the plugin list
- Name your service and click *Next*

- Configure the plugin
 - **Asset Name:** The asset name which will be used for all data read.
 - **GPIO Pin:** The GPIO pin on the Raspberry Pi to which the DHT11 serial pin is connected.
- Click *Next*
- Enable the service and click on *Done*

8.1.10 Digiducer Vibration Sensor



The *foglamp-south-digiducer* plugin allows a 333D01 USB Digital Accelerometer to be attached to FogLAMP for the collection of vibration data. The Digiducer is a piezoelectric accelerometer housed in a rugged enclosure complete with a data conditioning and acquisition interface that only requires a USB port on the FogLAMP device for connectivity.

The plugin allows for two modes of operation; continuous reading of the vibration data or sampled reading of the vibration data. In sampled mode the user configures a sample period and interval. The plugin will then read data for the sample period and forward it to the FogLAMP storage service. It will then pause collection for the sample interval before again collecting data. This repeats indefinitely.

To create a south service with the Digiducer

- Click on *South* in the left hand menu bar
- Select *digiducer* from the plugin list
- Name your service and click *Next*

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name: vibration

Sample Rate: 8000Hz

Block size: 256

Continuous Sampling: ☒

Sample Period: 10

Sample Interval: 30

Channel: 20 g pk

Previous Next

- Configure the plugin
 - **Asset Name:** The name of the asset that will be created in FogLAMP.
 - **Sample Rate:** The rate at which data will be sampled. A number of frequencies are supported in the range 8KHz to 48KHz.

Sample Rate

✓ 8000Hz

11025Hz

16000Hz

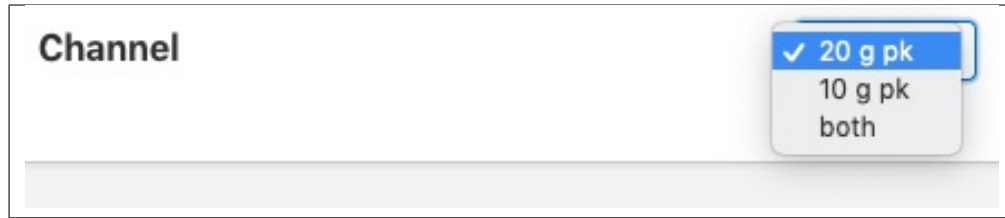
22050Hz

32000Hz

44100Hz

48000Hz

- **Block size:** To aid efficiency the plugin collects data in blocks, this allows the block size to be tuned. The value should be a power of 2.
- **Continuous Sampling:** This toggle supports the selection of continuous verses sampled collection.
- **Sample Period:** The duration of each sample period in seconds.
- **Sample Interval:** The time in seconds between each sample being taken.
- **Channel:** Select collection of the 10G Peak channel, the 20G Peak channel or both channels

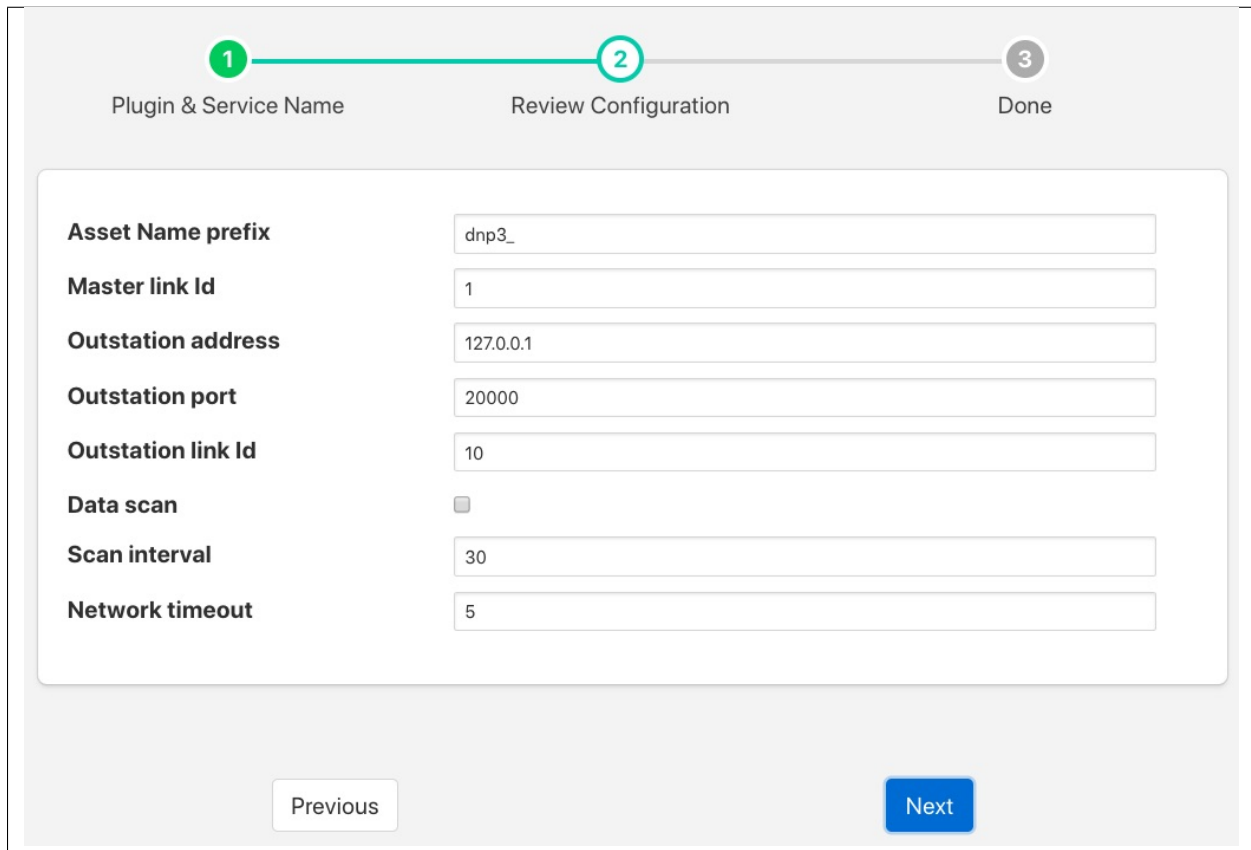


A screenshot of a web form with a label 'Channel' and a dropdown menu. The dropdown menu is open, showing three options: '20 g pk' (selected with a checkmark), '10 g pk', and 'both'.

- Click on *Next*
- Enable your south service and click on *Done*

8.1.11 DNP3 Master Plugin

The *foglamp-south-dnp3* allows FogLAMP to act as a DNP3 master and gather data from a DNP3 Out Station. The plugin will fetch all data types from the DNP3 Out Station and create assets for each in FogLAMP. The DNP3 plugin also handles unsolicited messages transmitted by the outstation.



A screenshot of the 'Review Configuration' step in the DNP3 Master Plugin setup. The form is titled 'Review Configuration' and is part of a three-step process: 1. Plugin & Service Name, 2. Review Configuration, and 3. Done. The form contains the following fields:

Field	Value
Asset Name prefix	dnp3_
Master link Id	1
Outstation address	127.0.0.1
Outstation port	20000
Outstation link Id	10
Data scan	<input type="checkbox"/>
Scan interval	30
Network timeout	5

At the bottom of the form, there are two buttons: 'Previous' and 'Next'.

- **Asset Name prefix:** An asset name prefix that is prepended to the DNP3 objects retrieved from the DNP3 outstations to create the FogLAMP asset name.
- **Master link id:** The master link id FogLAMP uses when implementing the DNP3 protocol.
- **Outstation address:** The IP address of the DNP3 Out Station to be connected.
- **Outstation port:** The port on the Out Station to which the connection is established.

- **Outstation link Id:** The Out Station link id.
- **Data scan:** Enable or disable the scanning of all objects and values in the Out Station. This is the Integrity Poll for all Classes.
- **Scan interval:** The interval between data scans of the Out Station.
- **Network timeout:** Timeout for fetching data from the Out Station expressed in seconds.

DNP3 Out Station Testing

The opendnp3 package contains a demo Out Station that can be used for test purposes. After building the opendnp3 package on your machine run the demo program as follows;

```
$ cd opendnp3/build
$ ./outstation-demo
```

This demo application listens on any IP address, port 20001 and has link Id set to 10. It also assumes master link Id is 1. Configuring your FogLAMP plugin with these parameters should allow FogLAMP to connect to this Out Station.

Once started it logs traffic and waits for use input to send unsolicited messages:


```
Enter one or more measurement changes then press <enter>
c = counter, b = binary, d = doublebit, a = analog, o = octet string, 'quit' = exit
```

Another option is the use of a DNP3 Out Station simulator, as an example:

<http://freyrscada.com/dnp3-ieee-1815-Client-Simulator.php#Download-DNP3-Development-Bundle>

Once the bundle has been downloaded, the **DNPOutstationSimulator.exe** application under the “Simulator” folder can be installed and run on a Windows 32bit platform.

8.1.12 Data Translation DT9837 Series



The *foglamp-south-dt9837* plugin is a south plugin that is designed to gather data from a Data Translation DT9837 Series DAQ.

To create a south service with the DT9837

- Click on *South* in the left hand menu bar
- Select *dt9837* from the plugin list
- Name your service and click *Next*

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name

Scan Rate

Input Mode

Range

First Channel

Last Channel

Sensitivity

IEPE Ch. 0 ☐

IEPE Ch. 1 ☐

IEPE Ch. 2 ☐

IEPE Ch. 3 ☐

Coupling Ch. 0

Coupling Ch. 1

Coupling Ch. 2

Coupling Ch. 3

- Configure the plugin
 - **Asset Name:** The name of the asset that will be created with the values read from the DT9837
 - **Scan Rate:** The rate at which each channel is read. This may be expressed as a numeric value, in which case it is the number of samples per second, or it may be expressed in KHz or MHz.
 - **Input Mode:** Defines how the input is treated, it may be either a differential pair or a single ended value with a reference ground.

Scan Rate

Input Mode

Range

- **Range:** This defines the voltage range for all channels. It may be defined as a bipolar value, in which case it is expected the signal can swing between + and - the specified voltage. A uni-polar

value, in which case the voltage swing is between ground and the specified voltage. Or it is a 0 to 20mA current loop.

The screenshot displays a configuration window with a list of settings on the left and a dropdown menu open on the right. The settings listed are:


- Range
- First Channel
- Last Channel
- Sensitivity
- IEPE Ch. 0
- IEPE Ch. 1
- IEPE Ch. 2
- IEPE Ch. 3
- Coupling Ch. 0
- Coupling Ch. 1
- Coupling Ch. 2
- Coupling Ch. 3

The dropdown menu for 'Range' is open, showing a list of options. The first option, 'BiPolar 60 Volts', is highlighted with a blue bar and a checkmark. The other options are:

- BiPolar 30 Volts
- BiPolar 20 Volts
- BiPolar 15 Volts
- BiPolar 10 Volts
- BiPolar 5 Volts
- BiPolar 4 Volts
- BiPolar 3 Volts
- BiPolar 2.5V Volts
- BiPolar 2 Volts
- BiPolar 1.25 Volts
- BiPolar 1 Volt
- BiPolar 625 mVolts
- BiPolar 500 mVolts
- BiPolar 250 mVolts
- BiPolar 312 mVolts
- BiPolar 200 mVolts
- BiPolar 156 mVlt
- BiPolar 125 mVolts
- BiPolar 100 mVolts
- BiPolar 78 mVolts
- BiPolar 50 mVolts
- BiPolar 10 mVolts
- BiPolar 5 mVolts
- 60 Volts
- 30 Volts
- 20 Volts
- 15 Volts
- 10 Volts
- 5 Volts
- 4 Volts
- 2.5 Volts
- 2 Volts
- 1.25 Volts
- 1 Volt
- 625 mVolts
- 500 mVolts
- 250 mVolts
- 200 mVolts
- 125 mVolts

A 'Previous' button is visible in the bottom right corner of the configuration area.

- **First Channel:** The DT9837 can scan a number of channel in a single operation, these must however be adjunct channels. This option sets the lowest numbered channel to be scanned.
- **Last Channel:** The DT9837 can scan a number of channel in a single operation, these must however be adjunct channels. This option sets the highest numbered channel to be scanned.
- **Sensitivity:** This sets the sensor sensitivity for IEPE sensors attached to any of the channels.
- **IEPE Ch. 0:** Specifies that a IEPE compatible sensor is attached to channel 0.
- **IEPE Ch. 1:** Specifies that a IEPE compatible sensor is attached to channel 1.
- **IEPE Ch. 2:** Specifies that a IEPE compatible sensor is attached to channel 2.
- **IEPE Ch. 3:** Specifies that a IEPE compatible sensor is attached to channel 3.
- **Coupling Ch. 0:** Specifies the input coupling to use for channel 0. This setting has no effect if the channel has been setup for IEPE as IEPE always uses AC coupling.



- **Coupling Ch. 1:** Specifies the input coupling to use for channel 1. This setting has no effect if the channel has been setup for IEPE as IEPE always uses AC coupling.
 - **Coupling Ch. 2:** Specifies the input coupling to use for channel 2. This setting has no effect if the channel has been setup for IEPE as IEPE always uses AC coupling.
 - **Coupling Ch. 3:** Specifies the input coupling to use for channel 3. This setting has no effect if the channel has been setup for IEPE as IEPE always uses AC coupling.
- Click on *Next*
 - Enable your service and click on *Done*

8.1.13 Edge ML Plugin

The plugin takes a video frame from a camera or stream , sends that to edgectl cluster running somewhere else. The Edge ML cluster returns a response in the form of json which contains information about detected objects, their bounding boxes and confidence score. This information is overlayed on the frame and saved onto disk in the form of images. The results are also streamed on a browser.

1 Plugin & Service Name
2 Review Configuration
 3 Done

?

Source of Data Generation	<div style="border: 1px solid #ccc; padding: 2px; display: inline-block;">stream ▼</div>
Directory Name	<div style="border: 1px solid #ccc; padding: 2px;">Directory For Using Images</div>
Camera ID	<div style="border: 1px solid #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> 0 ⬆ ⬇ ⬆ </div>
RTSP URL	<div style="border: 1px solid #ccc; padding: 2px;">rtsp://localhost:8554/clip</div>
Frames Per Minute	<div style="border: 1px solid #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> 1000 ⬆ ⬇ ⬆ </div>
Inference Choice	<div style="border: 1px solid #ccc; padding: 2px; display: inline-block;">k8 ▼</div>
Deployment Name	<div style="border: 1px solid #ccc; padding: 2px;">mledge-deployment</div>
The URL to post the images.	<div style="border: 1px solid #ccc; padding: 2px;">http://localhost:30163/v1/vision/detection</div>

- **‘source’:** type: enumeration default: **‘stream’**: Source of data being generated. Could be camera if camera is attached, stream if stream is to be used and directory if we have a directory of images.
- **‘sourceDirName’:** type: string default: **‘Directory For Using Images’**: If source is directory then the directory which contains images.
- **‘cameraId’:** type: integer default: **0**: If camera is to be used then enter the device id of camera. If you use 0 then the following command should be successful.
 v4l2-ctl -l -list-formats-ext -device /dev/video0 .
 In case you dont get output use camera id 1, 2 and so on.
- **‘rtspUrl’:** type: string default: **rtsp://localhost:8554/clip**: If source is stream, then enter the url of the rtsp stream.
- **‘fpm’:** type: integer default: **‘1000’**: No of frames to process in a minute.
- **‘inferenceChoice’:** type: enumeration default: **‘k8’**: If the edgeml cluster is running inside microk8’s on the same machine then use k8 or use URL if you want to send inference request to some other machine or some k8 cluster other than microk8’s.
- **‘deploymentName’:** type: string default: **‘mledge-deployment’**: If inferenceChoice is k8 then the name of the deployment inside microk8’s. The plugin will pick the ip and port from the deployment name itself.
- **‘restUrl’:** type: string default: **http://localhost:30163/v1/vision/detection**: If inferenceChoice is URL, then the URL where the post request will be sent.

Stream Results	<input checked="" type="checkbox"/>
Stream Port	<input type="text" value="8085"/>
Output Asset	<input type="text" value="Detection Results"/>
Destination Directory	<input type="text" value="detection"/>
Rotate after	<input type="text" value="120"/>
Rotation Data Amount	<input type="text" value="10"/>

Previous
Next

- **‘streamResults’**: type: boolean default: **‘true’**: Whether to stream detection results over HTTP(s)
- **‘streamPort’**: type: integer default: **‘8085’**: The port over which we can display detection results in browser.
- **‘outputAsset’**: type: string default: **‘Detected Results’**: The name of asset which contains detected results.
- **‘destinationDir’**: type: string default: **‘detection’**: The directory where resultant images will be stored.
- **‘rotateAfterMinutes’**: type: integer default: **‘120’**: The amount of time (in minutes) after which source images (with bounding boxes) are deleted”.
- **‘rotateDataMinutes’**: type: integer default: **‘10’**: The amount of jpeg files (with bounding boxes) in minutes to be rotated.

Installation

Part 1: Get the video feed

There are two ways to get the video feed.

1. Camera

To see the supported configuration of the camera run the following command.

```
$ v4l2-ctl --list-formats-ext --device /dev/video0
You will see something like
'YUYV' (YUYV 4:2:2)
  Size: Discrete 640x480
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 720x480
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 1280x720
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 1920x1080
    Interval: Discrete 0.067s (15.000 fps)
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 2592x1944
    Interval: Discrete 0.067s (15.000 fps)
  Size: Discrete 0x0
```

Above example uses Camera ID 0 to indicate use of /dev/video0 device, please use the applicable value for your setup

2. Network RTSP stream

To create a network stream follow the following steps

1. Install vlc

```
$ sudo add-apt-repository ppa:videolan/master-daily
$ sudo apt update
$ apt show vlc
$ sudo apt install vlc qtwayland5
$ sudo apt install libavcodec-extra
```

2. Download some sample files from here.

```
$ git clone https://github.com/intel-iot-devkit/sample-videos.git
```

3. Either stream a file using the following

```
$ vlc <name_of_file>.mp4 --sout '#gather:transcode{vcodec=h264,vb=512,
↳ scale=Auto,width=640,height=480,acodec=none,scodec=none}:rtp{sdp=rtsp://
↳ <ip_of_machine_streaming>:8554/clip}' --no-sout-all --sout-keep --loop --
↳ no-sout-audio --sout-x264-profile=baseline
```

Note : fill the <ip_of_the_machine> with ip of the machine which will be used to stream video. Also fill <name_of_file> with the name of mp4 file.

4. You can also stream from a camera using the following

```
$ vlc v4l2:///dev/video<index_of_video_device> --sout '#gather:transcode
↳ {vcodec=h264,vb=512,scale=Auto,width=<supported_width_of_camera_image>,
↳ height=<supported_height_of_camera_image>,acodec=none,scodec=none}:rtp
↳ {sdp=rtsp:///<ip_of_the_machine>:8554/clip}' --no-sout-all --sout-keep -
↳ no-sout-audio --sout-x264-profile=baseline
```

Fill the following :

- <index_of_video_device> The index with which you ran the v4l2 command mentioned above. for example video0.
- <supported_height_of_camera_image> Height you get when you ran v4l2 command mentioned above. For example Discrete 640x480. Here 480 is height.
- <supported_width_of_camera_image> Width you get when you ran v4l2 command mentioned above. For example Discrete 640x480. Here 640 is width.
- <ip_of_the_machine> ip of the machine which will be used to stream video.

Part 2: Start the Edge ML cluster

For starting the Edge ML cluster you should follow this [README](#) file.

Now run the plugin by filling parameters according to your setup.

8.1.14 EtherIP South Plugin

The *foglamp-south-etherip* plugin is a south plugin that supports PLC Tags read operation for Allen Bradley/Rockwell PLCs. This is based on an API level details can be read at

Configuration Parameters

An EtherIP south service is added in the same way as any other south service in FogLAMP,

- Select the *South* menu item
- Click on the + icon in the top right

You will be presented with the following page

The screenshot shows a configuration window titled 'Plugin & Service Name' (step 1 of 3). It features a 'South Plugin' dropdown menu with 'etherip' selected. Below the dropdown is a link for 'available plugins'. A 'Name' text input field contains the placeholder 'name'. At the bottom, there are 'Back' and 'Next' buttons. The top navigation bar shows steps 1, 2 ('Review Configuration'), and 3 ('Done').

- Select *etherip* from the plugin list
- Enter a name for your etherip service
- Click *Next*
- You will be presented with the following configuration page

Asset Name	PLCTags
PLC IP Address	127.0.0.1
PLC	micro800
Path	0
Timeout (ms)	5000
Tags to read	<pre> 1 { 2 "tags": [3 { 4 "name": "PLCTAG", 5 "type": "UINT32", 6 "program": "" 7 } 8] 9 }</pre>
Tags to write	<pre> 1 { 2 "tags": [3 { 4 "name": "PLCTAG", 5 "type": "UINT32", 6 "program": "" 7 } 8] 9 }</pre>

- **Asset Name:** This is the name of the asset that will be used for the data read by this service. Default Value PLCTags
- **PLC IP Address:** This is the IP Address of the PLC Tags server, Default value 127.0.0.1
- **PLC:** This is the type of the PLC from which the user wants to read the tags, Default value controllogix, the following have been supported:
 - Allen-Bradley ControlLogix & CompactLogix PLCs
 - Allen-Bradley Micro8x0 PLCs
 - Rockwell/Allen-Bradley PLCs accessed over a DH+ bridge (i.e. a LGX chassis with a DHRIO module) such as PLC/5, SLC 500 and MicroLogix
 - Omron NX/NJ series PLCs as for Allen-Bradley Micro8x0
- **Path:** This is the slot/rack(e.g. 0,1) to use to connect to your PLCTags device. This attribute is required for CompactLogix/ControlLogix tags and for tags using a DH+ protocol bridge (i.e. a DHRIO module) to get to a PLC/5, SLC 500, or MicroLogix PLC on a remote DH+ link. The attribute is ignored if it is not a DH+ bridge route, but will generate a warning if debugging is active. Note that Micro8x0 connections must not have a path attribute.
- **Timeout (ms):** The request timeout when communicating with a PLCTags server. Default value is 5000 (milliseconds).
- **Tags to read:** The map defines which PLCTags user wants to read. The map is a complex JSON object which is described in more detail below.
- **Tags to write:** The map defines which PLCTags user wants to write to. The map is a complex JSON object which is described in more detail below.

Map

The map of tags to read/write is a JSON object with a single array *tags*, each element of this array is a JSON object that defines a single item of data that will be stored in FogLAMP or written as set-point control. These objects support a number of properties and values, these are

Property	Description
name	The name of the PLCTag that we are reading/writing. In case of read, this becomes the name of the datapoint with the asset.
type	The datatype of the PLCTag that we are reading/writing.
program	This defines the scope of the PLCTag that we are reading/writing, possible values being Program level scope or Global/Controller level scope. The Program name needs to be specified here if it is a tag with program level scope. For global/controller level tags 'program' attribute in map can be omitted or set to JSON null value or empty.

Example JSON Map

In this example we will assume we have a Micro8x0 PLC and we want to read one tag with name Tag1, type UINT32 and having global/controller level scope. Also we want to write to a tag with name Tag2, type UDINT and having global/controller level scope.

The Tags to read Map attribute for this example would be as follow:

```
{
  "tags": [
    {
      "name": "Tag1",
      "type": "UINT32",
      "program": ""
    }
  ]
}
```

The Tags to write Map attribute for this example would be as follow:

```
{
  "tags": [
    {
      "name": "Tag2",
      "type": "UDINT",
      "program": ""
    }
  ]
}
```

8.1.15 Expression South Plugin

The *foglamp-south-expression* plugin is a plugin that is used to generate synthetic data using a mathematical expression to generate data that changes over time. The user may configure the plugin with an expression of their choice and define a period in terms of samples per period of the output and the increment between each sample.

The screenshot displays the configuration interface for the Expression South Plugin. At the top, a progress bar indicates three steps: 1. Plugin & Service Name, 2. Review Configuration (the current step, highlighted with a green circle), and 3. Done. Below the progress bar, a white box contains five input fields. The first field is 'Asset Name' with a placeholder 'Expression'. The second field is 'Expression' with the value 'clamp(-1.0, sin(2 * pi * x) + cos(x / 2 * pi), +1.0)'. The third field is 'Minimum Value' with the value '-5'. The fourth field is 'Maximum Value' with the value '5'. The fifth field is 'Step Value' with the value '0.001'. At the bottom of the interface, there are two buttons: 'Previous' and 'Next'.

The parameters that can be configured are;

- **Asset Name:** The name of the asset to be created inside FogLAMP.
- **Expression:** The expression that should be evaluated to create the asset value, see below.
- **Minimum Value:** The minimum value of x, where x is the value that sweeps over time.
- **Maximum Value:** The maximum value of x, where x is the value that sweeps over time.
- **Step Value:** The step in x for each call to the expression evaluation.

Expression Support

The *foglamp-south-expression* plugin makes use of the library to do run time expression evaluation. This library provides a rich mathematical operator set, the most useful of these in the context of this plugin are;


- Mathematical operators (+, -, *, /, %, ^)
- Functions (min, max, avg, sum, abs, ceil, floor, round, roundn, exp, log, log10, logn, pow, root, sqrt, clamp, inrange, swap)
- Trigonometry (sin, cos, tan, acos, asin, atan, atan2, cosh, cot, csc, sec, sinh, tanh, d2r, r2d, d2g, g2d, hyp)

8.1.16 Flir AX8 Thermal Imaging Camera



The *foglamp-south-FlirAX8* plugin is a south plugin that enables temperature data to be collected from Flir Thermal Imaging Devices, in particular the AX8 and other A Series cameras. The camera provides a number of temperatures for both spots and boxes defined within the field of view of the camera. In addition it can also provide deltas between two temperature readings.

The bounding boxes and spots to read are configured by connecting to the web interface of the camera and dropping the spots on a thermal imaging or pulling out rectangles for the bounding boxes. The camera will return a minimum, maximum and average temperature within each bounding box.



In order to configure a south service to obtain temperature data from a Flir camera select the *South* option from the left-hand menu bar and click on the Add icon in the top right corner of the South page that appears. Select the *FlirAX8* plugin from the list of south plugins, name your service and click on *Next*.

The screen that appears is the configuration screen for the *FlirAX8* plugin.

1

2

3

Plugin & Service Name

Review Configuration

Done

Asset Name

AX8

Server Address

192.168.0.48

Port

502

Slave ID

1

Previous

Next

There are four configuration parameters that can be set, usually it is only necessary to change the first two however;

- **Asset Name:** This is the asset name that the temperature data will be written to FogLAMP using. A single asset is used that will contain all of the values read from the camera.
- **Server Address:** This is the address of the Modbus server within the camera. This is the same IP address that is used to connect to the user interface of the camera.
- **Port:** The TCP port on which the cameras listens for Modbus requests. Unless changed in the camera the default port of 502 should be used.
- **Slave ID:** The Modbus Slave ID of the camera. By default the cameras are supplied with this set to 1, if changed within your camera setup you must also change the value here to match.

Once entered click on *Next*, enable the service on the next page and click on *Done*.

This will create a single asset that contains values for all boxes and spots that may be define. A filter *foglamp-filter-FlirValidity* can be added to the south service to remove data for boxes and spots not switched on in the camera user interface. See . This filter also allows you to name the boxes and hence have more meaningful names in the data points within the asset.

8.1.17 FLIR GW65 Vibration Sensors

The *foglamp-south-gw65* plugin provides a mechanism to connect the FLIR vibration sensors, via the GW65 gateway to FogLAMP. The plugin allows the GW65 to be used to connect sets of the SV87 vibration sensors. Raw vibration data is then collected from the sensors and may be process by one or more of the filters that offer vibration analysis.

The connection between the sensors and the plugin, via the GW65 gateway as established using the FLIR mobile application, before starting this process however you must install and configure the *foglamp-south-gw65* plugin. The plugin may be installed either by the user interface or by using the package manager of your Linux system to install it manually from a package.

You must also have an MQTT broker configured and running on your network. This should be configured to allow MQTTS and also have a username and password for the FLIR gateway to use.

Creating the GW65 South Service

Using the normal procedure for creating a new south service in FogLAMP,


- Select the *South* item from the menu on the left hand side of the screen
- Click on the *Add +* link on the top left of the South service screen
- In the list of available plugins choose the *gw65* entry, if it is not in the list click on the *available plugins* link and install it.
- Enter a name for your service and click on *Next*
- The configuration page will appear

The screenshot shows a configuration wizard with three steps: 1. Plugin & Service Name, 2. Review Configuration (current step), and 3. Done. A progress bar at the top indicates the current step. Below the progress bar, there is a form with two fields: 'Asset Name' with the value 'gw65' and 'MQTT Broker' with the value 'localhost'. A help icon (?) is visible next to the 'Asset Name' field. At the bottom of the form, there are two buttons: 'Previous' and 'Next'.

- **Asset Name:** The asset name to use as a fallback asset. This is normally unused.
- **MQTT Broker:** The address of the MQTT broker that will be used to communicate with the GW65 gateway.
- Enter the address of your MQTT broker, this **must** be the public address of the machine and not be localhost or the address 127.0.0.1 even when running your MQTT broker on the same machine as the FogLAMP. The information will be passed on to the GW65 gateway and must be an address to which the GW65 can route.
- Click on *Next*
- Enable the service and click on *Done*

You may now proceed to configure the GW65 using the FLIR mobile application. It is important that the south service within FogLAMP is running before this process is started.

- Open the FLIR application and select the *LOCAL SERVER* option.



- Click on the *Add Server* link

14:59 1

< Add Server >

To configure, please enter the Server IP, User Name, Password and Port.

Server IP

192.168.0.34

User name

flir

Password


123456

Port

8883

Continue

- **Server IP:** This is the IP address of your MQTT server. This should be the same as was entered into the south service that was created in FogLAMP.
- **User name:** The user name you configured in your MQTT server. The default value is *flir*.
- **Password:** The password assigned to the above user. The default is *12345678*.
- **Port:** Leave this as the default value 8883
- Click on *Continue*
- If the application responds with the error “Invalid IP Address” this could mean one of many things
 - The IP address you entered was incorrect
 - Your phone is not on the same network as FogLAMP
 - The username and password entered were incorrect
 - The south service for the GW65 is not running on the FogLAMP



- Click on the name of the discovered gateway and follow the steps to setup a connection to the WiFi network
- You will see the details of your GW65 gateway

14:18

<

Gateway Settings

Gateway Information

Gateway ID

GW657EA2

Name

GW657EA2

Server

unspecified

Sampling Rate

1 minute

WiFi Router

MILLEND2

Firmware Version

1.0.3

FW Update

Continue

- Click on *Continue* and follow the instructions to add your sensors

Installing an MQTT Broker

You may use any compatible MQTT broker with the plugin and FLIR GW65 gateway, during testing the Mosquitto MQTT broker was used, a package exists that allows this to be installed and configured for use with gateway, this package is called *foglamp-mqtt-broker*.

To use the package simply use your package manager to install the package, for example on a *apt* based system such as Ubuntu

```
apt install foglamp-mqtt-broker
```

This will

- Install the mosquitto MQTT service
- Configure it with a certificate
- Add a user with the username *foglamp* and the password *12345678*
- Start it listening on port 8883 for MQTTS

Alternatively you can manually configure the Mosquitto MQTT browser by using the following steps

- Edit the configuration file */etc/mosquitto/mosquitto.conf* and adding the following lines

```
# Start of MQTTS support
listener 8883
cafile /etc/mosquitto/certs/ca.crt
certfile /etc/mosquitto/certs/client.crt
keyfile /etc/mosquitto/certs/client.key

password_file /etc/mosquitto/passwordfile
# End of MQTTS support
```

- Create a password file by running the command

```
touch /etc/mosquitto/passwordfile
```

- Create the *foglamp* user, the password shown here is 12345678 but any password can be used but must be set in the GW65 configuration application.

```
mosquitto_passwd -b /etc/mosquitto/passwordfile foglamp 12345678
```

- Generate the required certificates

```
mkdir /etc/mosquitto/certs
cd /etc/mosquitto/certs
openssl req -new -x509 -days 365 -extensions v3_ca -keyout ca.key -out ca.crt -
↳subj "/C=RO/ST=Home/L=Home/O=Dianomic/OU=FogLAMP/CN=dianomic.com" -passout_
↳pass:foglamp
openssl genrsa -out client.key 2048
openssl req -new -out client.csr -key client.key -subj "/C=RO/ST=H/L=Home/O=MQTT_
↳Broker/OU=MQTT Broker/CN=mqtt-broker.local"
openssl x509 -req -in client.csr -CA ca.crt -CAkey ca.key -CAcreateserial -out_
↳client.crt -days 365 pass:foglamp
openssl rsa -in client.key -out client.key
```

- Set permissions for Mosquitto MQTT

```
chown -R mosquitto:/etc/mosquitto
chmod 700 /etc/mosquitto/certs
```

- Then restart your MQTT broker

```
systemctl restart mosquitto
```

8.1.18 South HTTP

The *foglamp-south-http* plugin allows data to be received from another FogLAMP instance or external system using a REST interface. The FogLAMP which is sending the data to the corresponding north task with the HTTP north plugin installed. There are two options for the HTTP north or , these serve the dual purpose of providing a data path between FogLAMP instances and also as examples of how other systems might use the REST interface from C/C++ or Python. The plugin supports both HTTP and HTTPS transport protocols and sends a JSON payload of reading data in the internal FogLAMP format.

The primary purpose of this plugin is for FogLAMP to FogLAMP communication however, there is no reason to prevent other applications that wish to send data into a FogLAMP system to not use this plugin also. The only requirement is that the application that is sending the data uses the same JSON payload structure as FogLAMP uses for passing reading data between different instances. Data should be sent to the URL defined in the configuration of the plugin using a POST request. The caller may choose to send one or many readings within a single POST request and those readings may be for multiple assets.

To create a south service you, as with any other south plugin

- Select *South* from the left hand menu bar.
- Click on the + icon in the top left
- Choose *http_south* from the plugin selection list
- Name your service
- Click on *Next*
- Configure the plugin

1 Plugin & Service Name 2 Review Configuration 3 Done

Host	0.0.0.0
Port	6683
URI	sensor-reading
Asset Name Prefix	http-
Enable HTTP	<input checked="" type="checkbox"/>
HTTPS Port	6684
Certificate Name	fledge

Previous Next

- **Host:** The host name or IP address to bind to. This may be left as default, in which case the plugin binds to any address. If you have a machine with multiple network interfaces you may use this parameter to select one of those interfaces to use.
 - **Port:** The port to listen for connection from another FogLAMP instance.
 - **URL:** URI that the plugin accepts data on. This should normally be left to the default.
 - **Asset Name Prefix:** A prefix to add to the incoming asset names. This may be left blank if you wish to preserve the same asset names.
 - **Enable HTTP:** This toggle specifies if HTTP connections should be accepted or not. If the toggle is off then only HTTPS connections can be used.
 - **Certificate Name:** The name of the certificate to use for the HTTPS encryption. This should be the name of a certificate that is stored in the FogLAMP .
- Click *Next*
 - Enable your service and click *Done*

JSON Payload

The payload that is expected by this plugin is a simple JSON presentation of a set of reading values. A JSON array is expected with one or more reading objects contained within it. Each reading object consists of a timestamp, an asset name and a set of data points within that asset. The data points are represented as name value pair JSON properties within the reading property.

The fixed part of every reading contains the following

Name	Description
times-tamp	The timestamp as an ASCII string in ISO 8601 extended format. If no time zone information is given it is assumed to indicate the use of UTC.
asset	The name of the asset this reading represents.
read-ings	A JSON object that contains the data points for this asset.

The content of the *readings* object is a set of JSON properties, each of which represents a data value. The type of these values may be integer, floating point, string, a JSON object or an array of floating point numbers.

A property

```
"voltage" : 239.4
```

would represent a numeric data value for the item *voltage* within the asset. Whereas

```
"voltageUnit" : "volts"
```

Is string data for that same asset. Other data may be presented as arrays

```
"acceleration" : [ 0.4, 0.8, 1.0 ]
```

would represent acceleration with the three components of the vector, x, y, and z. This may also be represented as an object


```
"acceleration" : { "X" : 0.4, "Y" : 0.8, "Z" : 1.0 }
```

both are valid formats within FogLAMP.

An example payload with a single reading would be as shown below

```
[
  {
    "timestamp" : "2020-07-08 16:16:07.263657+00:00",
    "asset"      : "motor1",
    "readings"   : {
      "voltage"  : 239.4,
      "current"   : 1003,
      "rpm"      : 120147
    }
  }
]
```

8.1.19 INA219 Voltage & Current Sensor



The *foglamp-south-ina219* plugin is a south plugin that uses an INA219 breakout board to measure current and voltage. The Texas Instruments INA219 is capable of measuring voltages up to 26 volts and currents up to 3.2 Amps. It connects via the I2C bus of the host and multiple sensors may be daisy chain on a single I2C bus. Breakout boards that mount the chip and its associate shunt resistor and connectors and easily available and attached to hosts with I2C buses.

The INA219 support three voltage/current ranges

- 32 Volts, 2 Amps
- 32 Volts, 1 Amp
- 16 Volts, 400 mAmps

Choosing the smallest range that is sufficient for your application will give you the best accuracy.

Note: This plugin is only available for the Raspberry Pi as it requires to be interfaced to the I2C bus on the Raspberry Pi GPIO header socket.

To create a south service with the INA219

- Click on *South* in the left hand menu bar
- Select *ina219* from the plugin list
- Name your service and click *Next*

- Configure the plugin
 - **Asset Name:** The asset name of the asst that will be written
 - **I2C Address:** The address of the INA219 device
 - **Voltage Range:** The voltage range that is to be used. This may be one of 32V2A, 32V1A or 16V400mA
- Click *Next*
- Enable the service and click on *Done*

Wiring The Sensor

The INA219 uses the I2C bus on the Raspberry PI, which requires two wires to connect the bus, it also requires power taking the total to four wires

INA219 Pin	Raspberry Pi Pin
Vin	3V3 pin 1
GND	GND pin 9
SDA	SDA pin 3
SCL	SCL pin 5

8.1.20 Lathe Simulation

The *foglamp-south-lathesim* plugin is a south plugin that simulates a lathe with a number of attached sensors. The purpose of this plugin is for test and demonstration only as it does not attach to any real device.

The plugin simulates four sensor devices attached to the virtual lathe

- The PLC controlling the lathe that gives details such as cutting depth, tool position, motor speed
- A current sensor that measures the current draw from the lathe
- A vibration sensor giving the RMS value of the vibration and the dominant vibration frequency
- A thermal imaging device that takes temperature readings every second from the motor, gearbox, headstock, tailstock and tool on the lathe

The vibration sensor reports at half the rate of the other sensors attached to the lathe in order to simulate handling data that is related to the same physical device but not available at the same rate as the other sensors.

The simulation runs a repeated pattern of operations;

- A spin-up period where the lathe spins up to speed from idle.
- A period where the lathe is doing some cutting of a work piece.
- A spin-down period where the lathe is slowing to a stop.
- An idle period where the work piece is removed and replaced with a new billet.

During the spin up period the lathe speed, expressed in revolutions per minute, will linearly increase from 0 to the maximum defined.

When the lathe is cutting the speed will remain predominantly constant, with a small random variation, whilst the depth of cut and X position of the cutting tool will change.

The lathe then spins down to rest and will remain idle for a short time whilst the worked item is removed and a new billet of material is installed.

During the cutting period the current draw and vibration will alter as load is applied to the piece.

Configuring the PLC

There are a number of configuration options that can be applied to the simulation.

Lathe1 South Service

Lathe Name: lathe

Spin up time: 5

Cutting time: 45

Idle time: 15

Spin down time: 6

RPM: 500

Current: 750

Enabled: ☒ [Show Advanced Config](#)

Applications +

Cancel Save

- **Lathe Name:** The name of the lathe in this configuration. This name is used to derive the assets returned from the three sets of sensors. The PLC data is returned with an asset name that machines the lathe name. The current data has *Current* appended to the lathe name and the asset id of the vibration name is the lathe name with *Vibration* appended to it. The temperature data uses the asset with the name of the lathe and *IR* appended to it.
- **Spin up time:** The time in seconds it takes the lathe to spin up to working speed from idle.
- **Cutting time:** The time in seconds for which the lathe is cutting material.
- **Spin Down time:** The time in seconds for which the lathe is spinning down from operating speed to stop.
- **Idle time:** The time in seconds for which the lathe is idle between jobs.

- **RPM:** The operating speed of the lathe, expressed in revolutions per minute.
- **Current:** The nominal operating current draw of the lathe.

8.1.21 Modbus South Plugin

The *foglamp-south-modbus-c* plugin is a south plugin that supports both TCP and RTU variants of Modbus. The plugin provides support for reading Modbus coils, input bits, registers and input registers, a flexible mechanism is provided to create a mapping between the Modbus registers and coils and the assets within FogLAMP. Multiple registers can be combined to allow larger values than register width to be mapped from devices that represent data in this way. Support is also included for floating point representation within the Modbus registers.

Configuration Parameters

A Modbus south service is added in the same way as any other south service in FogLAMP,

- Select the *South* menu item
- Click on the + icon in the top right
- Select *ModbusC* from the plugin list
- Enter a name for your Modbus service
- Click *Next*
- You will be presented with the following configuration page

Asset Name

modbus

Protocol

RTU

Server Address

127.0.0.1

Port

2222

Device

Baud Rate

9600

Number Of Data Bits

8

Number Of Stop Bits

1

Parity

none

Slave ID

1

Register Map

1

{

2

"values": [

3

{

4

"name": "temperature",

5

"slave": 1,

6

"assetName": "Booth1",

7

"register": 0,

8

"scale": 0.1,

9

"offset": 0

10

},

11

{

12

"name": "humidity",

}

}

Timeout

0.5

Control

None

Control Map

1

{

2

"values": []

3

}

- **Asset Name:** This is the name of the asset that will be used for the data read by this service. You can override this within the Modbus Map, so this should be treated as the default if no override is given.

- **Protocol:** This allows you to select either the *RTU* or *TCP* protocol. Modbus RTU is used whenever you have a serial connection, such as RS485 for connecting to your device. The TCP variant is used where you have a network connection to your device.
- **Server Address:** This is the network address of your Modbus device and is only valid if you selected the *TCP* protocol.
- **Port:** This is the port to use to connect to your Modbus device if you are using the TCP protocol.
- **Device:** This is the device to open if you are using the RTU protocol. This would be the name of a Linux device in `/dev`, for example `/dev/SERIAL0`
- **Baud Rate:** The baud rate used to communicate if you are using a serial connection with Modbus RTU.
- **Number of Data Bits:** The number of data bits to send on serial connections.
- **Number of Stop Bits:** The number of stop bits to send on the serial connections.
- **Parity:** The parity setting to use on the serial connection.
- **Slave ID:** The slave ID of the Modbus device from which you wish to pull data.
- **Register Map:** The register map defines which Modbus registers and coils you read, and how to map them to FogLAMP assets. The map is a complex JSON object which is described in more detail below.
- **Timeout:** The request timeout when communicating with a Modbus TCP client. This can be used to increase the timeout when a slow Modbus device or network is used.
- **Control:** Which register map should be used for mapping control entities to modbus registers.



If no control is required then this may be set to *None*. Setting this to *Use Register Map* will cause all the registers that are being read to also be targets for control. Setting this to *Use Control Map* will cause the separate *Control Map* to be used to map the control set points to modbus registers.

- **Control Map:** The register map that is used to map the set point names into Modbus registers for the purpose of set point control. The control map is the same JSON format document as the register map and uses the same set of properties.

Register Map

The register map is the most complex configuration parameter for this plugin and over time has supported a number of different variants. We will only document the latest of these here although previous variants are still supported. This latest variant is the most flexible to date and is thus the recommended approach to adopt.

The map is a JSON object with a single array *values*, each element of this array is a JSON object that defines a single item of data that will be stored in FogLAMP. These objects support a number of properties and values, these are

Property	Description
name	The name of the value that we are reading. This becomes the name of the data point with the asset. This may be either the default asset name defined plugin or an individual asset if an override is given.
slave	The Modbus slave ID of the device if it differs from the global Slave ID defined for the plugin. If not given the default Slave ID will be used.
asset-Name	This is an optional property that allows the asset name define for the plugin to be overridden on an individual basis. Multiple values in the values array may share the same AssetName, in which case the values read from the Modbus device are placed in the same asset. Note: This is unused in a control map.
register	This defines the Modbus register that is read. It may be a single register, in which case the value is the register number or it may be multiple registers in which case the value is a JSON array of numbers. If an array is given then the registers are read in the order of that array and combined into a single value by shifting each value up 16 bits and performing a logical OR operation with the next register in the array.
coil	This defines the number of the Modbus coil to read. Coils are single bit Modbus values.
input	This defines the number of the Modbus discrete input. Coils are single bit Modbus values.
inputRegister	This defines the Modbus input register that is read. It may be a single register, in which case the value is the register number or it may be multiple registers in which case the value is a JSON array of numbers. If an array is given then the registers are read in the order of that array and combined into a single value by shifting each value up 16 bits and performing a logical OR operation with the next register in the array.
scale	A scale factor to apply to the data that is read. The value read is multiplied by this scale. This is an optional property.
offset	An optional offset to add to the value read from the Modbus device.
type	This allows data to be cast to a different type. The only support type currently is <i>float</i> and is used to interpret data read from the one or more of the 16 bit registers as a floating point value. This property is optional.
swap	This is an optional property used to byte swap values read from a Modbus device. It may be set to one of <i>bytes</i> , <i>words</i> or <i>both</i> to control the swapping to apply to bytes in a 16 bit value, 16 bit words in a 32 bit value or both bytes and words in 32 bit values.

Every *value* object in the *values* array must have one and only one of *coil*, *input*, *register* or *inputRegister* included as this defines the source of the data in your Modbus device. These are the Modbus object types and each has an address space within a typical Modbus device.

Object Type	Size	Address Space	Map Property
Coil	1 bit	00001 - 09999	coil
Discrete Input	1 bit	10001 - 19999	input
Input Register	16 bits	30001 - 39999	inputRegister
Holding Register	16 bits	40001 - 49999	register

The values in the map for coils, inputs and registers are relative to the base of the address space for that object type rather than the global address space and each is 0 based. A map value that has the property *"coil" : 10* would return

the values of the tenth coil and “*register*” : 10 would return the tenth register.

Example Maps

In this example we will assume we have a cooling fan that has a Modbus interface and we want to extract three data items of interest. These items are

- Current temperature that is in Modbus holding register 10
- Current speed of the fan that is stored as a 32 bit value in Modbus holding registers 11 and 12
- The active state of the fan that is stored in a Modbus coil 1

The Modbus Map for this example would be as follow:

```
{
  "values" : [
    {
      "name"      : "temperature",
      "register"   : 10
    },
    {
      "name"      : "speed",
      "register"   : [ 11, 12 ]
    },
    {
      "name"      : "active",
      "coil"      : 1
    }
  ]
}
```

Since none of these values have an `assetName` defined all there values will be stored in a single asset, the name of which is the default asset name defined for the plugin as a whole. This asset will have three data points within it; *temperature*, *speed* and *active*.

Function Codes

The *foglamp-south-modbus-c* plugin attempts to make as few calls as possible to the underlying modbus device in order to collect the data. This is done in order to minimise the load that is placed on the modbus server. The modbus function codes used to read each coil or register type are as follows;

Object Type	Function Code	Size	Address Space	Map Property
Coil	01 Read Coils	1 bit	00001 - 09999	coil
Discrete Input	02 Read Discrete inputs	1 bit	10001 - 19999	input
Input Register	04 Read register	16 bits	30001 - 39999	inputRegister
Holding Register	16 Read multiple registers	16 bits	40001 - 49999	register

Set Point Control

The *foglamp-south-modbus-c* plugin supports the FogLAMP set point control mechanisms and allows a register map to be defined that maps the set point attributes to the underlying modbus registers. As an example a control map as follows

```
{
  "values" : [
    {
      "name" : "active",
      "coil" : 1
    }
  ]
}
```

Defines that a set point write operation can be instigated against the set point named *active* and this will map to the Modbus coil 1.

Set points may be defined for Modbus coils and registers, the read only input bits and input registers can not be used for set point control.

The *Control Map* can use the same swapping, scaling and offset properties as modbus *Register Map*, it can also map multiple registers to a single set point and floating point values.

Error Messages

The following are messages that may be produced by the *foglamp-south-modbus-c* plugin, these messages are written to the system log file and may be viewed by the *System* menu item in the FogLAMP user interface. This display may be filtered on the name of a particular south service in order to view just the messages that originate from that south service.

The value of slave in the modbus map should be an integer When a modbus slave identifier is defined within the JSON modbus map it should always be given as a integer value and should not be enclosed in quotes

```
"slave" : 0
```

The value of slave for item 'X' in the modbus map should be an integer A name entity in the modbus map is defined as a string and must be enclosed in double quotes. This error would indicate that a non-string value has been given.

```
"name" : "speed"
```

Each item in the modbus map must have a name property Each of the modbus entities that is read must define a name property for the entity.

```
"name" : "speed"
```

The value of assetName for item 'X' in the modbus map should be a string The optional property *assetName* must always be provided as a string in the modbus map.

```
"assetName" : "pumpSpeed"
```

The value of scale for item 'X' in the modbus map should be a floating point number The optional property *scale* must always be expressed as a numeric value in the JSON of the modbus map, and should not be enclosed in quotes.

```
"scale" : 1.4
```

The value of offset for item ‘X’ in the modbus map should be a floating point number The optional property *offset* must always be given as a numeric value in the JSON definition of the modbus item, and should not be enclosed in quotes.

```
"offset" : 2.0
```

The value of coil for item ‘X’ in the modbus map should be a number The coil number given in the modbus map of an item must be an integer number, and should not be enclosed in quotes.

```
"coil" : 22
```

The value of input for item ‘X’ in the modbus map must be either an integer The input number given in the modbus map of an item must be an integer number, and should not be enclosed in quotes.

```
"input" : 22
```

The value of register for item ‘X’ in the modbus map must be either an integer or an array The register to read for an entity must be either an integer number or in the case of values constructed from multiple registers it may be an array of integer numbers. Numeric values should not be enclosed on quotes.

```
"register" : 22
```

Or, if two registers are being combined

```
"register" : [ 18, 19 ]
```

The register array for item ‘X’ in the modbus map contain integer values When giving an array as the value of the register property for a modbus item, that array must only contain register numbers expressed as numeric values. Register numbers should not be enclosed in quotes.

```
"register" : [ 18, 19 ]
```

The value of inputRegister for item ‘X’ in the modbus map must be either an integer or an array The input register to read for an entity must be either an integer number or in the case of values constructed from multiple input registers it may be an array of integer numbers. Numeric values should not be enclosed on quotes.

```
"inputRegister" : 22
```

Or, if two input registers are being combined

```
"inputRegister" : [ 18, 19 ]
```

The type property of the item ‘X’ in the modbus map must be a string The optional *type* property for a modbus entity must be expressed as a string enclosed in double quotes.

```
"type" : "float"
```

The type property ‘Y’ of the item ‘X’ in the modbus map is not supported The *type* property of the item is not supported by the plugin. Only the type *float* is currently supported.

The swap property ‘Y’ of item ‘X’ in the modbus map must be one of bytes, words or both An unsupported option has been supplied as the value of the swap property, only *bytes*, *words* or *both* are supported values.

The swap property of the item ‘X’ in the modbus map must be a string The optional *swap* property of a modbus item must be given as a string in double quotes and must be one of the supported swap options.

```
"swap" : "bytes"
```

Item ‘X’ in the modbus map must have one of coil, input, register or inputRegister properties Each modbus item to be read from the modbus server must define how that item is addressed. This is done by adding a modbus property called *coil*, *input*, *register* or *inputRegister*.

Item ‘X’ in the modbus map must only have one of coil, input, register or inputRegister properties Each modbus item to be read from the modbus server must define how that item is addressed. This is done by adding a modbus property called *coil*, *input*, *register* or *inputRegister*, these are mutually exclusive and only one of them may be given per item in the modbus map.

N errors encountered in the modbus map A number of errors have been detected in the modbus map. These must be correct in order for the plugin to function correctly.

Parse error in modbus map, the map must be a valid JSON object. The modbus map JSON document has failed to parse. An additional text will be given that describes the error that has caused the parsing of the map to fail.

Parse error in control modbus map, the map must be a valid JSON object. The modbus control map JSON document has failed to parse. An additional text will be given that describes the error that has caused the parsing of the map to fail.

Failed to connect to Modbus device The plugin has failed to connect to a modbus device. In the case of a TCP modbus connection this could be because the address or port have been misconfigured or the modbus device is not currently reachable on the network. In the case of a modbus RTU device this may be a misconfiguration or a permissions issue on the entry in */dev* for the device. Additional information will be given in the error message to help identify the issue.

8.1.22 South MQTT

The *foglamp-south-mqtt-readings* plugin allows to create an MQTT subscriber service. MQTT Subscriber reads messages from topics on the MQTT broker.

To create a south service you, as with any other south plugin

- Select *South* from the left hand menu bar
- Click on the + icon in the top right
- Choose *mqtt-readings* from the plugin selection list
- Name your service
- Click on *Next*
- Configure the plugin

- **MQTT Broker host:** Hostname or IP address of the broker to connect to.
 - **MQTT Broker Port:** The network port of the broker.
 - **Keep Alive Interval:** Maximum period in seconds allowed between communications with the broker. If no other messages are being exchanged, this controls the rate at which the client will send ping messages to the broker.
 - **Topic To Subscribe:** The subscription topic to subscribe to receive messages.
 - **QoS Level:** The desired quality of service level for the subscription.
 - **Asset Name:** Name of Asset.
- Click *Next*
 - Enable your service and click *Done*

Message Payload

The content of the message payload published to the topic, to which the service is configured to subscribe, should be parsable to a JSON object.

e.g. `'{"humidity": 93.29, "temp": 16.82}'`

```
$ mosquitto_pub -h localhost -t "Room1/conditions" -m '{"humidity": 93.29, "temp": 16.82}'
```

The `mosquitto_pub` client utility comes with the `mosquitto` package, and a great tool for conducting quick tests and troubleshooting. https://mosquitto.org/man/mosquitto_pub-1.html

8.1.23 MQTT Sparkplug B

The `foglamp-south-mqtt-sparkplug` plugin implements the Sparkplug B payload format with an MQTT (Message Queue Telemetry Transport) transport. The plugin will subscribe to a configured topic and will process the Sparkplug B payloads, creating FogLAMP assets from those payloads. Sparkplug is an open source software specification of a payload format and set of conventions for transporting sensor data using MQTT as the transport mechanism.

Note: Sparkplug is bi-directional, however this plugin will only read data from the Sparkplug device.

To create a south service with the MQTT Sparkplug B plugin

- Click on *South* in the left hand menu bar
- Select *mqtt_sparkplug* from the plugin list
- Name your service and click *Next*

- Configure the plugin
 - **Asset Name:** The asset name which will be used for all data read.
 - **MQTT Host:** The MQTT host to connect to, this is the host that is running the MQTT broker.
 - **MQTT Port:** The MQTT port, this is the port the MQTT broker uses for unencrypted traffic, usually 1883 unless modified.
 - **Username:** The user name to be used when authenticating with the MQTT subsystem.
 - **Password:** The password to use when authenticating with the MQTT subsystem.
 - **Topic:** The MQTT topic to which the plugin will subscribe.
- Click *Next*
- Enable the service and click on *Done*

8.1.24 MQTT South with Payload Scripting

The *foglamp-south-mqtt-scripted* plugin uses MQTT to receive messages via an MQTT broker from sensors or other sources. It then uses an optional script, written in Python, that converts the message into a JSON document and pushes data to the FogLAMP System.

If the payload of the MQTT message is a JSON document with simple key/value pairs, e.g.

```
{ "temperature" : 23.1, "humidity" : 47.2 }
```

Then no translation script is required. Also if the payload is a simple numeric value the plugin will accept this and create an asset with the data point name matching the topic on which the value was given in the payload.

If the message format is not a simple JSON document or a single value, or is in some other format then a Python script should be provided that turns the message into a JSON format.

An example script, assuming the payload in the message is simply a value, might be as follows

```
def convert(message, topic):  
    return {  
        'temperature' : float(message)  
    }
```

Note that the message and topic are passed as a strings and the data we wish to ingest into FogLAMP in this case is assumed to be a floating point value. The example above of course is unnecessary as the plugin can consume this data without the need of a script.

The script could return either one or two values.

The script should return the JSON document as a Python DICT in the case of a single value.

The script should return a string and a JSON document as a Python DICT in the case of two values, the first of these values is the name of the asset to use and overrides the default asset naming defined in the plugin configuration.

First case sample:

```
def convert(message, topic):  
    return {"temperature_1": 10.2}
```

Second case sample:

```
def convert(message, topic):  
    return "ExternalTEMP", {"temperature_3": 11.3}
```

Limitations & Recommendations

When a script is provided it is best practice to do the minimum required to allow the data to be ingested into the FogLAMP data pipeline. Further processing to shape the data to exact requirements can often be done using an existing filter. The advantages of this are twofold; it simplifies the scripts required here and it simplifies maintenance should the data be required in a different format some time later.

Other filters exists, such as and that allow assets and data points to be renamed or to use regular expressions to substitute portions of asset names and data points. Filters such as these can be applied to the result of the MQTT scripted filter to convert the data into the form required and maintain a simpler Python script, or obviate the need for a Python script, in the MQTT scripted plugin.

Configuration

When adding a south service with this plugin the same flow is used as with any other south service. The configuration page for the plugin is as follows.

The screenshot displays the 'Review Configuration' step of the FogLAMP plugin configuration process. The interface features a progress bar at the top with three steps: '1 Plugin & Service Name', '2 Review Configuration' (the current step), and '3 Done'. Below the progress bar, the configuration form is organized into two columns. The left column lists the configuration parameters: Asset Name, MQTT Broker, Username, Password, Trusted Certificate, Client Certificate, Private Key, Key Password, Topic, Object Policy, and Script. The right column contains the input fields for these parameters. The 'Asset Name' field is populated with 'mqtt'. The 'MQTT Broker' field is populated with 'localhost'. The 'Username' field is empty. The 'Password' field is a greyed-out text box. The 'Trusted Certificate', 'Client Certificate', and 'Private Key' fields are also greyed-out text boxes. The 'Key Password' field is populated with 'password' and has a small icon to its right. The 'Topic' field is populated with 'sensor'. The 'Object Policy' field is a dropdown menu set to 'Single reading from root level'. The 'Script' field is a large text area with the number '1' in the top left corner. At the bottom of the form, there is a 'Choose files' button and the text 'No file chosen'.

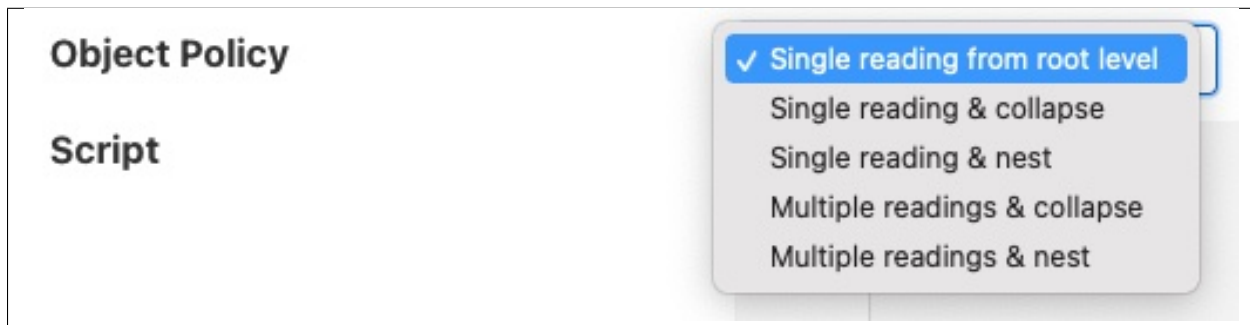
- **Asset Name:** The name of the asset the plugin will create for each message, unless the convert function returns an explicit asset name to be used.
- **MQTT Broker:** The IP address/hostname of the MQTT broker to use. Note FogLAMP requires an external MQTT broker is run currently and does not provide an internal broker in the current release.
- **Username:** The username to be used if required for authentication. This should be left blank if authentication is not required.
- **Password:** The password to use if username is to be used.
- **Trusted Certificate:** The trusted certificate of the MQTT broker. If MQTTS communication is not required

then this can be left blank.

- **Client Certificate:** The certificate that will be used by the MQTT plugin.
- **MQTTS Key:** The private key of the MQTT plugin. If the key is included in the PEM file of the client certificate this may be left blank.
- **Key Password:** The password used to encrypted the private key. This may be left blank if the private key was not encrypt.
- **Topic:** The MQTT topic to which to subscribe. The topic may include the usual MQTT wildcards; + for a single level wildcard and # for a multi-level wildcard
- **Object Policy:** Controls how the plugin deals with nested objects within the JSON payloads it receives or the return from the script that is executed. See below for a description of the various object policy values.
- **Time Format:** The format to both pass the timestamps into the query parameters using and also to interpret the timestamps returned in the payload.
- **Timezone:** The timezone to use for the start and end times that are sent in the API request and also when timestamps are read from the API response. Timezone is expressed as an offset in hours and minutes from UTC for the local timezone of the API. E.g. -08:00 for PST time zones.
- **Script:** The Python script to execute for message processing. Initially a file must be uploaded, however once uploaded the user may edit the script in the box provided. A script is optional.

Object Policy

The object policy is used by the plugin to determine how it deals with nested objects within the JSON that is in the MQTT payload or the JSON that is returned from the script that is executed, if present.



- **Single reading from root level:** This is the simple behavior of the plugin, it will only take numeric and string values that are in the root of the JSON document and ignore any objects contained in the root.
- **Single reading & collapse:** The plugin will create a single reading form the payload that will contain the string and numeric data in the root level. The plugin will also recursively traverse any child objects and add the string and numeric data from those to the reading as data points of the reading itself.
- **Single reading & nest:** As above, the plugin will create a single reading form the payload that will contain the string and numeric data in the root level. The plugin will also recursively traverse any child objects and add the string and numeric data from those objects and add them as nested data points.
- **Multiple readings & collapse:** The plugin will create one reading that contains any string and numeric data in the root of the JSON. It will then create one reading for each object in the root level. Each of these readings will contain the string and numeric data from those child objects along with the data found in the children of those objects. Any child data will be collapse into the base level of the readings.

- **Multiple readings & nest:** As above, but any data in the children of the readings found below the first level, which defines the reading names, will be created as nested data points rather than collapsed.

As an example of how the policy works assume we have an MQTT payload with a message as below

```
{
  "name" : "pump47",
  "motor" : {
    "current" : 0.75,
    "speed" : 1496
  },
  "flow" : 1.72,
  "temperatures" : {
    "bearing" : 21.5,
    "impeller" : 16.2,
    "motor" : {
      "casing" : 24.6,
      "gearbox" : 28.2
    }
  }
}
```

If the policy is set to *Single reading from root level* then a reading would be created, with the asset name given in the configuration of the plugin, that contained two data points *name* and *flow*.

If the policy is set to *Single reading & collapse* then the reading created would now have 8 data points; *name*, *current*, *speed*, *flow*, *bearing*, *impeller*, *casing* and *gearbox*. These would all be in a reading with the asset name defined in the configuration and in a flat structure.

If the policy is set to *Single reading & nest* there would still be a single reading, with the asset name set in the configuration, which would have data points for *name*, *motor*, *flow* and *temperature*. The *motor* data point would have two child data points called *current* and *speed*, the *temperature* data point would have three child data points called *bearing*, *impeller* and *motor*. This *motor* data point would itself have two children call *casing* and *gearbox*.

If the policy is set to *Multiple readings & collapse* there would be three readings created from this payload; one that is names as per the asset name in the configuration, a *motor* reading and a *temperature* reading. The first of these readings would have data points called *name* and *flow*, the *motor* reading would have data points *current* and *speed*. The *temperatures* reading would have data points *bearing*, *impeller*, *casing* and *gearbox*.

If the policy is set to *Multiple readings & nest* there would be three readings created from this payload; one that is names as per the asset name in the configuration, a *motor* reading and a *temperature* reading. The first of these readings would have data points called *name* and *flow*, the *motor* reading would have data points *current* and *speed*. The *temperatures* reading would have data points *bearing*, *impeller* and *motor*, the *motor* data point would have two child data points *casing* and *gearbox*.

Timestamp Treatment

The default timestamp for a reading collected via this plugin will be the time at which the reading was taken, however it is possible for the API that is being called to include a different timestamp.

Returning a data point called whose name is defined in the *Timestamp* configuration option will result in the value of that data point being used as the timestamp. This data point will not be added to the reading. The default name of the timestamp is *timestamp*.

The timestamp data point should be a string and the timestamp should be formatted to match the definition given in the *Time format* configuration parameter. The format is based on the standard Linux strptime formatting options and is discussed below in the section discussing the *Time Format* selection method.

The timezone may be set by using the *Timezone* configuration parameter to set the offset of the timezone in which the API is running.

Time Format

The format of the timestamps read in the message payload or by the script returned are defined by the *Time Format* configuration parameter and uses the standard Linux mechanism to define a time format. The following character sequences are supported.

% % The % character.

%a or %A The name of the day of the week according to the current locale, in abbreviated form or the full name.

%b or %B or %h

The month name according to the current locale, in abbreviated form or the full name.

%c The date and time representation for the current locale.

%C The century number (0–99).

%d or %e The day of month (1–31).

%D Equivalent to %m/%d/%y. (This is the American style date, very confusing to non- Americans, especially since %d/%m/%y is widely used in Europe. The ISO 8601 standard format is %Y-%m-%d.)

%H The hour (0–23).

%I The hour on a 12-hour clock (1–12).

%j The day number in the year (1–366).

%m The month number (1–12).

%M The minute (0–59).

%n Arbitrary white space.

%p The locale’s equivalent of AM or PM. (Note: there may be none.)

%r The 12-hour clock time (using the locale’s AM or PM). In the POSIX locale equivalent to %I:%M:%S %p. If t_fmt_ampm is empty in the LC_TIME part of the current locale, then the behavior is undefined.

%R Equivalent to %H:%M.

%S The second (0–60; 60 may occur for leap seconds; earlier also 61 was allowed).

%t Arbitrary white space.

%T Equivalent to %H:%M:%S.

%U The week number with Sunday the first day of the week (0–53). The first Sunday of January is the first day of week 1.

%w The ordinal number of the day of the week (0–6), with Sunday = 0.

%W The week number with Monday the first day of the week (0–53). The first Monday of January is the first day of week 1.

%x The date, using the locale’s date format.

%X The time, using the locale’s time format.

%y The year within century (0–99). When a century is not otherwise specified, values in the range 69–99 refer to years in the twentieth century (1969–1999); values in the range 00–68 refer to years in the twenty-first century (2000–2068).

%Y The year, including century (for example, 1991).


Script Error Handling

If an error occurs in the plugin or Python script, including script coding errors and Python exception, details will be logged to the error log and data will not flow through the pipeline to the next filter or into the storage service.

Warnings raised will also be logged to the error log but will not cause data to cease flowing through the pipeline.

To view the error log you may examine the file directly on your host machine, for example `/var/log/syslog` on a Ubuntu host, however it is also possible to view the error logs specific to Fledge from the Fledge user interface. Select the *System* option under *Logs* in the left hand menu pane. You may then filter the logs for a specific service to see only those logs that refer to the service which uses the filter you are interested in.

Alternatively if you open the dialog for the service in the *South* or *North* menu items you will see two icons displayed in the bottom left corner of the dialog that lets you alter the configuration of the service.



The left most icon, with the ? in a circle, allows you to view the documentation for the plugin, the right most icon, which looks like a page of text with a corner folded over, will open the log view page filtered to view the service.

Error Messages & Warnings

The following are some errors you may see within the log with some description of the cause and remedy for the error.

The supplied Python script does not define a valid “convert” function The script that has been supplied does not define a Python function called convert. The script must provide a single function called convert that accepts the MQTT payload and topic and will process these to provide the JSON DICT and an optional asset name to import.

Python error: IndentationError ‘expected an indented block’ in XXXX at line Y of script The script supplied does not conform to Python requirements for code block indentation. The text XXXX will be replaced with the line of text in error and Y with the line number within the script.

Python error: SyntaxError ‘invalid syntax’ in XXXX at line Y of script The script supplied does has invalid Python syntax. The text XXXX will be replaced with the line of text in error and Y with the line number within the script.

Python error: ModuleNotFoundError “No module named ‘nosuchpackage’” in supplied script The script supplied is attempting to import a Python module that is not available.

Python error: TypeError “convert() missing 1 required positional argument: ‘name’” in supplied script The type of the convert function has been incorrectly defined. The convert function should take a single argument which is the message to process.

Return from Python convert function is of an incorrect type, it should be a Python DICT object or a DICT object and a string The convert function is returning data of an incorrect type. It may either return a Python DICT, which may be empty, None or a string and a Python DICT.

The plugin is unable to process data without a valid ‘convert’ function in the script. This warning will periodically be logged following an earlier error that has resulted in an error which prevents the Python convert function from processing the messages. Fix the earlier error to stop this warning being logged.

Unable to process message ‘XXXX’ expecting a simple value This warning is logged if there is no script defended for the plugin and the message is not simply a numeric value. In this case a Python script should be added that processes the payload.

The returned asset name was None, either a valid string must be returned or the asset name may be omitted The python script has returned a pair of values, but the asset name returned is None. If an asset name is returned it must be a string. If no asset name is required then it can be omitted from the return value of the script.

8.1.25 OPC/UA South Plugin

The *foglamp-south-opcua* plugin allows FogLAMP to connect to an OPC/UA server and subscribe to changes in the objects within the OPC/UA server.

A south service to collect OPC/UA data is created in the same way as any other south service in FogLAMP.

- Use the *South* option in the left hand menu bar to display a list of your South services
- Click on the + add icon at the top right of the page
- Select the *opcua* plugin from the list of plugins you are provided with
- Enter a name for your south service
- Click on *Next* to configure the OPC/UA plugin

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name

OPCUA Server URL

OPCUA Object Subscriptions

```

1 {
2   "subscriptions": [
3     "ns=5;s=85/0:Simulation"
4   ]
5 }

```

Subscribe By ID ☒

Asset Name Source

Asset Path Delimiter

Min Reporting Interval

[Previous](#) [Next](#)

The configuration parameters that can be set on this page are;

- **Asset Name:** This is a prefix that will be applied to all assets that are created by this plugin. The OPC/UA plugin creates a separate asset for each data item read from the OPC/UA server. This is done since the OPC/UA server will deliver changes to individual data items only. Combining these into a complex asset would result in assets that contain only one of many data points in each update. This can cause problems in upstream systems with the ever-changing asset structure.
- **OPCUA Server URL:** This is the URL of the OPC/UA server from which data will be extracted. The URL should be of the form `opc.tcp://.../`
- **OPCUA Object Subscriptions:** The subscriptions are a set of locations in the OPC/UA object hierarchy that defined which data is subscribed to in the server and hence what assets get created within FogLAMP. A fuller description of how to configure subscriptions is shown below.
- **Subscribe By ID:** This toggle determines if the OPC/UA objects in the subscription are using names to identify the objects in the OPC/UA object hierarchy or using object ID's.
- **Asset Name Source:** This drop-down allows you to choose the source of the name for the Asset in FogLAMP. The choices are:
 - *NodeId*: the Node Id of the OPC UA node. This is the default.
 - *BrowseName*: the Browse Name of the OPC UA node.
 - *Subscription Path with NodeId*: the path to the node in the OPC/UA server's object hierarchy starting with the node specified in *Subscriptions*. Every node in the hierarchy is named with its Node Id. The Node Id is also used as the Data Point name.

- *Subscription Path with BrowseName*: same as *Subscription Path with NodeId* except that the Browse Name is used as the name of every node in the hierarchy. The Browse Name is also used as the Data Point name.
- *Full Path with NodeId*: the path to the node in the OPC/UA server's object hierarchy starting with the top-level *Objects* folder. Every node in the hierarchy is named with its Node Id. The *Objects* folder itself is not part of the full path. The Node Id is also used as the Data Point name.
- *Full Path with BrowseName*: same as *Full Path with NodeId* except that the Browse Name is used as the name of every node in the hierarchy. The Browse Name is also used as the Data Point name.
- **Asset Path Delimiter**: A character to separate segments of the Asset Path. The delimiter is a single character. If multiple characters are specified, the string will be truncated to one character. The default is the forward slash ("/").
- **Min Reporting Interval**: This controls the minimum interval between reports of data changes in subscriptions. It sets an upper limit to the rate that data will be ingested into the plugin and is expressed in milliseconds.

Subscriptions

Subscriptions to OPC/UA objects are stored as a JSON object that contains an array named "subscriptions". This array is a set of OPC/UA nodes that will control the subscription to variables in the OPC/UA server.

The array may be empty, in which case all variables are subscribed to in the server and will create assets in FogLAMP. Note that simply subscribing to everything will return a lot of data that may not be of use.

If the *Subscribe By ID* option is set then this is an array of Node Id's. Each Node Id should be of the form *ns=...;s=...*. Where *ns* is a namespace index and *s* is the Node Id string identifier. A subscription will be created with the OPC/UA server for the object with the specified Node Id and its children, resulting in data change messages from the server for those objects. Each data change received from the server will create an asset in FogLAMP with the name of the object prepended by the value set for *Asset Name*. An integer identifier is also supported by using a Node Id of the form *ns=...;i=...*.

If the *Subscribe By ID* option is not set then the array is an array of Browse Names. The format of the Browse Names is <namespace>:<name>. If the namespace is not required then the name can simply be given, in which case any name that matches in any namespace will have a subscription created. The plugin will traverse the node tree of the server from the *ObjectNodes* root and subscribe to all variables that live below the named nodes in the subscriptions array.

Configuration examples

```
{ "subscriptions": [ "5:Simulation", "2:MyLevel" ] }
```

We subscribe to

- 5:Simulation is a node name under ObjectsNode in namespace 5
- 2:MyLevel is a variable under ObjectsNode in namespace 2

```
{ "subscriptions": [ "5:Sinusoid1", "2:MyLevel", "5:Sawtooth1" ] }
```

We subscribe to

- 5:Sinusoid1 and 5:Sawtooth1 are variables under ObjectsNode/Simulation in namespace 5
- 2:MyLevel is a variable under ObjectsNode in namespace 2

```
{ "subscriptions": [ "2:Random.Double", "2:Random.Boolean" ] }
```

We subscribe to

- Random.Double and Random.Boolean are variables under ObjectsNode/Demo both in namespace 2

It's also possible to specify an empty subscription array:

```
{"subscriptions":[]}
```

Note: Depending on OPC/UA server configuration (number of objects, number of variables) this empty configuration might take a long time to create the subscriptions and hence delay the startup of the south service. It will also result in a large number of assets being created within FogLAMP.

Object names, variable names and NamespaceIndexes can be easily retrieved browsing the given OPC/UA server using OPC UA clients, such as .

8.1.26 Person Detection Plugin

The *foglamp-south-person-detection* detects a person on a live video feed from either a camera or on a network stream. It uses Google's Mobilenet SSD v2 to detect a person. The bounding boxes and confidence scores are displayed on the same video frame itself. Also FPS (frames per second) are also displayed on the same frame. The detection results are also converted into readings. The readings have mainly three things:

1. *Count* : The number of people detected.
2. *Coordinates* : It consists of coordinates (x,y) of top-left and bottom right corners of bounding box for each detected person.
3. *Confidence* : Confidence with which the model detected each person.

rtrr South Service	
TFlite Model File	detect.tflite
Labels File	coco_labels.txt
Asset Name	person_detection_4
Enable Edge TPU	<input type="checkbox"/>
Minimum Confidence Threshold	0.5
Source For Detection Camera/Stream	stream
Stream URL	rtsp://192.168.0.109:8554/clip
Opencv Backend	ffmpeg
Stream Protocol	udp
Camera ID	0
Enable Detection Window	<input type="checkbox"/>

- **TFlite Model File:** This is the name of the tflite model file that should be placed in `python/foglamp/plugins/south/person_detection/model` directory. Its default value is `detect_edgetpu.tflite`. If a Coral Edge TPU is not being used, the file name will be different (i.e. `detect.tflite`).
- **Labels File:** This is the name of the labels file that was used when training the above model, this file should also be placed in same directory as the model.

- **Asset Name:** The name of the asset used for the readings generated by this plugin.
- **Enable Edge TPU:** Indicates whether to use edge TPU for inference. If you don't want to use Coral Edge TPU then disable this configuration parameter. Also ensure to change the name of the model file to detect.tflite if disabled. Default is set to enabled.
- **Minimum Confidence Threshold:** The detection results from the model will be filtered out, if the score is below this value.
- **Source:** Either use a stream over a network or use a local camera device. Default is set to stream.
- **Streaming URL:** The URL of the RTSP stream, if stream is to be used. Only RTSP streams are supported for now.
- **OpenCV Backend:** The backend required by OpenCV to process the stream, if stream is to be used. Default is set to ffmpeg.
- **Streaming Protocol:** The protocol over which live frames are being transported over the network, if stream is to be used. Default is set to udp.
- **Camera ID:** The number associated with your video device. See /dev in your filesystem you will see video0 or video1. It is required when source is set to camera. Default is set to 0.
- **Enable Detection Window:** Show detection results in a native window. Default is set to disabled.

rtrr South Service

Stream Protocol

udp

Camera ID

0

Enable Detection Window

☐

Enable Web Streaming

☒

Web Streaming Port

8085

Enabled

☐

[Show Advanced Config](#)

Applications

Cancel

Save

?

Export Readings

Delete Service

- **Enable Web Streaming:** Whether to stream the detected results in a browser or not. Default is set to enabled.
- **Web Streaming Port:** Port number where web streaming server should run, if web streaming is enabled. Default is set to 8085.

Installation

1. First run requirements.sh

There are two ways to get the video feed.

1. Camera

To see the supported configuration of the camera run the following command.

```
$ v4l2-ctl --list-formats-ext --device /dev/video0
You will see something like
'YUYV' (YUYV 4:2:2)
  Size: Discrete 640x480
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 720x480
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 1280x720
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 1920x1080
    Interval: Discrete 0.067s (15.000 fps)
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 2592x1944
    Interval: Discrete 0.067s (15.000 fps)
  Size: Discrete 0x0
```

Above example uses Camera ID 0 to indicate use of /dev/video0 device, please use the applicable value for your setup

2. Network RTSP stream

To create a network stream follow the following steps

1. Install vlc

```
$ sudo add-apt-repository ppa:videolan/master-daily
$ sudo apt update
$ apt show vlc
$ sudo apt install vlc qtwayland5
$ sudo apt install libavcodec-extra
```

2. Download some sample files from here.

```
$ git clone https://github.com/intel-iot-devkit/sample-videos.git
```

3. Either stream a file using the following

```
$ vlc <name_of_file>.mp4 --sout '#gather:transcode{vcodec=h264,vb=512,
↪scale=Auto,width=640,height=480,acodec=none,scodec=none}:rtp{sdp=rtsp://
↪<ip_of_machine_streaming>:8554/clip}' --no-sout-all --sout-keep --loop --
↪no-sout-audio --sout-x264-profile=baseline
```

Note : fill the <ip_of_the_machine> with ip of the machine which will be used to stream video. Also fill <name_of_file> with the name of mp4 file.

4. You can also stream from a camera using the following

```
$ vlc v4l2:///dev/video<index_of_video_device> --sout '#gather:transcode
↪{vcodec=h264,vb=512,scale=Auto,width=<supported_width_of_camera_image>,
↪height=<supported_height_of_camera_image>,acodec=none,scodec=none}:rtp
↪{sdp=rtsp://<ip_of_the_machine>:8554/clip}' --no-sout-all --sout-keep --
↪no-sout-audio --sout-x264-profile=baseline
```

(continues on next page)

(continued from previous page)

Fill the following :

<index_of_video_device> The index with which you ran the v4l2 command mentioned above. for example video0.

<supported_height_of_camera_image> Height you get when you ran v4l2 command mentioned above. For example Discrete 640x480. Here 480 is height.

<supported_width_of_camera_image> Width you get when you ran v4l2 command mentioned above. For example Discrete 640x480. Here 640 is width.

<ip_of_the_machine> ip of the machine which will be used to stream video.

Once you have run the plugin by filling appropriate parameters Now go to your browser and enter *ip_where_foglamp_is_running:the_port_for_web_streaming*

8.1.27 Person Detector Plugin

The *foglamp-south-person-detector* detects a person on a live video feed from either a camera or on a network stream. It uses Google's Mobilenet SSD v2 to detect a person. The bounding boxes and confidence scores are displayed on the same video frame itself. Also FPS (frames per second) are also displayed on the same frame. The detection results are also converted into readings. The readings have mainly three things:

1. *Count* : The number of people detected.
2. *Coordinates* : It consists of coordinates (x,y) of top-left and bottom right corners of bounding box for each detected person.
3. *Confidence* : Confidence with which the model detected each person.

pd3 South Service

TFLite model name

Model version

Asset Name

Enable Edge TPU

Minimum Confidence Threshold

Source of video feed

Stream URL

OpenCV Backend

Stream Protocol

Camera ID

Enable Detection Window

Enable Web Streaming

Web Streaming Port

Enabled

People

1.2

person_detector

☒

0.5

stream

rtsp://ip:port

ffmpeg

udp

0

☐

☒

8085

☒

Show Advanced Config

Applications

- **TFlite model name:** This is the name of the tflite model as stored with Bucket service Its default value is “People”
- **Model version:** Model version as stored with Bucket service Its default value is “1.2”
- **Asset Name:** The name of the asset used for the readings generated by this plugin.
- **Enable Edge TPU:** Indicates whether to use edge TPU for inference. If you don’t want to use Coral Edge TPU then disable this configuration parameter. Also ensure to change the name of the model file to detect.tflite if disabled. Default is set to enabled.
- **Minimum Confidence Threshold:** The detection results from the model will be filtered out, if the score is below this value.
- **Source:** Either use a stream over a network or use a local camera device. Default is set to stream.
- **Streaming URL:** The URL of the RTSP stream, if stream is to be used. Only RTSP streams are supported for now.
- **OpenCV Backend:** The backend required by OpenCV to process the stream, if stream is to be used. Default is set to ffmpeg.
- **Streaming Protocol:** The protocol over which live frames are being transported over the network, if stream is to be used. Default is set to udp.
- **Camera ID:** The number associated with your video device. See /dev in your filesystem you will see video0 or video1. It is required when source is set to camera. Default is set to 0.
- **Enable Detection Window:** Show detection results in a native window. Default is set to disabled.
- **Enable Web Streaming:** Whether to stream the detected results in a browser or not. Default is set to enabled.
- **Web Streaming Port:** Port number where web streaming server should run, if web streaming is enabled. Default is set to 8085.

Installation

1. First run requirements.sh

There are two ways to get the video feed.

1. **Camera**

To see the supported configuration of the camera run the following command.

```
$ v4l2-ctl --list-formats-ext --device /dev/video0
You will see something like
'YUYV' (YUYV 4:2:2)
  Size: Discrete 640x480
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 720x480
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 1280x720
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 1920x1080
    Interval: Discrete 0.067s (15.000 fps)
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 2592x1944
    Interval: Discrete 0.067s (15.000 fps)
  Size: Discrete 0x0
```

Above example uses Camera ID 0 to indicate use of /dev/video0 device, please use the applicable value for your setup

2. Network RTSP stream

To create a network stream follow the following steps

1. Install vlc

```
$ sudo add-apt-repository ppa:videolan/master-daily
$ sudo apt update
$ apt show vlc
$ sudo apt install vlc qtwayland5
$ sudo apt install libavcodec-extra
```

2. Download some sample files from here.

```
$ git clone https://github.com/intel-iot-devkit/sample-videos.git
```

3. Either stream a file using the following

```
$ vlc <name_of_file>.mp4 --sout '#gather:transcode{vcodec=h264,vb=512,
↪scale=Auto,width=640,height=480,acodec=none,scodec=none}:rtp{sdp=rtsp://
↪<ip_of_machine_streaming>:8554/clip}' --no-sout-all --sout-keep --loop --
↪no-sout-audio --sout-x264-profile=baseline
```

Note : fill the <ip_of_the_machine> with ip of the machine which will be used to stream video. Also fill <name_of_file> with the name of mp4 file.

4. You can also stream from a camera using the following

```
$ vlc v4l2:///dev/video<index_of_video_device> --sout '#gather:transcode
↪{vcodec=h264,vb=512,scale=Auto,width=<supported_width_of_camera_image>,
↪height=<supported_height_of_camera_image>,acodec=none,scodec=none}:rtp
↪{sdp=rtsp:///<ip_of_the_machine>:8554/clip}' --no-sout-all --sout-keep -
↪no-sout-audio --sout-x264-profile=baseline
```

Fill the following :

<index_of_video_device> The index with which you ran the v4l2 command mentioned above. for example video0.

<supported_height_of_camera_image> Height you get when you ran v4l2 command mentioned above. For example Discrete 640x480. Here 480 is height.

<supported_width_of_camera_image> Width you get when you ran v4l2 command mentioned above. For example Discrete 640x480. Here 640 is width.

<ip_of_the_machine> ip of the machine which will be used to stream video.

Once you have run the plugin by filling appropriate parameters Now go to your browser and enter *ip_where_foglamp_is_running:the_port_for_web_streaming*

8.1.28 PI Web API south Plugin

The *foglamp-south-piwebapi* plugin is a south plugin that reads a PI Point and the related attributes from PI Web API. The plugin extracts the last value stored in the PI Point.

Using the Plugin

To create a south service with the PI Web API plugin

- Click on *South* in the left hand menu bar
- Select *PIWebAPI* from the plugin list
- Name your service and click *Next*

1 Plugin & Service Name 2 Review Configuration 3 Done

Hostname

Server port, 0=use the default

Authentication Method

User Id

Password

PIPoint

Attributes

```
1 {  
2   "items": []  
3 }
```

PI Server type

Server instance name

Database to use

Path on the server

Previous Next

- Configure the plugin

- **Hostname:** The name or IP address of the PI Web API server.
- **Server port:** The port on which the PI Web API server is listening. 0 means to use the default 443 port.
- **Authentication Method:** The authentication method requested by the PI Web API server, it could be either *basic* or *anonymous*, if *basic* is selected *user id* and *password* are required.

Authentication Method	basic ▼
User Id	anonymous
	basic

- **User Id:** The user id on the PI Web API server to allow the *basic* authentication.
- **Password:** The password associated to the user on the PI Web API server.
- **PIPoint:** The name of the PI Point on PI Web API for which the data should be extracted.
- **Attributes:** The attributes of the PI Point to extract. It can be either a single attribute or multiple attributes expressed as a json array, an example:

Attributes	<pre> 1 { 2 "items": [3 "attr_4_1", 4 "attr_4_2", 5 "attr_4_3" 6] 7 }</pre>
-------------------	---

- **Server type:** It allows to select the PI Server type either PI Asset Framework or PI Data Archive.

PI Server type	PI Asset Framework ▼
Server instance name	PI Asset Framework
	PI Data Archive

- **Server instance name:** It specifies server instance to be used.
- **Database to use:** Available only in case of PI Asset Framework, it specifies the Asset Framework database from which the data should be extracted.
- **Path on the server:** Available only in case of PI Asset Framework, the path of the PI Web API hierarchy that should be traversed to identify the position from which the data should be extracted, an example:

Path on the server	foglamp/r4dev2_4758
---------------------------	---------------------

8.1.29 Playback Plugin

The *foglamp-south-playback* plugin is a feature rich plugin for playing back comma separated variable (CSV) files. It supports features such as;

- Header rows
- User defined column names
- Use of historic or current timestamps
- Multiple timestamp formats
- Pick and optionally rename columns
- Looped or single pass readings of the data

To create a south service with the playback plugin

- Click on *South* in the left hand menu bar
- Select *playback* from the plugin list
- Name your service and click *Next*

The screenshot shows the 'Review Configuration' step (step 2) of the plugin setup. The interface includes a progress bar at the top with three steps: 1. Plugin & Service Name, 2. Review Configuration, and 3. Done. The configuration form contains the following fields and options:

- Asset Name:** sample
- CSV file name with extension:** some.csv
- Header Row:** ☒
- Header columns:** None
- Cherry pick column with same/new name:** A table with one row: 1 | {}
- Historic timestamps:** ☐
- Pick timestamp delta from file:** ☐
- Timestamp column name:** ts
- Timestamp format:** %Y-%m-%d %H:%M:%S.%f
- Ingest mode:** batch (dropdown menu)
- Sample Rate:** 100
- Burst Interval (ms):** 1000
- Burst size:** 1
- Read file in a loop:** ☐

At the bottom of the form are 'Previous' and 'Next' buttons.

- Configure the plugin
 - **Asset Name:** An asset name to use for the content of the file.

- **CSV file name with extension:** The name of the file that is to be processed, the file must be located in the foglamp data directory.
 - **Header Row:** Toggle to indicate the first row is a header row that contains the names that should be used for the data points within the asset.
 - **Header Columns:** Only used if *Header Row* is not enabled. This parameter should be a column separated list of column names that will be used to name the data points within the asset.
 - **Cherry pick column with same/new name:** This is a JSON document that can define a set of columns to include and optionally names to give those columns. If left empty then all columns, are included.
 - **Historic timestamps:** A toggle field to control if the timestamp data should be the current time or a date and time taken from the file itself.
 - **Pick timestamp delta from file:** If current timestamps are used then this option can be used to maintain the same relative times between successive timestamps added to the data as it is ingested.
 - **Timestamp column name:** The name of the column that should be used for reading timestamp value. This must be given if either historic timestamps are used or the interval between readings is to be maintained.
 - **Timestamp format:** The format of the timestamp within the file.
 - **Ingest mode:** Determine if ingest should be in batch or burst mode. In burst mode data is ingested as a set of bursts of rows, defined by *Burst size*, every *Burst Interval*, this allows simulation of sensors that have internal buffering within them. Batch mode is the normal, regular rate ingest of data.
 - **Sample Rate:** The data sampling rate that should be used, this is defined in readings per second.
 - **Burst Interval (ms):** The time interval between consecutive bursts when burst mode is used.
 - **Burst size:** The number of readings to be sent in each burst.
 - **Read file in a loop:** Once the end of the file is reached then the plugin will go back to the start and resend the data if this toggle is on.
- Click *Next*
 - Enable the service and click on *Done*

Picking Columns

The *Cherry pick column with same/new name* entry is a JSON document with a set of key/value pairs. The key is the name of the column in the file and the value is the name which should appear in the final asset. To illustrate this let's assume we have a CSV file as follows

```
X, Y, Z, Amps, Volts
1.3, 0.1, 0.3, 2.1, 240
1.2, 0.3, 0.2, 2.2, 235
....
```

We want to create an asset that has the *X* and *Y* values, *Amps* and *Volts* but we want to name them *X*, *Y*, *Current*, *Voltage*. We can do this by creating a JSON document that maps the columns.

```
{
  "X" : "X",
  "Y" : "Y",
  "Amps" : "Current",
  "Volts" : "Voltage"
}
```

Since we only mention the columns *X*, *Y*, *Amps* and *Volts*, only these will be included in the asset and we will not include the column *Z*. We map the column name *X* to *X*, so it will be unchanged. As will the column *Y*, the column *Amps* will become the data point *Current* and *Volts* will become *Voltage*.

8.1.30 PT100 Temperature Sensor

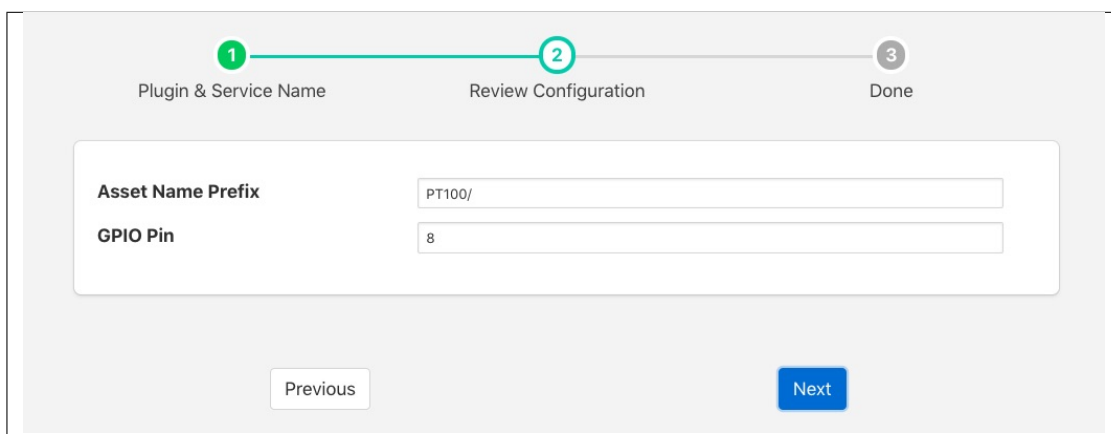


The *foglamp-south-pt100* is a south plugin for the PT-100 temperature sensor. The PT100 is a resistance temperature detectors (RTDs) consist of a fine wire (typically platinum) wrapped around a ceramic core, exhibiting a linear increase in resistance as temperature rises. The sensor connects via a MAX31865 converter to a GPIO pins for I2C bus and a chip select pin.

Note: This plugin is only available for the Raspberry Pi as it requires to be interfaced to the I2C bus on the Raspberry Pi GPIO header socket.

To create a south service with the PT100

- Click on *South* in the left hand menu bar
- Select *pt100* from the plugin list
- Name your service and click *Next*

A screenshot of the FogLAMP web interface showing the configuration steps for a new south service. At the top, a progress bar has three steps: 1. Plugin & Service Name, 2. Review Configuration (which is the current step and highlighted with a green circle), and 3. Done. Below the progress bar, there is a form with two fields: 'Asset Name Prefix' with the value 'PT100/' and 'GPIO Pin' with the value '8'. At the bottom of the form, there are two buttons: 'Previous' and 'Next'.


- Configure the plugin
 - **Asset Name Prefix:** A prefix to add to the asset name
 - **GPIO Pin:** The GPIO pin on the Raspberry PI to which the MAX31865 chip select is connected.
- Click *Next*
- Enable the service and click on *Done*

Wiring The Sensor


The MAX31865 uses the I2C bus on the Raspberry PI, which requires three wires to connect the bus, it also requires a chip select pin to be wired to a general GPIO pin and power.

MAX 31865 Pin	Raspberry Pi Pin
Vin	3V3
GND	GND
SDI	MOSI
SDO	MISO
CLK	SCLK
CS	GPIO (default GPIO8)

There are two options for connecting a PT100 to the MAX31865, a three wire PT100 or a four wire PT100.



To connect a four wire PT100 to the MAX 31865 the wires are connected in pairs, the two red wires are connected to the RTD- connector pair on the MAX31865 and the two remaining wires are connected to the RTD+ connector pair. If your PT100 does not have red wires or you wish to verify the colours are correct use a multimeter to measure the resistance across the pair of wires. Each pair should show 2 ohms between them and the difference between the two pairs should be 102 ohms, but will vary with temperature.



To connect a three wire sensor connect the red pair of wires across the RTD+ pair of connectors and the third wire on the RTD- block. If your PT100 does not have a pair of red wires, or you wish to verify the colours and have access to a multimeter, the resistance between the red wires should be 2 ohms. ~The resistance to the third wire, from the red pair, will be approximately 102 ohms but will vary with temperature.

If using the 3 wire sensor you must also modify the jumpers on the MAX31865.



Create a solder bridge across the 2/3 Wire jumper, outlined in red in the picture above.

You must also cut the thin wire trace on the jumper block outlined in yellow that runs between the 2 and 4.

Then create a new connection between the 4 and 3 side of this jumper block. This is probably best done with a solder bridge.

8.1.31 Random

The *foglamp-south-random* plugin is a plugin that will create random data.

To create a south service with the Random plugin


- Click on *South* in the left hand menu bar
- Select *Random* from the plugin list
- Name your service and click *Next*

The screenshot shows a configuration wizard for the 'Random' plugin. At the top, a progress bar has three steps: '1 Plugin & Service Name', '2 Review Configuration' (which is the current step), and '3 Done'. Below the progress bar, there is a form with a label 'Asset Name' and a text input field containing the word 'Random'. At the bottom of the form, there are two buttons: 'Previous' and 'Next'.

- Configure the plugin
 - **Asset name:** The name of the asset that will be created
- Click *Next*
- Enable the service and click on *Done*

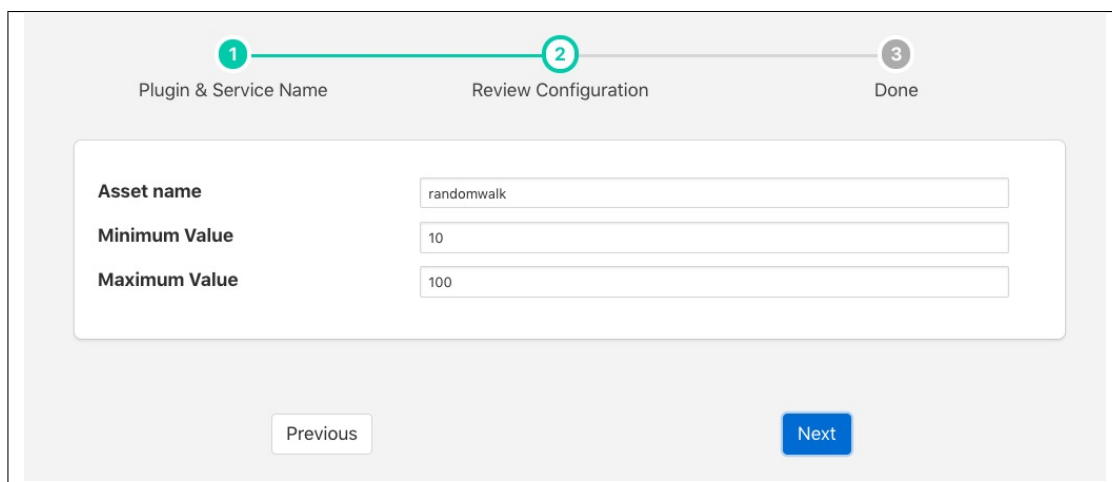
8.1.32 Random Walk

The *foglamp-south-randomwalk* plugin is a plugin that will create random data between a pair of values. Each new value is based on a random increment or decrement of the previous. This results in an output that appears as follows



To create a south service with the Random Walk plugin

- Click on *South* in the left hand menu bar
- Select *randomwalk* from the plugin list
- Name your service and click *Next*



1 Plugin & Service Name 2 Review Configuration 3 Done

Asset name randomwalk

Minimum Value 10

Maximum Value 100

Previous Next

- Configure the plugin
 - **Asset name:** The name of the asset that will be created
 - **Minimum Value:** The minimum value to include in the output
 - **Maximum Value:** The maximum value to include in the output
- Click *Next*
- Enable the service and click on *Done*

8.1.33 Enviro pHAT Plugin



The *foglamp-south-rpienviro* is a plugin that uses the Pimoroni Enviro pHAT sensor board. The Enviro pHAT board is an environmental sensing board populated with multiple sensors, the plugin pulls data from the;

- RGB light sensor
- Magnetometer
- Accelerometer
- Temperature/pressure Sensor

Individual sensors can be enabled or disabled separately in the configuration. Separate assets are created for each sensor within FogLAMP with individual controls over the naming of these assets.

Note: The Enviro pHAT plugin is only available on the Raspberry Pi as it is specific the GPIO pins of that device.

To create a south service with the Enviro pHAT

- Click on *South* in the left hand menu bar
- Select *rpienviro* from the plugin list
- Name your service and click *Next*

The screenshot shows the 'Review Configuration' step of a three-step wizard. The steps are labeled at the top: 1. Plugin & Service Name, 2. Review Configuration, and 3. Done. The configuration form contains the following fields and controls:

- Asset Name Prefix:** Text input field with the value 'e_'.
- RGB Sensor:** Checkmark control, which is checked.
- RGB Sensor Name:** Text input field with the value 'rgb'.
- Magnetometer Sensor:** Checkmark control, which is checked.
- Magnetometer Sensor Name:** Text input field with the value 'magnetometer'.
- Accelerometer Sensor:** Checkmark control, which is checked.
- Accelerometer Sensor Name:** Text input field with the value 'accelerometer'.
- Weather Sensor:** Checkmark control, which is checked.
- Weather Sensor Name:** Text input field with the value 'weather'.

At the bottom of the form, there are two buttons: 'Previous' (disabled) and 'Next' (active).

- Configure the plugin
 - **Asset Name Prefix:** An optional prefix to add to the asset names. The asset names created by the plugin are; rgb, magnetometer, accelerometer and weather. Using the prefix you can add an identifier to the front of each such that it becomes easier to differentiate between multiple instances of the sensor.
 - **RGB Sensor:** A toggle control to turn on or off collection of RGB light level information
 - **RGB Sensor Name:** Set a name for the RGB sensor asset
 - **Magnetometer Sensor:** A toggle control to turn on or off collection of magnetometer data
 - **Magnetometer Sensor Name:** Set a name for the magnetometer sensor asset
 - **Accelerometer Sensor:** A toggle to turn on or off collection of accelerometer data
 - **Accelerometer Sensor Name:** Set a name for the accelerometer sensor asset
 - **Weather Sensor:** A toggle to turn on or off collection of weather data
 - **Weather Sensor Name:** Set a name for the weather sensor asset
- Click *Next*
- Enable the service and click on *Done*

8.1.34 OPC/UA Safe & Secure South Plugin

The *foglamp-south-s2opcua* plugin allows FogLAMP to connect to an OPC/UA server and subscribe to changes in the objects within the OPC/UA server. This plugin is very similar to the *foglamp-south-opcua* plugin but is implemented using a different underlying OPC/UA open source library, from Systemrel. The major difference between the two is the ability of this plugin to support secure endpoints with the OPC/UA server.

A south service to collect OPC/UA data is created in the same way as any other south service in FogLAMP.

- Use the *South* option in the left hand menu bar to display a list of your South services
- Click on the + add icon at the top right of the page
- Select the *s2opcua* plugin from the list of plugins you are provided with
- Enter a name for your south service
- Click on *Next* to configure the OPC/UA plugin

1

2

3

Plugin & Service NameReview ConfigurationDone

Asset Name

s2opcua

OPCUA Server URL

opc.tcp://localhost:53530/OPCUA/SimulationServer

OPCUA Object Subscriptions

1 ▾

{

2 ▾

"subscriptions": [

3

"ns=3;i=1001",

4

"ns=3;i=1002"

5

]

6

}

Min Reporting Interval (millisec)

1000

Security Mode

None ▾

Security Policy


None ▾

User Authentication Policy

anonymous ▾

Username

Password

password 

CA Certificate Authority

Server Public Certificate

Client Public Certificate

Client Private Key

Certificate Revocation List

Debug Trace File

☐

Previous

Next

The configuration parameters that can be set on this page are;

- **Asset Name:** This is a prefix that will be applied to all assets that are created by this plugin. The OPC/UA plugin creates a separate asset for each data item read from the OPC/UA server. This is done since the OPC/UA server will deliver changes to individual data items only. Combining these into a complex asset would result in assets that do only contain one of many data points in each update. This can cause upstream systems problems with the every changing asset structure.
- **OPCUA Server URL:** This is the URL of the OPC/UA server from which data will be extracted. The URL should be of the form `opc.tcp://.../`
- **OPCUA Object Subscriptions:** The subscriptions are a set of locations in the OPC/UA object hierarchy that defined which data is subscribed to in the server and hence what assets get created within FogLAMP. A fuller description of how to configure subscriptions is shown below.
- **Min Reporting Interval:** This control the minimum interval between reports of data changes in subscriptions. It sets an upper limit to the rate that data will be ingested into the plugin and is expressed in milliseconds.

- **Security Mode:** Specify the OPC/UA security mode that will be used to communicate with the OPC/UA server.

- **Security Policy:** Specify the OPC/UA security policy that will be used to communicate with the OPC/UA server.

- **User Authentication Policy:** Specify the user authentication policy that will be used when authenticating the connection to the OPC/UA server.
- **Username:** Specify the username to use for authentication. This is only used if the *User authentication policy* is set to *username*.
- **Password:** Specify the password to use for authentication. This is only used if the *User authentication policy* is set to *username*.

- **CA Certificate Authority:** The name of the root certificate authorities certificate file in DER format. This is the certificate authority that forms the root of trust and signs the certificates that will be trusted. If using self-signed certificates this should be left blank.
- **Server Public Certificate:** The name of the public certificate of the OPC/UA server specified in the *OPCUA Server URL*. This must be a DER format certificate file. It must be signed by the certificate authority unless you are using self-signed certificates.
- **Client Public Certificate:** The name of the public certificate of the OPC/UA client application, that is, this plugin. This must be a DER format certificate file. It must be signed by the certificate authority unless you are using self-signed certificates.
- **Client Private Key:** The name of the private key of the client application, that is, the private key the plugin will use. This must be a PEM format key file.
- **Certificate Revocation List:** The name of the certificate authority's Certificate Revocation List. This is a DER format certificate. If using self-signed certificates this should be left blank.
- **Debug Trace File:** Enable the S2OPCUA OPCUA Toolkit trace file for debugging. If enabled, log files will appear in the directory */usr/local/foglamp/data/logs*.

Subscriptions

Subscriptions to OPC/UA objects are stored as a JSON object that contains an array named “subscriptions.” This array is a set of OPC/UA nodes that will control the subscription to variables in the OPC/UA server. Each element in the array is an OPC/UA node id, if that node is the id of a variable then that single variable will be added to the subscription list. If the node id is not a variable, then the plugin will recurse down the object tree below that node and add every variable it finds in this tree to the subscription list.

A subscription list which gives the root node of the OPC/UA server will cause all variables within the server to be added to the subscription list. Care however should be taken as this may be a large number of assets.

Subscription examples

```
{ "subscriptions": [ "5:Simulation", "2:MyLevel" ] }
```

We subscribe to

- 5:Simulation is a node name under ObjectsNode in namespace 5
- 2:MyLevel is a variable under ObjectsNode in namespace 2

```
{ "subscriptions": [ "5:Sinusoid1", "2:MyLevel", "5:Sawtooth1" ] }
```

We subscribe to

- 5:Sinusoid1 and 5:Sawtooth1 are variables under ObjectsNode/Simulation in namespace 5
- 2:MyLevel is a variable under ObjectsNode in namespace 2

```
{ "subscriptions": [ "2:Random.Double", "2:Random.Boolean" ] }
```

We subscribe to

- Random.Double and Random.Boolean are variables under ObjectsNode/Demo both in namespace 2

Object names, variable names and namespace indices can be easily retrieved browsing the given OPC/UA server using OPC UA clients, such as .

Certificate Management

OPC UA clients and servers use X509 certificates to confirm each other's identities and to enable digital signing and data encryption. Certificates are often issued by a Certificate Authority (CA) which means either the client or the server could reach out to the CA to confirm the validity of the certificate if it chooses to.

The configuration described above uses the names of certificates that will be used by the plugin. These certificates must be loaded into the FogLAMP Certificate Store manually and named to match the names used in the configuration before the plugin is started. When entering certificate and key file names, do not include directory names or file extensions (*.der* or *.pem*).

Typically the Certificate Authorities certificate is retrieved and uploaded to the FogLAMP Certificate Store along with the certificate from the OPC/UA server that has been signed by that Certificate Authority. A public/private key pair must also be created for the plugin and signed by the Certificate Authority. These are uploaded to the FogLAMP Certificate Store.

[OpenSSL](#) may be used to generate and convert the keys and certificates required. An to do this is available as part of the underlying library.

Certificate Requirements

Certificates must be X509 Version 3 certificates and must have the following field values:

Certificate Field	Value
Version	V3
Subject	This field must include a Common Name (CN=) which is a human-readable name such as <i>S2OPCUA South Plugin</i> . Do not use your device hostname.
Subject Alternative Name	URI= foglamp:south:s2opcua, DNS= <i>deviceHostname</i>
Key Usage	Digital Signature, Key Encipherment, Non Repudiation, Data Encipherment
Extended Key Usage	Client Authentication

Self-Signed Certificates

A common configuration is to use self-signed certificates which are issued by your own systems and cannot be validated against a CA. For this to work, the OPC UA client and server must each have a copy of the other's certificate in their Trusted Certificate stores. This task must be done by a system manager who is creating the device configuration. By copying certificates, the system manager is confirming that the client and server can legitimately communicate with each other.

Creating a Self-Signed Certificate

There is a very useful online tool for creating self-signed certificates called [CertificateTools](#). You can watch a demonstration of CertificateTools on [YouTube](#). This section will walk you through the necessary steps to create a self-signed certificate for the S2OPCUA South plugin which is the OPC UA Client.

The [CertificateTools](#) main page is divided into sections. You can leave many of the sections at their default values. Here are the required entries for each section:

Private Key

Leave the default values as-is: *Generate PKCS#8 RSA Private Key* and *2048 Bit*. Leave *Encrypt* unchecked.

Subject Attributes

In *Common Names*, enter a human-readable name such as *S2OPCUA South Plugin*. Click *Add*.

Edit *Country*, *State*, *Locality* and *Organization* as you wish. We recommend:

- Country: US
- State: CA
- Locality: Menlo Park
- Organization: Dianomic

Subject Alternative Name

Set the drop-down to *DNS*. Enter the hostname of your FogLAMP device. This can be an unqualified name, that is, the device hostname without domain name. Click *Add*.

Set the drop-down to *URI*. Enter *foglamp:south:s2opcua*. Click *Add*.

x509v3 Extensions

Key Usage

Click the check boxes to enable *Critical*, *Digital Signature*, *Key Encipherment*, *Non Repudiation* and *Data Encipherment*.

Extended Key Usage

Click the check boxes to enable *Critical* and *TLS Web Client Authentication*.

Encoding Options

Leave at Default.

CSR Options

Leave the first drop-down at *SHA256*. Change the second drop-down from *CSR Only* to *Self-Sign*. Doing this will expose drop-downs to set the self-signed certificate expiration time.

Generating the Certificate and Private Key

Click *Submit*. This will create a new section marked by a blue bar labelled *Certificate 0*.

Open *Certificate 0*. This will reveal a subsection called *Download*. You will need only two of these files:

- PEM Certificate (filename *cert.crt*)
- PKCS#12 Certificate and Key (filename *cert.pfx*)

When you click the *PKCS#12 Certificate and Key* link, you will be prompted for a password for the private key. It is acceptable to click *Cancel* to proceed without a password. Download these two files to a working directory on any computer with OpenSSL installed (you will need OpenSSL to post-process the downloaded files). You do not need to do this on your FogLAMP device. You must do this on a machine that can run the FogLAMP GUI in a browser; you will need the browser to import the certificate and key into the FogLAMP Certificate Store.

Note: The CertificateTools webpage can show you the equivalent OpenSSL commands to perform the self-signed certificate and key generation. Look for *OpenSSL Commands* below the blue *Certificate 0* bar.

Post-Processing the Certificate and Private Key

Use the OpenSSL command-line utility to convert the certificate and key files to the formats needed for the S2OPCUA South Plugin.

Converting the Certificate File

The *PEM Certificate* file (*cert.crt*) is in PEM format. It must be converted to DER format. The command is:

```
openssl x509 -inform pem -outform der -in cert.crt -out myclientcert.der
```

Converting the Private Key File

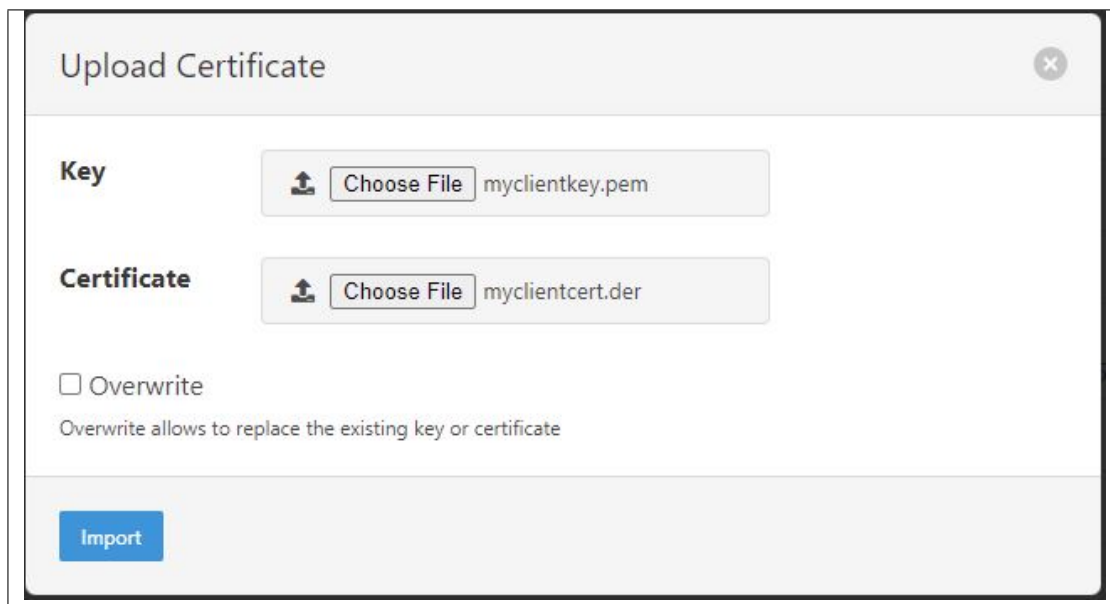
The *PKCS#12 Certificate and Key* file (*cert.pfx*) is in Public-Key Cryptography Standards [PKCS#12](#) format. It must be converted to PEM format. The command is:

```
openssl pkcs12 -in cert.pfx -out myclientkey.pem -nodes
```


This command will prompt for the Import Password. If you created a password when you downloaded the PKCS#12 Certificate and Key file, enter it now. If you did not create a password, hit Enter.


Importing the Certificate and Key Files

Launch the FogLAMP GUI. Navigate to the Certificate Store. In the upper right corner of the screen, click *Import*.



Upload Certificate

Key  Choose File myclientkey.pem

Certificate  Choose File myclientcert.der

☐ Overwrite
Overwrite allows to replace the existing key or certificate

Import

In the *Key* section, click *Choose File* and navigate to the location of the key file *myclientkey.pem*.

In the *Certificate* section, click *Choose File* and navigate to the location of the certificate file *myclientcert.der*.

Click *Import*.

You should use the Certificate Store in the FogLAMP GUI to import your OPC UA server certificate. In this case, enter the server certificate file name in the *Certificate* portion of the Import dialog and then click *Import*.

8.1.35 Siemens S7 PLC



The *foglamp-south-s7* plugin is a south plugin that reads data from a Siemens S7 PLC using the S7 communication protocol. Data can be read from a number of sources within the PLC

- Data blocks - The data blocks store the state of the PLC

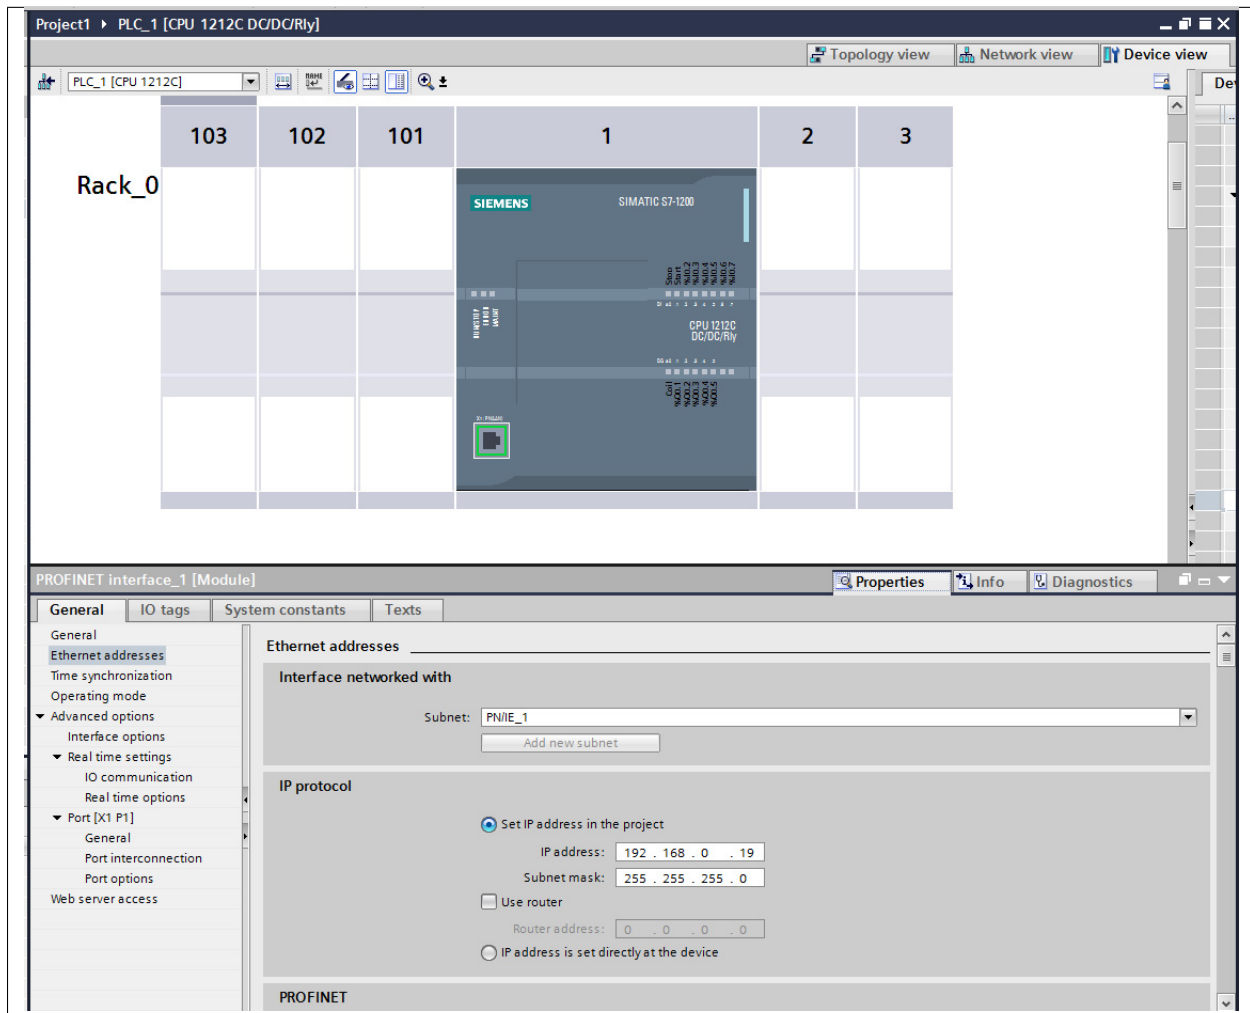
- Inputs - Read the state of the inputs to the PLC
- Outputs - Read the state of the outputs from the PLC
- Merkers - Read from the single bit flag store
- Counters - Read a counter
- Timers - Read a timer

Configuring the PLC

There are a number of configuration steps that must be taken on the PLC itself to support the use of the S7 protocol.

Assigning an IP Address

Using the Siemens TIA console assign an IP address to your PLC. Connect to your PLC and locate the display of the PLC device. Double click on the network connector to bring up the properties for the network interface.




Assign an IP address to your interface and if you require it you may also assign a router to use.

Enable PUT/GET operations

The S7 1200 and 1500 series PLC's require the PUT/GET communication from partners to be enabled in order to retrieve data using the S7 protocol. To permit the PUT/GET network operations on your PLC use the Siemens VIA tool. Note you must be sure that you are offline when you do this. Locate you PLC in the tool and right click on the device select properties and the following dialog will be displayed.

The older S7-300 and S7-400 series do not require this to be done.



Select the protection tab and scroll down to find the checkbox that enables the use of GET/PUT operations. Make sure it is selected for your PLC.

Using the Plugin

To create a south service with the Siemens S7 plugin

- Click on *South* in the left hand menu bar
- Select *S7* from the plugin list
- Name your service and click *Next*

1 Plugin & Service Name 2 Review Configuration 3 Done

Default Asset Name:

PLC IP Address:

Rack:

Slot:

Map:

```

1 {
2   "items": [
3     {
4       "datapoint": "dbl",
5       "area": "DB",
6       "DBnumber": 1,
7       "start": 1,
8       "type": "byte"
9     }
10  ]
11 }

```

Previous Next

- Configure the plugin
 - **Default Asset Name:** The name of the asset to use if none is given in each of the data mapping items.
 - **PLC IP Address:** The IP address assigned to your PLC.
 - **Rack:** The rack number to address, usually this is 0 for a standalone PLC.
 - **Slot:** The slot within the rack, most CPU's are in slot 1 of the rack.
 - **Map:** The data mapping for the plugin. This tells the plugin what data to fetch from the PLC.

Map Format

The data mapping uses a JSON document to define the data that should be read. The document is an array of items to read from the PLC, each item is a datapoint within either the default asset or it may be defined as a different asset within the item. An item contains a number of properties

- **asset:** An optional property that can be used to put this item into an asset other than than one defined as the default asset for the service.
- **datapoint:** The name of the data point that the data will be placed in. All items must have a datapoint defined.
- **area:** The area in the PLC that data will be read from. There are a number of areas available
 - **PE:** Process Input - these are the inputs to the PLC
 - **PA:** Process Output - these are the outputs from the PLC
 - **MK:** Merker - a single bit memory used to store flags
 - **DB:** Data Block - the data blocks within the PLC used to store state within the PLC code
 - **CT:** Counter - The counters within the PLC

- **TM:** Timer - The timers within the PLC

You may use the abbreviated area name, e.g. *PA* or the longer name *Process Inputs* interchangeably in the map.

- **DBnumber:** The data block number, this is only required for data blocks and is used to define the block to read.
- **start:** The offset of the start of the item within the data block
- **type:** The type of the data item to read. A number of different types are supported
 - **bit:** A single bit value, mostly used to retrieve the state of a digital input to the PLC
 - **byte:** A 8 bit integer value.
 - **word:** A 16 bit integer value.
 - **dword:** A 32 bit integer value.
 - **real:** A 32 bit floating point value.
 - **counter:** A 16 bit counter.
 - **timer:** a 16 bit timer.

A simple data mapping that wanted to read the state of two digital inputs to the PLC, say DI0 and DI2, and wanted to labeled these as datapoints “Stop” and “Start” within the default asset would consist of two items as follows

```
{
  "items" : [
    {
      "datapoint": "Stop",
      "area": "PE",
      "start": 0,
      "type": "bit"
    },
    {
      "datapoint": "Start",
      "area": "PE",
      "start": 2,
      "type": "bit"
    }
  ]
}
```

In this case we set start to 0 for DI0 as it is the first digital input in the set. DI2 has a start of 2 as it is the second input. We use the type of *bit* to return a simple 0 or 1 to indicate the state of the input. We could use *byte* instead, this would return the 8 inputs states encoded as a binary number.

```
{
  "datapoint": "Inputs",
  "area": "PE",
  "start": 0,
  "type": "byte"
}
```

Since *start* is set to 0 and *type* is byte, then we return the state of the 8 inputs. We can do the same thing using the longer name form of the area as follows.

```
{
  "datapoint": "Inputs",
  "area": "Process Inputs",
  "start": 0,
  "type": "byte"
}
```

To add in a digital output, say DO4 and label that running, we would add another item to the map

```
{
  "datapoint": "Running",
  "area": "PA",
  "start": 4,
  "type": "bit"
}
```

If we assume we have a data block that we wish to read data from that appears as follows

	Name	Data type	Offset	Start value	Retain	Accessible f...	Writa...	Visible in ...	Setpoint	Comment
1	Static	Int	0.0	0	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
2	count	Int	2.0	1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
3	state	Int	4.0	3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
4	failures	DWord	8.0	1.3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
5	rate	Bool	12.0	false	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
6	running	Time	14.0	T#575_656MS	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
7	downtime	Time	14.0	T#575_656MS	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	

Then we can setup a number of items in the map to retrieve these values and place them in data points. The items that would read this data block would be

```
{
  "datapoint": "count",
  "area": "DB",
  "DBnumber": 1,
  "start": 0,
  "type": "word"
},
{
  "datapoint": "state",
  "area": "DB",
  "DBnumber": 1,
  "start": 2,
  "type": "word"
},
{
  "datapoint": "failures",
  "area": "DB",
  "DBnumber": 1,
  "start": 4,
  "type": "dword"
},
{
  "datapoint": "rate",
  "area": "DB",
```

(continues on next page)

(continued from previous page)

```

        "DBnumber" : 1,
        "start": 8,
        "type": "word"
    },
    {
        "datapoint": "running",
        "area": "DB",
        "DBnumber" : 1,
        "start": 12,
        "type": "word"
    },
    {
        "datapoint": "downtime",
        "area": "DB",
        "DBnumber" : 1,
        "start": 14,
        "type": "timer"
    }
}

```

For clarity we have used the name in the data block as the datapoint name, but these need not be the same.

If there is an error in the map definition for a given item then that item is ignored and a message is written to the error log. For example if a bad area name is given

```

Jun 25 08:53:04 foglamp-18 FogLAMP S7[6121]: ERROR: Invalid area Data specified in ↵
↵device mapping for S7 db1-bad
Jun 25 08:53:04 foglamp-18 FogLAMP S7[6121]: ERROR: Discarded invalid item in map for ↵
↵datapoint db1-bad

```

If a Data Block is missing it's DBnumber property then the following style of error will be produced.

```

Jun 25 08:39:07 foglamp-18 FogLAMP S7[6121]: ERROR: Missing data block number in map ↵
↵for S7, db1-bad. A data block number must be specified for a data block area read.
Jun 25 08:39:07 foglamp-18 FogLAMP S7[6121]: ERROR: Discarded invalid item in map for ↵
↵datapoint db1-bad

```

Other errors that can occur include

```

Jun 25 08:57:28 foglamp-18 FogLAMP S7[6121]: ERROR: Missing start in map for ↵
↵datapoint db1-bad
Jun 25 08:57:46 foglamp-18 FogLAMP S7[6121]: ERROR: Missing type in map for datapoint ↵
↵db1-bad

```

8.1.36 Samotics4 South Plugin

The *foglamp-south-samotics4* plugin is a south plugin that pulls data from the SAM4 Data Integration API of .

This is based on Samotics API

The data retrieved via this API includes

- Motors metadata
- Motors metrics
- Motor incidents

Configuration Parameters

A Samotics4 south service is added in the same way as any other south service in FogLAMP,

- Select the *South* menu item
- Click on the + icon in the top right

You will be presented with the following page

- Select *samotics4* from the plugin list
- Enter a name for your Samotics4 service
- Click *Next*
- You will be presented with the following configuration page

- **Asset Name Prefix:** This is the prefix of the assets that will be added for the data read by this service. The default value is *sam4_*.
- **OAuth2 URL:** This is the base URL for OAuth2 authentication.
- **OAuth2 client id:** This is OAuth2 client id needed for authentication.
- **OAuth2 secret:** This is OAuth2 secret needed for authentication.
- **Motor metadata in metrics:** This parameter defines whether to store motors metadata along with motor metrics data or to create new assets for motor metadata. The default value is true.

8.1.37 SenseHAT



The *foglamp-south-sensehat* is a plugin that uses the Raspberry Pi Sense HAT sensor board. The Sense HAT has an 8x8 RGB LED matrix, a five-button joystick and includes the following sensors:

- Gyroscope
- Accelerometer
- Magnetometer
- Temperature
- Barometric pressure
- Humidity

In addition it has an 8x8 matrix for RGB LED's, these are not included in the devices the plugin supports.

Individual sensors can be enabled or disabled separately in the configuration. Separate assets are created for each sensor within FogLAMP with individual controls over the naming of these assets.

Note: The Sense HAT plugin is only available on the Raspberry Pi as it is specific the GPIO pins of that device.

To create a south service with the Sense HAT

- Click on *South* in the left hand menu bar
- Select *sensehat* from the plugin list
- Name your service and click *Next*

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name Prefix	sensehat/
Pressure Sensor	<input checked="" type="checkbox"/>
Pressure Sensor Name	pressure
Temperature Sensor	<input checked="" type="checkbox"/>
Temperature Sensor Name	temperature
Humidity Sensor	<input checked="" type="checkbox"/>
Humidity Sensor Name	humidity
Gyroscope Sensor	<input checked="" type="checkbox"/>
Gyroscope Sensor Name	gyroscope
Accelerometer Sensor	<input checked="" type="checkbox"/>
Accelerometer Sensor Name	accelerometer
Magnetometer Sensor	<input checked="" type="checkbox"/>
Magnetometer Sensor Name	magnetometer
Joystick Sensor	<input checked="" type="checkbox"/>
Joystick Sensor Name	joystick

Previous Next

- Configure the plugin
 - **Asset Name Prefix:** An optional prefix to add to the asset names.
 - **Pressure Sensor:** A toggle control to turn on or off collection of pressure information
 - **Pressure Sensor Name:** Set a name for the Pressure sensor asset
 - **Temperature Sensor:** A toggle control to turn on or off collection of temperature information
 - **Temperature Sensor Name:** Set a name for the temperature sensor asset
 - **Humidity Sensor:** A toggle control to turn on or off collection of humidity information
 - **Humidity Sensor Name:** Set a name for the humidity sensor asset
 - **Gyroscope Sensor:** A toggle control to turn on or off collection of gyroscope information
 - **Gyroscope Sensor Name:** Set a name for the gyroscope sensor asset
 - **Accelerometer Sensor:** A toggle to turn on or off collection of accelerometer data
 - **Accelerometer Sensor Name:** Set a name for the accelerometer sensor asset
 - **Magnetometer Sensor:** A toggle control to turn on or off collection of magnetometer data
 - **Magnetometer Sensor Name:** Set a name for the magnetometer sensor asset
 - **Joystick Sensor:** A toggle control to turn on or off collection of joystick data
 - **Joystick Sensor Name:** Set a name for the joystick sensor asset
- Click *Next*
- Enable the service and click on *Done*

8.1.38 Simple REST with Payload Scripting

The *foglamp-south-simple-rest* plugin uses REST calls to receive API responses from sensors or other sources. The plugin make HTTP or HTTPS GET requests to retrieve API responses, HTTP header fields can also be added via the plugin configuration. It then uses an optional script, written in Python, that converts the message into a JSON document and pushes data to the FogLAMP System. However it also has a set of built in rules for interpreting some common payload formats which enable it to be used without providing a script in a large number of common cases.

When a script is provided it is best practice to do the minimum required to allow the data to be ingested into the FogLAMP data pipeline. Further processing to shape the data to exact requirements can often be done using an existing filter. The advantages of this are twofold; it simplifies the scripts required here and it simplifies maintenance should the data be required in a different format some time later.

Configuration

When adding a south service with this plugin the same flow is used as with any other south service. The configuration page for the plugin is as follows.

The screenshot shows a configuration form for a REST API asset. The form is organized into several sections:

- Asset Name:** A text input field containing "rest".
- URL:** A text input field containing "http://server/location".
- Headers:** A table with one row containing the JSON object "{}".
- Selection Method:** A dropdown menu set to "None".
- ID Parameter:** An empty text input field.
- Initial ID:** An empty text input field.
- ID Field:** An empty text input field.
- Start:** An empty text input field.
- End:** An empty text input field.
- Timestamp:** An empty text input field.
- Time Format:** An empty text input field.
- Timezone:** A text input field containing "+00:00".
- Collapse:** A checkbox that is checked.
- Asset Field:** An empty text input field.
- Script:** A large text area for entering a script.

At the bottom of the form, there is a "Choose files" button and a "No file chosen" message.

- **Asset Name:** The name of the asset the plugin will create for each message, unless the convert function returns an explicit asset name to be used.
- **URL:** The URL of the REST API to be called. This should be a complete URL, including the http or https protocol to use.
- **Headers:** An optional set of headers to include in the REST API call. The headers are encoded as a JSON document as a set of name/value pairs within a JSON object.
- **Selection Method:** The plugin supports a number of methods for selecting the data should be returned. The choices are return all the data, return data based on an ID or return data based on time. See [Selection Method](#) for more details.



The image shows a web form with three input fields: "Selection Method", "ID Parameter", and "Initial ID". A dropdown menu is open for the "Selection Method" field, showing three options: "None" (selected with a checkmark), "ID Based", and "Time Based".

- **ID Parameter:** An optional URL query parameter to add to each call to the URL. This is expected to be a numeric value that gets passed to the API and is used for implementing ID passing to calls. This is only valid if the selection method *ID Based* has been chosen.
- **Initial ID:** The initial value to pass for the query parameter. This may be used on the first call only if the *ID Based* selection method is chosen.
- **ID Field:** This defines a data field that is ingested that will be used for second and subsequent calls to the API as the new value of ID. This is only used with a selection method of *ID Based*.
- **Start:** The name of the query parameter to add to the URL to indicate the start time if a selection method of *Time Based* has been chosen.
- **End:** The name of the query parameter to add to the URL to indicate the end time if a selection method of *Time Based* has been chosen.
- **Time Format:** The format to both pass the timestamps into the query parameters using and also to interpret the timestamps returned in the payload.
- **Timezone:** The timezone to use for the start and end times that are sent in the API request and also when timestamps are read from the API response. Timezone is expressed as an offset in hours and minutes from UTC for the local timezone of the API. E.g. -08:00 for PST timezones.
- **Collapse:** Collapse the returned returning to a flat structure, if not enabled a nested reading will be produced.
- **Timestamp:** The name of the item in the response payload that should be treated as the timestamp for the reading.
- **Asset Field:** The name of a field in the response payload that should be treated as the asset name to use for the reading. If this is left empty or the data does not contain a field with this name then the default asset name configured in the *Asset Name* configuration item will be used.
- **Script:** The Python script to execute for message processing. Initially a file must be uploaded, however once uploaded the user may edit the script in the box provided. A script is optional.

Selection Method

The plugin supports two methods to select data to be retrieved from the API that is called, these methods are designed for use with an API that is maintaining historic data and provided a mechanism to present the same historic data being read multiple times. If your API does not store historic data then you may select the method *None* to simply retrieve all the data available via the API.

ID Based

The select mechanism *ID Based* is designed for API that give each value some form of ID that increases over time. When a call is made you pass the value of the ID for the next data item you wish to read. This method is used in conjunction with 3 other parameters. These parameters are used to control the name of the query parameter to add to the URL, *ID Parameter*. This name will be used to pass in the ID to be read and is added to the URL that is configured. In the first call using the *ID Based* method the value of *Initial ID* will be passed as the value. In subsequent calls the maximum value of the data field name as per the *ID Field* configuration parameter will be used as the value of the ID parameter.

ID Parameter is the name of the parameter that is passed in the requests, it is appended to the configured URL along with the current value for the parameter.

For example if the URL is configured as `http://api-server.com/api/v1/data?user=dianomic` and the *ID Parameter* is defined as *requestID* with the *Initial ID* of 100, then the full URL that is used in the call will be

```
http://api-server.com/api/v1/data?user=dianomic&requestID=100
```

The URL used in the next call will be dependent on the setting of *ID Field*. If it is left empty then the value last used will be incremented for the next call, provided the previous call was successful. In our above example this would result in the next call using the URL

```
http://api-server.com/api/v1/data?user=dianomic&requestID=101
```

If the first call had failed, then the next call would use the same value for our parameter.

A more common case is when the data returned contains ID values for each returned value, in this case the *ID Field* configuration option is set and the values taken from the response will generate the next ID to use. For example, if the response payload returns sets of readings, each identified by a field called *id*, then set *ID Field* to *id*. A response payload that returned *id*'s 125, 126 & 127 would then cause the next request to send a value for the parameter of 128.

```
http://api-server.com/api/v1/data?user=dianomic&requestID=128
```

Time Based

The selection mechanism *Time Based* is designed for an API that returns values for a time window. It requires two parameters to be passed in the request, *Start* and *End*, to specify the time window to the server. The format of the timestamps passed to the server are defined by the *Time Format* configuration parameter.

% The % character.

%a or %A The name of the day of the week according to the current locale, in abbreviated form or the full name.

%b or %B or %h

The month name according to the current locale, in abbreviated form or the full name.

%c The date and time representation for the current locale.

%C The century number (0–99).

%d or %e The day of month (1–31).

%D Equivalent to `%m/%d/%y`. (This is the American style date, very confusing to non- Americans, especially since `%d/%m/%y` is widely used in Europe. The ISO 8601 standard format is `%Y-%m-%d`.)

%H The hour (0–23).

- %I** The hour on a 12-hour clock (1–12).
- %j** The day number in the year (1–366).
- %m** The month number (1–12).
- %M** The minute (0–59).
- %n** Arbitrary white space.
- %p** The locale’s equivalent of AM or PM. (Note: there may be none.)
- %r** The 12-hour clock time (using the locale’s AM or PM). In the POSIX locale equivalent to **%I:%M:%S %p**. If **t_fmt_ampm** is empty in the **LC_TIME** part of the current locale, then the behavior is undefined.
- %R** Equivalent to **%H:%M**.
- %S** The second (0–60; 60 may occur for leap seconds; earlier also 61 was allowed).
- %t** Arbitrary white space.
- %T** Equivalent to **%H:%M:%S**.
- %U** The week number with Sunday the first day of the week (0–53). The first Sunday of January is the first day of week 1.
- %w** The ordinal number of the day of the week (0–6), with Sunday = 0.
- %W** The week number with Monday the first day of the week (0–53). The first Monday of January is the first day of week 1.
- %x** The date, using the locale’s date format.
- %X** The time, using the locale’s time format.
- %y** The year within century (0–99). When a century is not otherwise specified, values in the range 69–99 refer to years in the twentieth century (1969–1999); values in the range 00–68 refer to years in the twenty-first century (2000–2068).
- %Y** The year, including century (for example, 1991).

When using the *Time Based* selection mechanism two parameters may be appended. For example if the *URL* is configured as *http://api-server.com/api/v1/data?user=dianomic*, the configuration option *Start* is defined as *startTime* and *End* as *endTime*, with the *Time Format* set to be *%Y-%m-%dT%H:%M:%s*, then the full URL that is used in the call will be

```
http://api-server.com/api/v1/data?user=dianomic&startTime=2021-07-11T15:12:34&
↪endTime=2021-07-12T12:45:12
```

If this call succeeds then the next call will use the *endTime* from this call as the *startTime* for the next call. The *endTime* is always the current time.

Request URL Handling

The plugin makes HTTP (or HTTPS) GET requests to the configured URL, this may include parameter passing. Parameters used be encoded within the URL of in the plugin configuration, however if a *Selection Method* other than *None* is selected extra parameters will be added to the request URL.

Response Payload Handling

If the payload of the REST response is a JSON document with simple key/value pairs, e.g.

```
{ "temperature" : 23.1, "humidity" : 47.2 }
```

Then no translation script is required, each key/value pair will become a data point within an asset whose name is set in the configuration of the plugin. A working example of this is the /foglamp/ping API call of FogLAMP itself, it produces a response,

```
{
  "uptime": 27,
  "dataRead": 1063459,
  "dataSent": 617310,
  "dataPurged": 1063024,
  "authenticationOptional": true,
  "serviceName": "FogLAMP",
  "hostName": "foglamp-18",
  "ipAddresses": [
    "192.168.0.173"
  ],
  "health": "green",
  "safeMode": false,
  "version": "1.9.1"
}
```

This results in an asset which has data points for all the string, numeric and boolean items with the response. In this case the *ipAddresses* item is ignored as FogLAMP does not currently support string array type data.

If the response payload included nested JSON objects then these will be included also. For example if the response payload was

```
{
  "motor" : {
    "speed" : 12345,
    "current" : 1.4
  },
  "gearbox" : {
    "ratio" : 64,
  }
}
```

The three values would be extracted, *speed*, *current* and *ratio*. How these values are represented will depend on the setting of the *Collapse Data* configuration option. If this is set to true then a flat reading will be created for each of the three values. If it is set to false then a reading with two objects will be created, one for the motor and one for the gearbox. Within these they will contain the values for the appropriate object. The choice of flattening the data will depend on how the user wishes to use this data upstream within FogLAMP.

If the payload is a JSON document that is an array rather than an object, then it is interpreted as a set of readings. For example a payload that is an array of numbers such as

```
[
    12.4, 15.8, 18.2
]
```

Will result in a set of readings, one per value being created. The asset name and data point name will be taken for the *Asset* configuration option. This same rule applies for arrays of integers, floating point numbers, string or booleans.

The payload may also be an array of objects, in which case each object will be an asset, with the members of the object becoming the data points. These objects may be nested in which case they will follow the same rules as the motor and gearbox example above.

```
[
  {
    "motor" : {
      "speed" : 12345,
      "current" : 1.4
    },
    "gearbox" : {
      "ratio" : 64,
    }
  },
  {
    "motor" : {
      "speed" : 12345,
      "current" : 1.4
    },
    "gearbox" : {
      "ratio" : 64,
    }
  },
]
```

This will create two assets with the name of the asset as the *Asset* configuration option.

The default asset naming can be overridden by setting a value for the configuration item *Asset Field*. This can be used to extract a value from the data returned by the API as the name to use for the resultant asset. If the *Asset Field* configuration item is set to *machine* and the payload returned by the API calls is as follows.

```
[
  {
    "machine" : "CNC14698",
    "motor" : {
      "speed" : 12345,
      "current" : 1.4
    },
    "gearbox" : {
      "ratio" : 64,
    }
  },
  {
    "machine" : "CNC15217",
    "motor" : {
      "speed" : 12345,
      "current" : 1.4
    },
    "gearbox" : {
      "ratio" : 64,
    }
  }
]
```

(continues on next page)

(continued from previous page)

```
    },
]
```

The result would be two assets called *CNC14698* and *CNC15217*.

Also if the payload is a simple numeric value the plugin will accept this and create an asset with the data point name matching the topic on which the value was given in the payload.

If the message format is not a JSON document that can be parsed using the built in rules or is in some other format then a Python script should be provided that turns the message into a JSON format.

An example script, assuming the payload in the message is simply a value, might be as follows

```
def convert(message):
    return {
        'temperature' : float(message)
    }
```

Note that the message is passed as a string and the data we wish to ingest into FogLAMP in this case is assumed to be a floating point value. The example above of course is unnecessary as the plugin can consume this data without the need of a script.

The script could return either one or two values.

The script should return the JSON document as a Python DICT in the case of a single value.

The script should return a string and a JSON document as a Python DICT in the case of two values, the first of these values is the name of the asset to use and overrides the default asset naming defined in the plugin configuration.

First case sample:

```
def convert(message):
    return {'temperature_1': 10.2}
```

Second case sample:

```
def convert(message):
    return "ExternalTEMP", {'temperature_3': 11.3}
```

A single API call can return reading data for multiple assets. In this case the script can return a more complex JSON document that contains both the asset name and the data points for that asset. The return document should return a single JSON object called *readings* and within that a set of readings, one per asset, expressed as a number of reading objects. The key of each of these objects becomes the asset name and the value of each is the data points within the asset.

As an example, if a single API call gives us back both data on a motor and on a machine tool, we can process that API response into two distinct assets; *motor* and *tool*, each with its own set of data points. In this case the *rpm* and *current* of the motor and the *temperature* and *coolant* flow rate for the machine tool.

```
{
  "readings" :
  {
    "motor" : {
      "rpm"      : 8450,
      "current"  : 1.3
    },
    "tool" : {
      "temperature" : 32.1,

```

(continues on next page)

(continued from previous page)

```
        "coolant" : 147
    }
}
```

If a script is returning this more complex JSON object it should not return an asset name, if it does return an asset name with this JSON format then the asset name will be ignored.

Timestamp Treatment

The default timestamp for a reading collected via this plugin will be the time at which the reading was taken, however it is possible for the API that is being called to include a different timestamp.

Returning a data point called whose name is defined in the *Timestamp* configuration option will result in the value of that data point being used as the timestamp. This data point will not be added to the reading. The default name of the timestamp is *timestamp*.

The timestamp data point should be a string and the timestamp should be formatted to match the definition given in the *Time format* configuration parameter. The format is based on the standard Linux *strptime* formatting options and is discussed above in the section discussing the *Time Based* selection method. It should be noted however that this timestamp handling in the payload is independent of the selection method chosen.

The timezone may be set by using the *Timezone* configuration parameter to set the offset of the timezone in which the API is running.

The plugin will automatically filter out a second or subsequent readings that have the same timestamp value as previous reading for that same asset. This allows an API which returns the timestamp of the data to be called multiple times and the data will only be ingested once for the given timestamp. The result is the polling rate of the south service can be set independently of the rate the data changes.

Script Error Handling

If an error occurs in the plugin or Python script, including script coding errors and Python exception, details will be logged to the error log and data will not flow through the pipeline to the next filter or into the storage service.


Warnings raised will also be logged to the error log but will not cause data to cease flowing through the pipeline.

To view the error log you may examine the file directly on your host machine, for example */var/log/syslog* on a Ubuntu host, however it is also possible to view the error logs specific to Fledge from the Fledge user interface. Select the *System* option under *Logs* in the left hand menu pane. You may then filter the logs for a specific service to see only those logs that refer to the service which uses the filter you are interested in.

The screenshot shows the FogLAMP System Logs interface. On the left is a sidebar with navigation items: Dashboard, Assets & Readings, South, North, Notifications, Control Dispatcher, Configuration, Schedules, Certificate Store, Backup & Restore, Logs, Audit, Notifications, Packages, System (highlighted), Tasks, Support, Settings, and Help. The main area is titled 'System Logs' and includes an 'Auto Refresh' checkbox. Below this are filters for 'Service' (set to 'Simple') and 'Severity' (set to 'Info and above'), along with a search bar. The log entries are as follows:

Timestamp	Host	Service	Severity	Message
Jul 7 17:06:48	foglamp-18	FogLAMP Simple[22518]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:06:38	foglamp-18	FogLAMP Simple[22518]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:06:37	foglamp-18	FogLAMP Simple[22518]	ERROR	Service is not registered with the core - not registering configuration interest
Jul 7 17:06:37	foglamp-18	FogLAMP Simple[22518]	ERROR	Failed to register service Simple
Jul 7 17:06:37	foglamp-18	FogLAMP Simple[22518]	ERROR	HTTP error during service registration: 400: A Service with the same name already exists
Jul 7 17:06:37	foglamp-18	FogLAMP Simple[22518]	ERROR	Python error in script: in delete message#012 at line 2
Jul 7 17:06:37	foglamp-18	FogLAMP Simple[22518]	ERROR	Failed to import script simple_script_syntax
Jul 7 17:06:18	foglamp-18	FogLAMP Simple[22140]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:06:03	foglamp-18	FogLAMP Simple[22140]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:05:48	foglamp-18	FogLAMP Simple[22140]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:05:44	foglamp-18	FogLAMP Simple[22140]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:05:43	foglamp-18	FogLAMP Simple[22140]	ERROR	Python error in script: in delete message#012 at line 2
Jul 7 17:05:43	foglamp-18	FogLAMP Simple[22140]	ERROR	Failed to import script simple_script_syntax
Jul 7 17:05:31	foglamp-18	FogLAMP Simple[21369]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:05:30	foglamp-18	FogLAMP Simple[21369]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:05:25	foglamp-18	FogLAMP Simple[21369]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:05:20	foglamp-18	FogLAMP Simple[21369]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:05:19	foglamp-18	FogLAMP Simple[21369]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:05:18	foglamp-18	FogLAMP Simple[21369]	ERROR	The supplied Python script does not define a valid "convert" function
Jul 7 17:05:17	foglamp-18	FogLAMP Simple[21369]	ERROR	The supplied Python script does not define a valid "convert" function

Alternatively if you open the dialog for the service in the *South* or *North* menu items you will see two icons displayed in the bottom left corner of the dialog that lets you alter the configuration of the service.



The left most icon, with the ? in a circle, allows you to view the documentation for the plugin, the right most icon, which looks like a page of text with a corner folded over, will open the log view page filtered to view the service.

Error Messages & Warnings

The following are some errors you may see within the log with some description of the cause and remedy for the error.

The supplied Python script does not define a valid “convert” function The script that has been supplied does not define a Python function called convert. The script must provide a single function called convert that accepts the HTTP response payload and will process that to provide the JSON DICT and an optional asset name to import.

Python error: IndentationError ‘expected an indented block’ in XXXX at line Y of script The script supplied does not conform to Python requirements for code block indentation. The text XXXX will be replaced with the line of text in error and Y with the line number within the script.

Python error: SyntaxError ‘invalid syntax’ in XXXX at line Y of script The script supplied does has invalid Python syntax. The text XXXX will be replaced with the line of text in error and Y with the line number within the script.

Python error: ModuleNotFoundError “No module named ‘nosuchpackage’” in supplied script The script supplied is attempting to import a Python module that is not available.

Python error: TypeError “convert() missing 1 required positional argument: ‘name’” in supplied script The type of the convert function has been incorrectly defined. The convert function should take a single argument which is the message to process.


Return from Python convert function is of an incorrect type, it should be a Python DICT object or a DICT object and a string

The convert function is returning data of an incorrect type. It may either return a Python DICT, which may be empty, None or a string and a Python DICT.

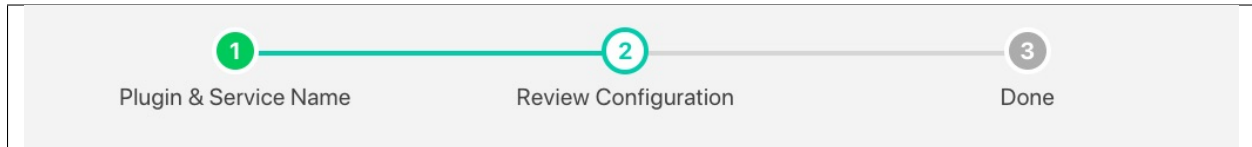
The plugin is unable to process data without a valid ‘convert’ function in the script. This warning will periodically be logged following an earlier error that has resulted in an error which prevents the Python convert function from processing the messages. Fix the earlier error to stop this warning being logged.

8.1.39 Sinusoid

The *foglamp-south-sinusoid* plugin is a south plugin that is primarily designed for testing purposes. It produces as it's output a simple sine wave, the period of which can be adjusted by changing the poll rate in the advanced settings of the south service into which it is loaded.



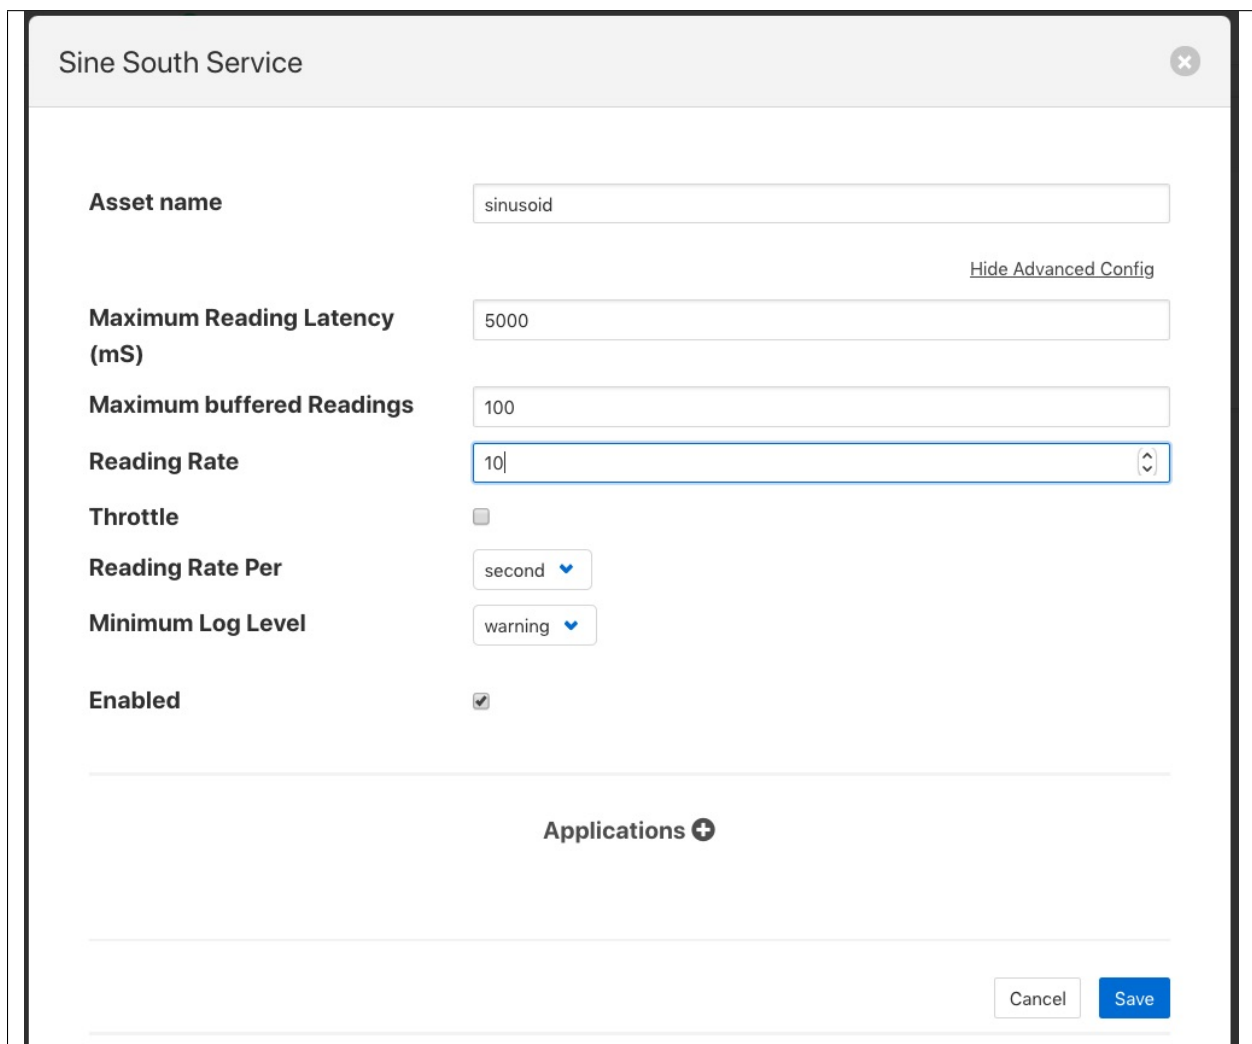
There is very little configuration required for the *sinusoid* plugin, merely the name of the asset that should be written. This can be useful if you wish to have multiple sinusoid in your FogLAMP system.



1 Plugin & Service Name 2 Review Configuration 3 Done

Asset name

The frequency of the sinusoid can be adjusted by changing the poll rate of the sinusoid plugin. To do this select the *South* item from the left-hand menu bar and then click on the name of your sinusoid service. You will see a link labeled *Show Advanced Config*, click on this to reveal the advanced configuration.



Sine South Service ✕

Asset name [Hide Advanced Config](#)

Maximum Reading Latency (mS)

Maximum buffered Readings

Reading Rate

Throttle ☐

Reading Rate Per

Minimum Log Level

Enabled ☒

Applications +

Amongst the advanced setting you will see one labeled *Reading Rate*. This defaults to 1 per second. The sinusoid takes 60 samples to complete one cycle of the sine wave, therefore it has a periodicity of 1 minute, or 0.0166Hz. If the *Reading Rate* is set to 60, then the frequency of the output becomes 1Hz.

8.1.40 Radiometric Data Capture for FLIR cameras



The *foglamp-south-spinnaker* plugin is a south plugin that uses the FLIR Spinnaker library to access the interface to cameras in order to retrieve radiometric or image data from compliant cameras.

The plugin can support both thermal cameras and regular image cameras that support the interface. It is however optimized for FLIR thermal cameras and designed to return radiometric data and other scalar values that are used to convert that radiometric data to other formats.

If used with a visual camera the regular image is returned in the *radiometric* data point.

Configuration

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name

Cameras

Frame Rate

Trigger

Trigger Edge

Output

Normalise ☐

Custom Configuration

1	
---	--

Focus Method

Auto Focus Method

Focus Direction

Focus Distance

Auto Focus Area

Focus Speed

- **Asset Name:** The name of the asset that will be used to store the data produced by the service.
- **Camera:** A drop down list of all the cameras that have been discovered on the network. Use this to select the camera to ingest data from.
- **Frame Rate:** Only used when continuous capture of frames is enabled. This defines how many frames per second will be captured.
- **Trigger:** The trigger source that will enable capture.

Trigger

Trigger Edge

✓ Continuous

Software

Line0

- *Continuous*: This trigger mode allows the camera to continually capture frames without waiting for a trigger source.
- *Software*: Software triggering is used. The camera can be triggered by executing the *trigger* operation via the FogLAMP control interface.
- *Line0*: This uses the hardware trigger via input 0 of the cameras M12 connector. A single frame will be captured for each trigger event.
- **Trigger Edge**: Used only if the value of *Trigger* is set to *Line0*, this defines if the triggering occurs on the rising edge, falling edge or either edge of the trigger input.

Trigger Edge

✓ RisingEdge

FallingEdge

AnyEdge

Output

- **Output**: Define what image data points are written to the asset.

Output

✓ Radiometric data

Heatmap image

Both

Normalise

- *Radiometric Data*: Include a single data point called *radiometric* that contains the raw radiometric data captured.
- *Heatmap*: Include a single data point called *heatmap* that contains a converted heat map image that uses scales of grey to show temperature values.
- *Both*: Include both *radiometric* and *heatmap* data points in the asset created.
- **Normalise**: Only used if the output includes heat map data. Selecting this option will cause the range of grey scale values to be normalised to show only temperature in the range detected by the camera. This will result in greater contrast within the heat map produced.
- **Custom Configuration**: This entry allows a set of configuration nodes to be set. It is provided to allow custom control of cameras and represents the nodes within a JSON document.

Custom Configuration

```

1 {
2   "NoiseReduction": "Off"
3 }

```


Each key/value pair within the JSON document represents the name of a node and its corresponding value. Values may be numeric, boolean, string or enumeration values. Enumerations are expressed as JSON strings.

- **Image Timestamp:** Some cameras support returning a timestamp with each image from the camera. If this is supported by the camera then this configuration option defines how to use that timestamp. The options are
 - *Ignore*: Do not use the image timestamp.
 - *Use as timestamp*: Use the timestamp reported by the camera as the timestamp for the reading.
 - *Add as additional information*: Add the timestamp as an additional data point for the reading.
- **Focus Method:** Set the camera to use either manual or automatic focus as the method of focusing. Not all cameras support automatic focus mode.
- **Auto Focus Method:** Allows the selection of course of fine focusing of the camera. Fine focusing will be a longer process than course focusing but yields a better focus result.
- **Focus Direction:** Control the automatic focus mode, the options of Stop focus, Near Focus or Far Focus are available.
- **Focus Distance:** If the manual focus mode is selected this sets the focal distance of the camera in meters.
- **Auto Focus Area:** Allows the user to set the area within the image that the camera will use when performing automatic focus operations. The area is expressed as a pair of co-ordinates that define the corners of a rectangle.
- **Focus Speed:** Set the speed of the motor used to focus the camera.

Control

This plugin supports the FogLAMP control mechanism, offering a two operation called *trigger* and *focus*.

The *trigger* operation takes no arguments and is used in conjunction with the trigger mode, *Software*. Calling this operation entry point will cause the plugin to capture a frame from the input.

The *focus* operation causes the camera to perform an auto focus operation, this is only valid if the camera has been setup to perform automatic focusing rather than manual focusing. Note all FLIR cameras support automatic focusing.

Data

The assets created by the plugin will consist of a number of data points.

Datapoint	Description
width	The width of the images captured
height	The height of the images captured
depth	The depth of the pixels in the radiometric data, bits per pixel
radiometric	The radiometric image itself
heatmap	The converted heatmap image, if configured
serialNo	The serial number of the camera
reflectedTemperature	The value of reflected temperature returned by the camera.
reflectedAtmosphere	The value of of reflected atmospheric temperature returned by the camera.
objectDistance	The cameras value for the distance to the object
objectEmissivity	The emissivity value of the observed object
relativeHumidity	The relative humidity
externalOpticsTemperature	The temperature of the external optics
FPS	The current setting of frames per second
FocusPosition	The focus position value read from the camera
FocusDistance	The focus distance read from the camera

Note, some of these values are not returned for non-thermal cameras. The data that is returned is design to allow the conversion of the radiometric data to other forms externally from the south plugin.

8.1.41 Suez Water Cloud Service Plugin

The *foglamp-south-SuezWater* plugin is a south plugin designed to enable the pulling of data from the cloud service using the InSight API. The plugin will interrogate the Suez Water InSight API to determine the set of sites and assets that are monitored by Suez Water and will automatically construct unique asset names within FogLAMP for each of the assets discovered.

The plugin will then fetch data for each of the discovered assets each time a poll request is made. The plugin maintains information on previously reported data in order to ensure that data is not duplicated if the poll interval of the south service is less than the rate at which the Suez Water system updates.

Configuration of the Suez Water plugin is similar in nature to all other south services, create a new south service and select the plugin names *SuezWater* from the list of plugins. Click on next and the configuration screen for the plugin will be displayed.

The screenshot displays the configuration interface for the Suez Water Cloud Service Plugin. At the top, a progress bar indicates three steps: 1. Plugin & Service Name, 2. Review Configuration (the active step), and 3. Done. The main configuration area contains the following elements:

- Asset Name Separator:** A text input field containing a period (.)
- Include Site Name:** A checkbox that is currently unchecked.
- Auth. Key:** A password input field showing masked characters (.....) with a copy icon to its right.
- Include Alarm Thresholds:** A checkbox that is currently unchecked.

At the bottom of the form, there are two buttons: 'Previous' and 'Next'.

- **Asset Name Separator:** the string to use to separator the components of the asset name in FogLAMP.
- **Include Site Name:** a toggle that controls if the site name should be included in the FogLAMP asset name.
- **Auth. Key:** the authentication key for the Suez Water InSight API. This is issued by Suez Water and allows access form the FogLAMP instance to the InSight API.
- **Include Alarm Thresholds:** a toggle to allow for the inclusion of alarm thresholds in the data ingested by the plugin

Error Messages

The following are messages that may be produced by the *foglamp-south-SuezWater* plugin, these messages are written to the system log file and may be viewed by the *System* menu item in the FogLAMP user interface. This display may be filtered on the name of a particular south service in order to view just the messages that originate from that south service.




































API request failed This is usually followed by a more detailed explanation, the most common causes of this error are that there is no network connection between the FogLAMP instance and the SuezWater Insight API or an invalid authentication key has been given in the configuration.

No available date range: This is followed by a URL and indicates that an asset has been discovered but that no data can be found for the asset. This should never happen as it indicates that the InSight API is reporting an asset for which it holds no data.

Failed to parse timeseries data response The plugin has received a response from the InSight API that it can not parse. This may occur if there is an interruption in the connection and should be resolved at the next call to the API. If this persists then action may be required to improve the communication link.

Failed to retrieve list of sites The plugin failed to obtain a list of available sites from the InSight API. This may occur if the API has been enabled but no data exists yet within the Suez Water system.

8.1.42 System Information

Asset	Readings		
system/cpuUsage_all	94		
system/diskTraffic_loop0	94		
system/diskTraffic_loop1	94		
system/diskTraffic_loop2	94		
system/diskTraffic_sda	94		
system/diskTraffic_sdb	94		
system/diskUsage_dev/loop0	94		
system/diskUsage_dev/loop1	94		
system/diskUsage_dev/sda2	94		
system/diskUsage_dev/sdb1	94		
system/diskUsage_tmpfs	470		
system/diskUsage_udev	94		
system/hostname	94		
system/loadAverage	94		
system/memInfo	94		
system/networkTraffic_enp0s3	94		
system/networkTraffic_lo	94		
system/pagingAndSwappingEvents	94		
system/platform	94		
system/processes	94		
system/uptime	94		


The *foglamp-south-systeminfo* plugin implements a that collects data about the machine that the FogLAMP instance

is running on. The plugin is designed to allow the monitoring of the edge devices themselves to be included in the monitoring of the equipment involved in processing environment.

The plugin will create a number of assets, in general there are one or more assets per device connected in the case of disks and network interfaces. There are also some generic assets to measure;

- CPU Usage
- Host name
- Load Average
- Memory Usage
- Paging and swapping
- Process information
- System Uptime

A typical output for one of these assets, in this case the processes asset is shown below



To create a south service with the systeminfo plugin

- Click on *South* in the left hand menu bar
- Select *systeminfo* from the plugin list
- Name your service and click *Next*

- Configure the plugin
 - **Asset Name Prefix:** The asset name prefix for the assets created by this plugin. The plugin will create a number of assets, the exact number is dependent on the number of devices attached to the machine.
- Click *Next*
- Enable the service and click on *Done*

8.1.43 Advantech USB-4704



The *foglamp-south-usb4704* plugin is a south plugin that is designed to gather data from an Advantech USB-4704 data acquisition module. The module supports 8 digital inputs and 8 analogue inputs. It is possible to configure the plugin to combine multiple digital input to create a single numeric data point or have each input as a boolean data point. Each analogue input, which is a 14 bit analogue to digital converter, becomes a single numeric data point in the range 0 to 16383, although a scale and offset may be applied to these values.

To create a south service with the USB-4704

- Click on *South* in the left hand menu bar
- Select *usb4704* from the plugin list
- Name your service and click *Next*

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name

Connections

```

1 {
2   "analogue_example": {
3     "type": "analogue",
4     "pin": "AI0",
5     "name": "value1",
6     "scale": 0.1
7   },
8   "digital_example": {
9     "type": "digital",
10    "pins": [
11      "DI0",
12      "DI1",
13      "DI2"

```

- Configure the plugin
 - **Asset Name:** The name of the asset that will be created with the values read from the USB-4704
 - **Connections:** A JSON document that describes the connections to the USB-4704 and the data points within the asset that they map to. The JSON document is a set of objects, one per data point. The objects contain a number of key/value pairs as follow

Key	Description
type	The type of connection, this may be either digital or analogue.
pin	The analogue pin used for the connection.
pins	An array of pins for a digital connection, the first element in the array becomes the most significant bit of the numeric value created.
name	The data point name within the asset.
scale	An optional scale value that may be applied to the value.

- Click on *Next*
- Enable your service and click on *Done*

8.1.44 Video4Linux

The *foglamp-south-video4linux* plugin is a south plugin that is primarily designed for use in computer vision pipelines with FogLAMP. It allows the collection of still frames from any compatible camera that is connected to the Linux machine that is hosting the FogLAMP service.

South services that ingest images using this plugin are created in the same way as any other FogLAMP south service. The plugin has 4 configuration parameters that can be set.

- **Asset Name:** The asset name given to the assets ingested by the plugin
- **Capture Devices:** A list of all the compatible devices found attached to the host.

Select the desired capture device from this list.

- **Images per interval:** The number of frames to collect in a given time interval.
- **Interval:** Interval in seconds over which to collect the number of frames defined above.

Asset Ingestion

A single asset is ingested by the plugin which contains a number of data points

- *img*: The frame itself
- *width*: The width of the captured frame in pixels
- *height*: The height of the captured frame in pixels
- *depth*: The number of bits per pixels within the image

8.1.45 South Webcam Media Plugin

The plugin keeps on taking a video frame from directory or webcam and saves into a directory. It also appends the name of the saved files in the reading generated.

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name

Media type

Media storage directory

File data format

Repeat loop ☐

Camera number

Frames per minute processed

- **‘assetName’**: type: string default: **‘WebcamImages’**: Name of Asset output.
- **‘mediaType’**: type: string default: **‘directory’**: Source from which the media files are generated
- **‘mediaDir’**: type: string default: **‘webcam’**: If the mediaType is camera then the directory where media will be stored. If the mediaType is directory then it is the name of directory inside FOGAMP_ROOT/data where images are stored.
- **‘dataFormat’**: type: enumeration default: **‘IMG’**: Format of files in ‘mediaDir’
- **‘repeatLoop’**: type: boolean default: **‘false’**: If the mediaType is directory is reload when you hit the end playing the images from directory.
- **‘cameraNumber’**: type: integer default: **‘0’**: Number associated with /dev/video in your file system. for example /dev/video0 then use 0.
- **‘fpm’**: type: float default: **‘10.0’**: frames to save per minute.

Execution

To run the south webcam media plugin you can either

1. Copy some images inside some directory in FOGAMP_ROOT/data. Let’s say the directory name is pics. Run the following command.

```
curl -sX POST http://localhost:8081/foglamp/service -d '{"name":"My_web_cam","type":
↳ "south","plugin":"webcam_media","enabled":true,"config":{"assetName":{"value":
↳ "WebcamImages"}, "imageDir":{"value":"pics"}, "mediaType":{"value":"directory"},
↳ "fpm":{"value":"10.0"}}}' |
```

2. Connect a camera to the machine and run the following command.

```
$ v4l2-ctl --list-formats-ext --device /dev/video0
You will see something like
'YUYV' (YUYV 4:2:2)
  Size: Discrete 640x480
    Interval: Discrete 0.033s (30.000 fps)
  Size: Discrete 720x480
```

(continues on next page)

(continued from previous page)

```

Interval: Discrete 0.033s (30.000 fps)
Size: Discrete 1280x720
Interval: Discrete 0.033s (30.000 fps)
Size: Discrete 1920x1080
Interval: Discrete 0.067s (15.000 fps)
Interval: Discrete 0.033s (30.000 fps)
Size: Discrete 2592x1944
Interval: Discrete 0.067s (15.000 fps)
Size: Discrete 0x0

```

Now we know that the id 0 is functional. If no output then try 1,2,3 and so on.

Finally launch the plugin using

```

curl -sX POST http://localhost:8081/foglamp/service -d '{"name":"My_web_cam","type":
↪":"south","plugin":"webcam_media","enabled":true,"config":{"assetName":{"value":
↪"WebcamImages"}, "imageDir":{"value":"webcam"}, "mediaType":{"value":"camera"},
↪"cameraNumber":{"value":"0"}, "fpm":{"value":"10.0"}}}' |jq

```

8.2 FogLAMP North Plugins

8.2.1 OMF

The *OMF* north plugin is included in all distributions of the FogLAMP core and provides the north bound interface to the OSIsoft data historians in all its forms; PI Server, Edge Data Store and OSIsoft Cloud Services.

PI Web API OMF Endpoint

To use the PI Web API OMF endpoint first ensure the OMF option was included in your PI Server when it was installed.

Now go to the FogLAMP user interface, create a new North instance and select the “OMF” plugin on the first screen. The second screen will request the following information:

?

Endpoint

PI Web API

Send full structure

☒

Naming Scheme

Concise

Server hostname

localhost

Server port, 0=use the default

0

Producer Token

omf_north_0001

Data Source

readings

Static Data

Location: Palo Alto, Company: Dianomic

Sleep Time Retry

1

Maximum Retry

3

HTTP Timeout

10

Integer Format

int64

Number Format

float64

Compression

☒

Default Asset Framework Location

/fledge/data_piwebapi/default

Asset Framework hierarchy rules

1

{}

PI Web API Authentication Method

anonymous

PI Web API User Id

user_id

PI Web API Password

PI Web API Kerberos keytab file

piwebapi_kerberos_https.keytab

OCS Namespace

name_space

OCS Tenant ID

ocs_tenant_id

OCS Client ID

ocs_client_id

OCS Client Secret

Select PI Web API from the Endpoint options.

- **Basic Information**

- **Endpoint:** This is the type of OMF endpoint. In this case, choose PI Web API.
- **Send full structure:** Used to control if Asset Framework structure messages are sent to the PI Server. If this is turned off then the data will not be placed in the Asset Framework.
- **Naming scheme:** Defines the naming scheme to be used when creating the PI points in the PI Data Archive. See [Naming Scheme](#).
- **Server hostname:** The hostname or address of the PI Web API server. This is normally the same address as the PI Server.
- **Server port:** The port the PI Web API OMF endpoint is listening on. Leave as 0 if you are using the default port.
- **Data Source:** Defines which data is sent to the PI Server. Choices are: readings or statistics (that is, FogLAMP's internal statistics).
- **Static Data:** Data to include in every reading sent to PI. For example, you can use this to specify the location of the devices being monitored by the FogLAMP server.

- **Asset Framework**

- **Default Asset Framework Location:** The location in the Asset Framework hierarchy into which the data will be inserted. All data will be inserted at this point in the Asset Framework hierarchy unless a later rule overrides this. Note this field does not include the name of the target Asset Framework Database; the target database is defined on the PI Web API server by the PI Web API Admin Utility.
- **Asset Framework Hierarchies Rules:** A set of rules that allow specific readings to be placed elsewhere in the Asset Framework. These rules can be based on the name of the asset itself or some metadata associated with the asset. See [Asset Framework Hierarchy Rules](#).

- **PI Web API authentication**

- **PI Web API Authentication Method:** The authentication method to be used: anonymous, basic or kerberos. Anonymous equates to no authentication, basic authentication requires a user name and password, and Kerberos allows integration with your single signon environment.
- **PI Web API User Id:** For Basic authentication, the user name to authenticate with the PI Web API.
- **PI Web API Password:** For Basic authentication, the password of the user we are using to authenticate.
- **PI Web API Kerberos keytab file:** The Kerberos keytab file used to authenticate.

- **Connection management (These should only be changed with guidance from support)**

- **Sleep Time Retry:** Number of seconds to wait before retrying the HTTP connection (FogLAMP doubles this time after each failed attempt).
- **Maximum Retry:** Maximum number of times to retry connecting to the PI Server.
- **HTTP Timeout:** Number of seconds to wait before FogLAMP will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match FogLAMP data types to the data type configured in PI. This defaults to int64 but may be set to any OMF data type compatible with integer data, e.g. int32.
- **Number Format:** Used to match FogLAMP data types to the data type configured in PI. The default is float64 but may be set to any OMF datatype that supports floating point values.

- **Compression:** Compress the readings data before sending them to the PI Web API OMF endpoint. This setting is not related to data compression in the PI Data Archive.

Edge Data Store OMF Endpoint

To use the OSIsoft Edge Data Store first install Edge Data Store on the same machine as your FogLAMP instance. It is a limitation of Edge Data Store that it must reside on the same host as any system that connects to it with OMF.

Now go to the FogLAMP user interface, create a new North instance and select the “OMF” plugin on the first screen. The second screen will request the following information:

?

Endpoint

Edge Data Store

Send full structure

☒

Naming Scheme

Concise

Server hostname

localhost

Server port, 0=use the default

0

Producer Token

omf_north_0001

Data Source

readings

Static Data

Location: Palo Alto, Company: Dianomic

Sleep Time Retry

1

Maximum Retry

3

HTTP Timeout

10

Integer Format

int64

Number Format

float64

Compression

☒

Default Asset Framework Location

/fledge/data_piwebapi/default

Asset Framework hierarchy rules

1

{ }

PI Web API Authentication Method

anonymous

PI Web API User Id

user_id

PI Web API Password

.....

PI Web API Kerberos keytab file

piwebapi_kerberos_https.keytab

OCS Namespace

name_space

OCS Tenant ID

ocs_tenant_id

OCS Client ID

ocs_client_id

OCS Client Secret

.....

Select Edge Data Store from the Endpoint options.

- **Basic Information**

- **Endpoint:** This is the type of OMF endpoint. In this case, choose Edge Data Store.
- **Naming scheme:** Defines the naming scheme to be used when creating the PI points within the PI Server. See [Naming Scheme](#).
- **Server hostname:** Normally the hostname or address of the OMF endpoint. For Edge Data Store, this must be *localhost*.
- **Server port:** The port the Edge Data Store is listening on. Leave as 0 if you are using the default port.
- **Data Source:** Defines which data is sent to the Edge Data Store. Choices are: readings or statistics (that is, FogLAMP's internal statistics).
- **Static Data:** Data to include in every reading sent to PI. For example, you can use this to specify the location of the devices being monitored by the FogLAMP server.

- **Connection management (These should only be changed with guidance from support)**

- **Sleep Time Retry:** Number of seconds to wait before retrying the HTTP connection (FogLAMP doubles this time after each failed attempt).
- **Maximum Retry:** Maximum number of times to retry connecting to the PI server.
- **HTTP Timeout:** Number of seconds to wait before FogLAMP will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match FogLAMP data types to the data type configured in PI. This defaults to int64 but may be set to any OMF data type compatible with integer data, e.g. int32.
- **Number Format:** Used to match FogLAMP data types to the data type configured in PI. The default is float64 but may be set to any OMF datatype that supports floating point values.
- **Compression:** Compress the readings data before sending them to the Edge Data Store.

OSIsoft Cloud Services OMF Endpoint

Go to the FogLAMP user interface, create a new North instance and select the “OMF” plugin on the first screen. The second screen will request the following information:

?

Endpoint

OSIsoft Cloud Services

Send full structure

☒

Naming Scheme

Concise

Server hostname

localhost

Server port, 0=use the default

0

Producer Token

omf_north_0001

Data Source

readings

Static Data

Location: Palo Alto, Company: Dianomic

Sleep Time Retry

1

Maximum Retry

3

HTTP Timeout

10

Integer Format

int64

Number Format

float64

Compression

☒

Default Asset Framework Location

/fledge/data_piwebapi/default

Asset Framework hierarchy rules

1

{}

PI Web API Authentication Method

anonymous

PI Web API User Id

user_id

PI Web API Password

.....

PI Web API Kerberos keytab file

piwebapi_kerberos_https.keytab

OCS Namespace

name_space

OCS Tenant ID

ocs_tenant_id

OCS Client ID

ocs_client_id

OCS Client Secret

.....

Select OSIsoft Cloud Services from the Endpoint options.

- **Basic Information**

- **Endpoint:** This is the type of OMF endpoint. In this case, choose OSIsoft Cloud Services.
- **Naming scheme:** Defines the naming scheme to be used when creating the PI points within the PI Server. See [Naming Scheme](#).
- **Data Source:** Defines which data is sent to OSIsoft Cloud Services. Choices are: readings or statistics (that is, FogLAMP's internal statistics).
- **Static Data:** Data to include in every reading sent to OSIsoft Cloud Services. For example, you can use this to specify the location of the devices being monitored by the FogLAMP server.

- **Authentication**

- **OCS Namespace:** Your namespace within OSIsoft Cloud Services.
- **OCS Tenant ID:** Your OSIsoft Cloud Services Tenant ID for your account.
- **OCS Client ID:** Your OSIsoft Cloud Services Client ID for your account.
- **OCS Client Secret:** Your OSIsoft Cloud Services Client Secret.

- **Connection management (These should only be changed with guidance from support)**


- **Sleep Time Retry:** Number of seconds to wait before retrying the HTTP connection (FogLAMP doubles this time after each failed attempt).
- **Maximum Retry:** Maximum number of times to retry connecting to the PI server.
- **HTTP Timeout:** Number of seconds to wait before FogLAMP will time out an HTTP connection attempt.

- **Other (Rarely changed)**

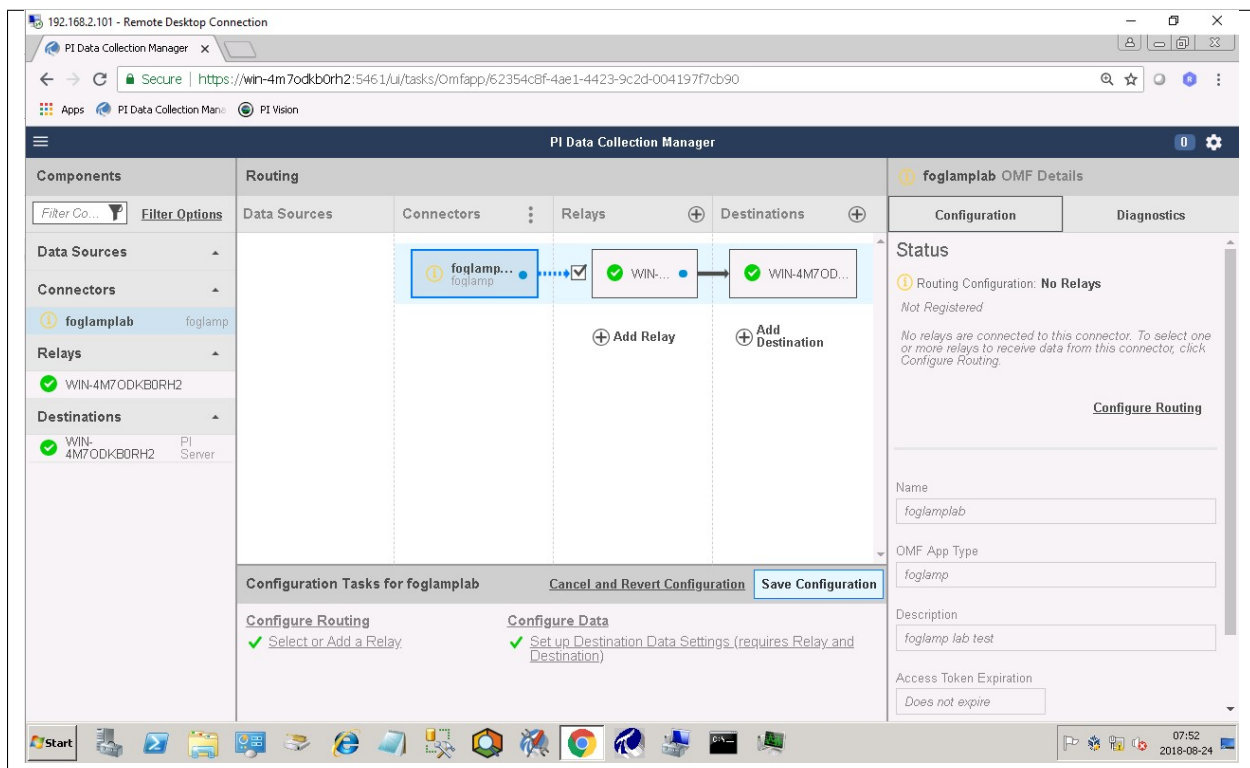
- **Integer Format:** Used to match FogLAMP data types to the data type configured in PI. This defaults to int64 but may be set to any OMF data type compatible with integer data, e.g. int32.
- **Number Format:** Used to match FogLAMP data types to the data type configured in PI. The default is float64 but may be set to any OMF datatype that supports floating point values.
- **Compression:** Compress the readings data before sending them to OSIsoft Cloud Services.

PI Connector Relay

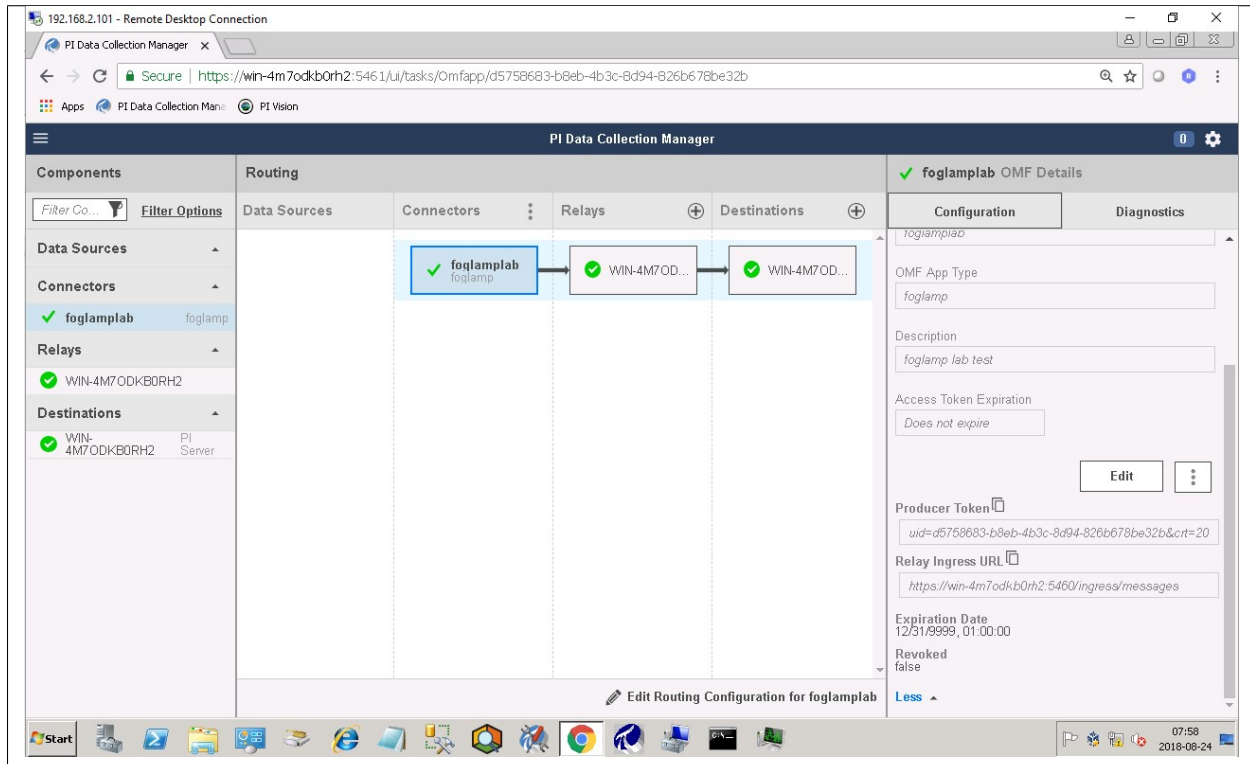
The PI Connector Relay has been discontinued by OSIsoft. All new deployments should use the PI Web API endpoint. Existing installations will still be supported. The PI Connector Relay was the original mechanism by which OMF data could be ingesting into a PI Server. To use the PI Connector Relay, open and sign into the PI Relay Data Connection Manager.



To add a new connector for the FogLAMP system, click on the drop down menu to the right of “Connectors” and select “Add an OMF application”. Add and save the requested configuration information.



Connect the new application to the PI Connector Relay by selecting the new FogLAMP application, clicking the check box for the PI Connector Relay and then clicking “Save Configuration”.



Finally, select the new FogLAMP application. Click “More” at the bottom of the Configuration panel. Make note of the Producer Token and Relay Ingress URL.

Now go to the FogLAMP user interface, create a new North instance and select the “OMF” plugin on the first screen. The second screen will request the following information:

?

Endpoint

Connector Relay

Send full structure

☒

Naming Scheme

Concise

Server hostname

localhost

Server port, 0=use the default

0

Producer Token

omf_north_0001

Data Source

readings

Static Data

Location: Palo Alto, Company: Dianomic

Sleep Time Retry

1

Maximum Retry

3

HTTP Timeout

10

Integer Format

int64

Number Format

float64

Compression

☒

Default Asset Framework Location

/fledge/data_piwebapi/default

Asset Framework hierarchy rules

1

{}

PI Web API Authentication Method

anonymous

PI Web API User Id

user_id

PI Web API Password

.....

PI Web API Kerberos keytab file

piwebapi_kerberos_https.keytab

OCS Namespace

name_space

OCS Tenant ID

ocs_tenant_id

OCS Client ID

ocs_client_id

OCS Client Secret

.....

- **Basic Information**

- **Endpoint:** This is the type of OMF endpoint. In this case, choose Connector Relay.
- **Server hostname:** The hostname or address of the PI Connector Relay.
- **Server port:** The port the PI Connector Relay is listening on. Leave as 0 if you are using the default port.
- **Producer Token:** The Producer Token provided by the PI Relay Data Connection Manager.
- **Data Source:** Defines which data is sent to the PI Connector Relay. Choices are: readings or statistics (that is, FogLAMP's internal statistics).
- **Static Data:** Data to include in every reading sent to PI. For example, you can use this to specify the location of the devices being monitored by the FogLAMP server.

- **Connection management (These should only be changed with guidance from support)**

- **Sleep Time Retry:** Number of seconds to wait before retrying the HTTP connection (FogLAMP doubles this time after each failed attempt).
- **Maximum Retry:** Maximum number of times to retry connecting to the PI server.
- **HTTP Timeout:** Number of seconds to wait before FogLAMP will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match FogLAMP data types to the data type configured in PI. This defaults to int64 but may be set to any OMF data type compatible with integer data, e.g. int32.
- **Number Format:** Used to match FogLAMP data types to the data type configured in PI. The default is float64 but may be set to any OMF datatype that supports floating point values.
- **Compression:** Compress the readings data before sending it to the PI System.

Naming Scheme

The naming of objects in the Asset Framework and of the attributes of those objects has a number of constraints that need to be understood when storing data into a PI Server using OMF. An important factor in this is the stability of your data structures. If you have objects in your environment that are likely to change, you may wish to take a different naming approach. Examples of changes are a difference in the number of attributes between readings, and a change in the data types of attributes.

This occurs because of a limitation of the OMF interface to the PI Server. Data is sent to OMF in a number of stages. One of these is the definition of the Types used to create AF Element Templates. OMF uses a Type to define an AF Element Template but once defined it cannot be changed. If an updated Type definition is sent to OMF, it will be used to create a new AF Element Template rather than changing the existing one. This means a new AF Element Template is created each time a Type changes.

The OMF plugin names objects in the Asset Framework based upon the asset name in the reading within FogLAMP. Asset names are typically added to the readings in the south plugins, however they may be altered by filters between the south ingest and the north egress points in the data pipeline. Asset names can be overridden using the *OMF Hints* mechanism described below.

The attribute names used within the objects in the PI System are based on the names of the datapoints within each Reading within FogLAMP. Again *OMF Hints* can be used to override this mechanism.

The naming used within the objects in the Asset Framework is controlled by the *Naming Scheme* option:

Concise No suffix or prefix is added to the asset name and property name when creating objects in the Asset Framework and PI Points in the PI Data Archive. However, if the structure of an asset changes a new AF Element Template will be created which will have the suffix `-type*x*` appended to it.

Use Type Suffix The AF Element names will be created from the asset names by appending the suffix `-type*x*` to the asset name. If the structure of an asset changes a new AF Element name will be created with an updated suffix.

Use Attribute Hash AF Attribute names will be created using a numerical hash as a prefix.

Backward Compatibility The naming reverts to the rules that were used by version 1.9.1 and earlier of FogLAMP: both type suffixes and attribute hashes will be applied to the name.

Asset Framework Hierarchy Rules

The Asset Framework rules allow the location of specific assets within the Asset Framework to be controlled. There are two basic types of hint:

- Asset name placement: the name of the asset determines where in the Asset Framework the asset is placed,
- Meta data placement: metadata within the reading determines where the asset is placed in the Asset Framework.

The rules are encoded within a JSON document. This document contains two properties in the root of the document: one for name-based rules and the other for metadata based rules.

```
{
  "names" :
  {
    "asset1" : "/Building1/EastWing/GroundFloor/Room4",
    "asset2" : "Room14"
  },
  "metadata" :
  {
    "exist" :
    {
      "temperature" : "temperatures",
      "power"       : "/Electrical/Power"
    },
    "nonexist" :
    {
      "unit" : "Uncalibrated"
    },
    "equal" :
    {
      "room" :
      {
        "4" : "ElecticalLab",
        "6" : "FluidLab"
      }
    },
    "notequal" :
    {
      "building" :
      {
        "plant" : "/Office/Environment"
      }
    }
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

The name type rules are simply a set of asset name and Asset Framework location pairs. The asset names must be complete names; there is no pattern matching within the names.

The metadata rules are more complex. Four different tests can be applied:

- **exists:** This test looks for the existence of the named datapoint within the asset.
- **nonexist:** This test looks for the lack of a named datapoint within the asset.
- **equal:** This test looks for a named datapoint having a given value.
- **notequal:** This test looks for a name datapoint having a value different from that specified.

The *exist* and *nonexist* tests take a set of name/value pairs that are tested. The name is the datapoint name to examine and the value is the Asset Framework location to use. For example

```
"exist" :
{
    "temperature" : "temperatures",
    "power"       : "/Electrical/Power"
}
```

If an asset has a datapoint called *temperature* it will be stored in the AF hierarchy *temperatures*, if the asset had a datapoint called *power* the asset will be placed in the AF hierarchy */Electrical/Power*.

The *equal* and *notequal* tests take an object as a child, the name of the object is datapoint to examine, the child nodes are sets of values and locations. For example

```
"equal" :
{
    "room" :
    {
        "4" : "ElectricalLab",
        "6" : "FluidLab"
    }
}
```

In this case if the asset has a datapoint called *room* with a value of *4* then the asset will be placed in the AF location *ElectricalLab*, if it has a value of *6* then it is placed in the AF location *FluidLab*.

If an asset matches multiple rules in the ruleset it will appear in multiple locations in the hierarchy, the data is shared between each of the locations.

If an OMF Hint exists within a particular reading this will take precedence over generic rules.

The AF location may be a simple string or it may also include substitutions from other datapoints within the reading. For example if the reading has a datapoint called *room* that contains the room in which the readings were taken, an AF location of */BuildingA/\${room}* would put the reading in the Asset Framework using the value of the room datapoint. The reading

```
"reading" : {
    "temperature" : 23.4,
    "room"        : "B114"
}
```

would be put in the AF at */BuildingA/B114* whereas a reading of the form

```
"reading" : {
  "temperature" : 24.6,
  "room"       : "2016"
}
```

would be put at the location */BuildingA/2016*.

It is also possible to define defaults if the referenced datapoint is missing. In our example above if we used the location */BuildingA/\${room:unknown}* a reading without a *room* datapoint would be placed in */BuildingA/unknown*. If no default is given and the data point is missing then the level in the hierarchy is ignored. E.g. if we use our original location */BuildingA/\${room}* and we have the reading

```
"reading" : {
  "temperature" : 22.8,
}
```

this reading would be stored in */BuildingA*.

OMF Hints

The OMF plugin also supports the concept of hints in the actual data that determine how the data should be treated by the plugin. Hints are encoded in a specially named datapoint within the asset, *OMFHint*. The hints themselves are encoded as JSON within a string.

Number Format Hints

A number format hint tells the plugin what number format to use when inserting data into the PI Server. The following will cause all numeric data within the asset to be written using the format *float32*.

```
"OMFHint" : { "number" : "float32" }
```

The value of the *number* hint may be any numeric format that is supported by the PI Server.

Integer Format Hints

An integer format hint tells the plugin what integer format to use when inserting data into the PI Server. The following will cause all integer data within the asset to be written using the format *integer32*.

```
"OMFHint" : { "number" : "integer32" }
```

The value of the *number* hint may be any numeric format that is supported by the PI Server.

Type Name Hints

A type name hint specifies that a particular name should be used when defining the name of the type that will be created to store the object in the Asset Framework. This will override the *Naming Scheme* currently configured.

```
"OMFHint" : { "typeName" : "substation" }
```


Type Hint

A type hint is similar to a type name hint, but instead of defining the name of a type to create it defines the name of an existing type to use. The structure of the asset *must* match the structure of the existing type with the PI Server, it is the responsibility of the person that adds this hint to ensure this is the case.

```
"OMFHint" : { "type" : "pump" }
```

Tag Name Hint

Specifies that a specific tag name should be used when storing data in the PI Server.

```
"OMFHint" : { "tagName" : "AC1246" }
```

Datapoint Specific Hint

Hints may also be targeted to specific data points within an asset by using the datapoint hint. A *datapoint* hint takes a JSON object as its value; the object defines the name of the datapoint and the hint to apply.

```
"OMFHint" : { "datapoint" : { "name" : "voltage:", "number" : "float32" } }
```

The above hint applies to the datapoint *voltage* in the asset and applies a *number format* hint to that datapoint.

Asset Framework Location Hint

An Asset Framework location hint can be added to a reading to control the placement of the asset within the Asset Framework. An Asset Framework hint would be as follows:

```
"OMFHint" : { "AFLocation" : "/UK/London/TowerHill/Floor4" }
```

Note the following when defining an *AFLocation* hint:

- An asset in a FogLAMP Reading is used to create a [Container in the OSIsoft Asset Framework](#). A *Container* is an AF Element with one or more AF Attributes that are mapped to PI Points using the OSIsoft PI Point Data Reference. The name of the AF Element comes from the FogLAMP Reading asset name. The names of the AF Attributes come from the FogLAMP Reading datapoint names.
- If you edit the AF Location hint, the Container will be moved to the new location in the AF hierarchy.
- If you disable the OMF Hint filter, the Container will not move.
- If you wish to move a Container, you can do this with the PI System Explorer. Right-click on the AF Element that represents the Container. Choose Copy. Select the AF Element that will serve as the new parent of the Container. Right-click and choose *Paste*. You can then return to the original Container and delete it. *Note that PI System Explorer does not have the traditional Cut function for AF Elements.*
- If you move a Container, OMF North will not recreate it. If you then edit the AF Location hint, the Container will appear in the new location.

Adding OMF Hints

An OMF Hint is implemented as a string data point on a reading with the data point name of *OMFHint*. It can be added at any point in the processing of the data, however a specific plugin is available for adding the hints, the .

8.2.2 Azure IoT Hub

The *foglamp-north-azure* plugin sends data from FogLAMP to the Microsoft Azure IoT Core service.

The configuration of the *Azure* plugin requires a few simple configuration parameters to be set.

- **Primary Connection String:** The primary connection string to connect to your Azure IoT project. The connection string should contain
 - The hostname to connect to
 - The DeviceID of the device you are using
 - The shared access key generated on your Azure login
- **MQTT over websockets:** Enable if you wish to run MQTT over websockets.
- **Data Source:** Which FogLAMP data to send to Azure; Readings or FogLAMP Statistics.
- **Apply Filter:** This allows a simple jq format filter rule to be applied to the connection. This should not be confused with FogLAMP filters and exists for backward compatibility reason only.
- **Filter Rule:** A jq filter rule to apply. Since the introduction of FogLAMP filters in the north task this has become deprecated and should not be used.

JSON Payload

The payload that is sent by this plugin to Azure is a simple JSON representation of a set of reading values. A JSON array is sent with one or more reading objects contained within it. Each reading object consists of a timestamp, an asset name and a set of data points within that asset. The data points are represented as name value pair JSON properties within the reading property.

The fixed part of every reading contains the following

Name	Description
timestamp	The timestamp as an ASCII string in ISO 8601 extended format. If no time zone information is given it is assumed to indicate the use of UTC.
asset	The name of the asset this reading represents.
readings	A JSON object that contains the data points for this asset.

The content of the *readings* object is a set of JSON properties, each of which represents a data value. The type of these values may be integer, floating point, string, a JSON object or an array of floating point numbers.

A property

```
"voltage" : 239.4
```

would represent a numeric data value for the item *voltage* within the asset. Whereas

```
"voltageUnit" : "volts"
```

Is string data for that same asset. Other data may be presented as arrays

```
"acceleration" : [ 0.4, 0.8, 1.0 ]
```

would represent acceleration with the three components of the vector, x, y, and z. This may also be represented as an object

```
"acceleration" : { "X" : 0.4, "Y" : 0.8, "Z" : 1.0 }
```

both are valid formats within FogLAMP.

An example payload with a single reading would be as shown below

```
[
  {
    "timestamp" : "2020-07-08 16:16:07.263657+00:00",
    "asset"      : "motor1",
    "readings"   : {
      "voltage"  : 239.4,
      "current"  : 1003,
      "rpm"      : 120147
    }
  }
]
```

8.2.3 Google Cloud Platform North Plugin

The *foglamp-north-gcp* plugin provide connectivity from a FogLAMP system to the Google Cloud Platform. The plugin connects to the IoT Core in Google Cloud using MQTT and is fully compliant with the security features of the Google Cloud Platform. See for a tutorial on setting up a FogLAMP system and getting it to send data to Google Cloud.

Prerequisites


A number of things must be done in the Google Cloud before you can create your north connection to GCP. You must

- Create a GCP IoT Core project
- Download the *roots.pem* certificate from your GCP account
- Create a registry
- Create a device ID and configure a key pair for that device
- Upload the certificates to the FogLAMP certificate store

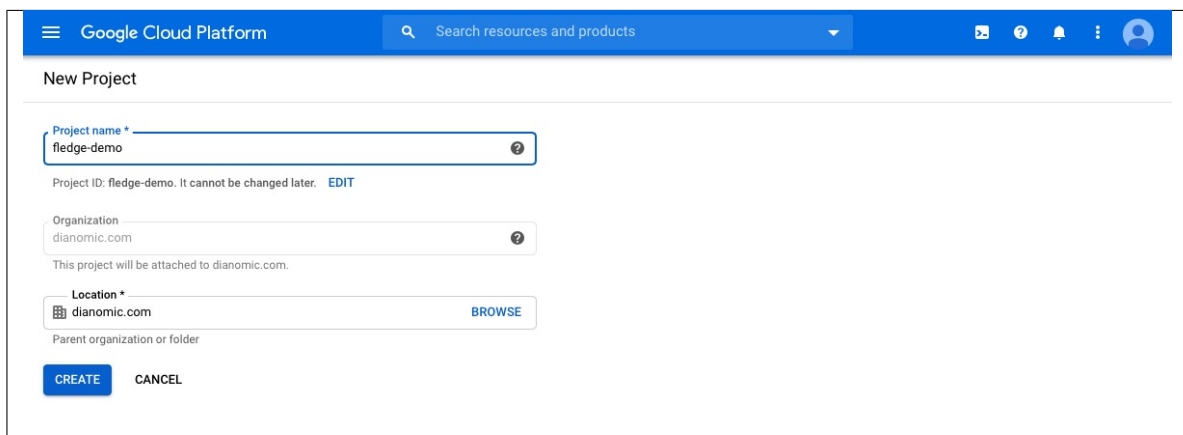
Create GCP IoT Core Project

To create a new project

- Goto the
- Select the Projects page and select the *Create New Project* option



- Enter your project details



Download roots.pem

To download the roots.pem security certificate

- From the command line shell of your machine run the command

```
$ wget https://pki.goog/roots.pem
```


Create a Registry

To create a registry in your project

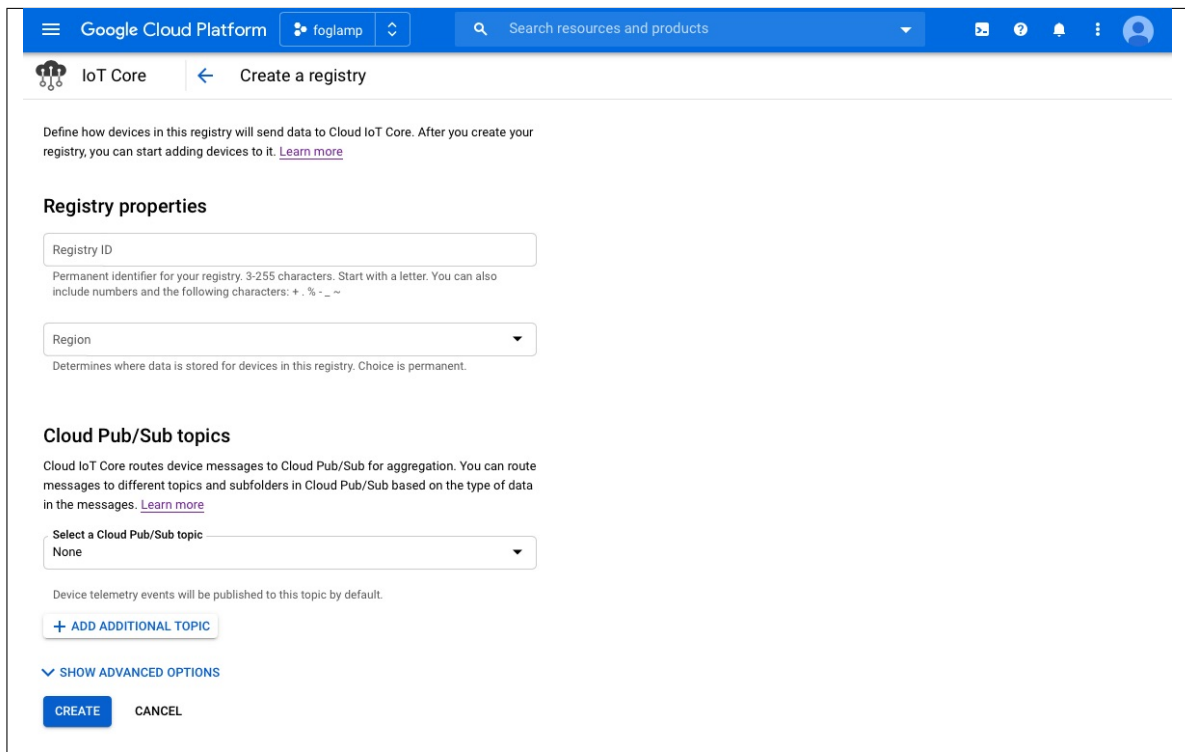
- Goto the



- Click on the menu icon in the top left corner of the page
- Select the *Create Registry* option



- A new screen is shown that allows you to create a new registry



- Note the Registry ID and region as you will need these later
- Select an existing telemetry topic or create a new topic (for example, projects/[YOUR_PROJECT_ID]/topics/[REGISTRY_ID])
- Click on *Create*

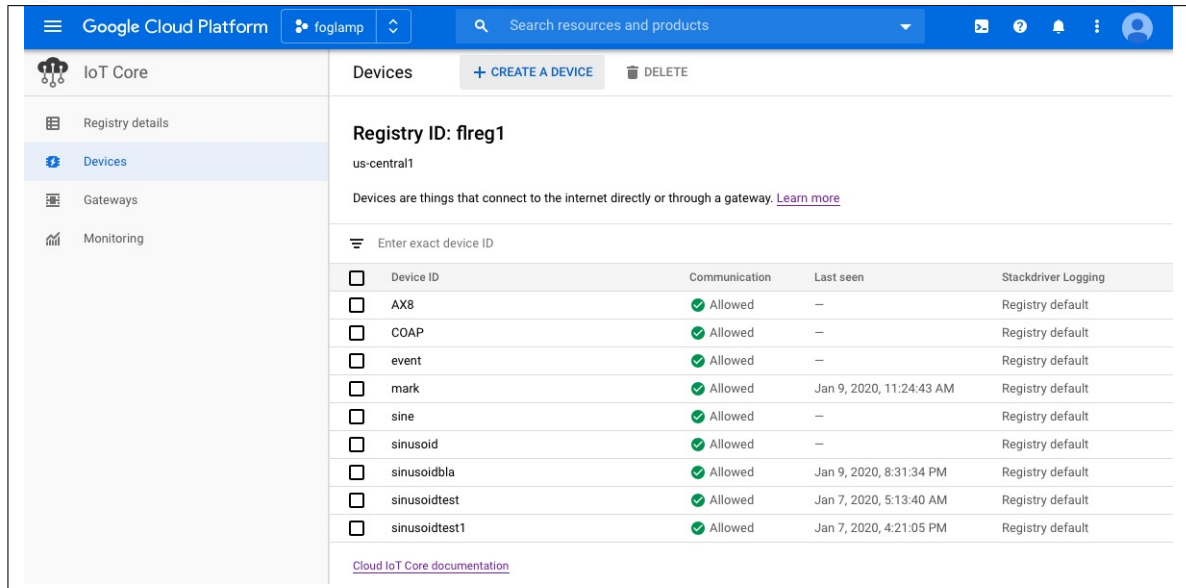
Create a Device ID

To create a device in your Google Cloud Project

- Create an RSA public/private key pair on your local machine

```
openssl genpkey -algorithm RSA -out rsa_foglamp.pem -pkeyopt rsa_keygen_bits:2048
openssl rsa -in rsa_foglamp.pem -pubout -out rsa_foglamp.pem
```

- Goto the
- In the left pane of the IoT Core page in the Cloud Console, click Devices



The screenshot shows the Google Cloud Platform IoT Core console. The left sidebar contains navigation links: Registry details, Devices (selected), Gateways, and Monitoring. The main content area is titled 'Devices' and shows the Registry ID 'freg1' and region 'us-central1'. Below this, there is a search bar and a table of devices. The table has columns for Device ID, Communication, Last seen, and Stackdriver Logging. The devices listed are AX8, COAP, event, mark, sine, sinusoid, sinusoidbla, sinusoidtest, and sinusoidtest1. All devices show 'Allowed' communication status.

Device ID	Communication	Last seen	Stackdriver Logging
AX8	Allowed	—	Registry default
COAP	Allowed	—	Registry default
event	Allowed	—	Registry default
mark	Allowed	Jan 9, 2020, 11:24:43 AM	Registry default
sine	Allowed	—	Registry default
sinusoid	Allowed	—	Registry default
sinusoidbla	Allowed	Jan 9, 2020, 8:31:34 PM	Registry default
sinusoidtest	Allowed	Jan 7, 2020, 5:13:40 AM	Registry default
sinusoidtest1	Allowed	Jan 7, 2020, 4:21:05 PM	Registry default

- At the top of the Devices page, click *Create a device*

- Enter a device ID, you will need to add this in the north plugin configuration later
- Click on the *ADD ATTRIBUTE COMMUNICATION, STACKDRIVER LOGGING, AUTHENTICATION* link to open the remainder of the inputs
- Make sure the public key format matches the type of key that you created in the first step of this section (for example, RS256)
- Paste the contents of your public key in the Public key value field.

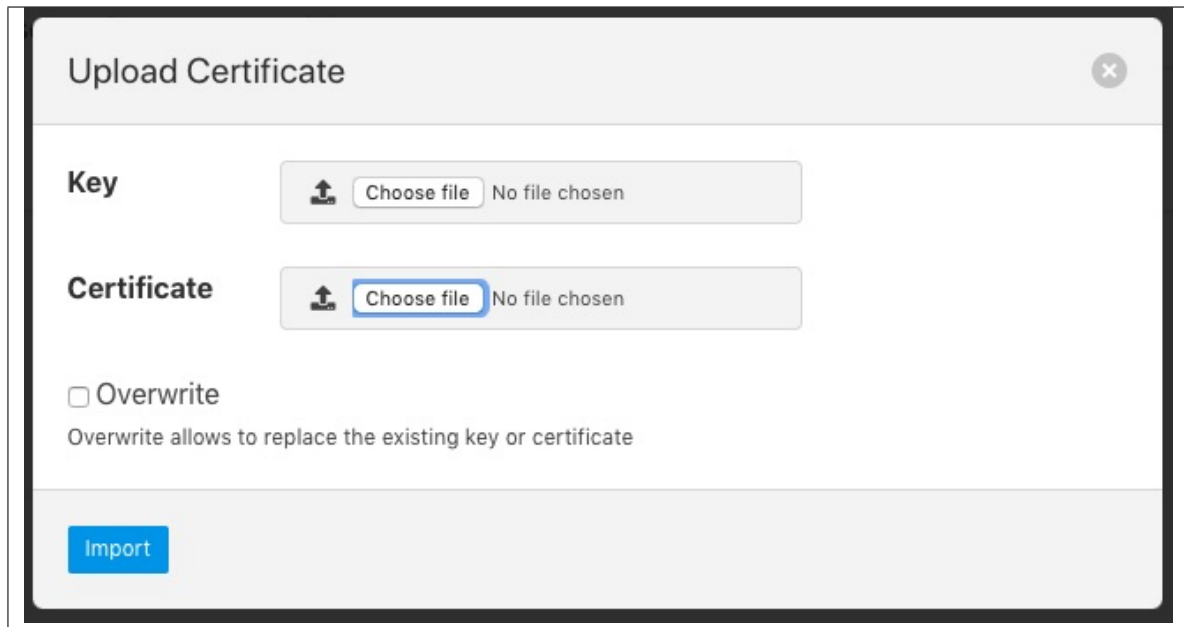
Upload Your Certificates

You should upload your certificates to FogLAMP


- From the FogLAMP user interface select the *Certificate Store* from the left-hand menu bar


Key	Extension		Certificate	Extension	
admin	key	delete	admin	cert	delete
ca	key	delete	ca	cert	delete
fledge	key		fledge	cert	
user	key	delete	user	cert	delete
			route	cert	delete

- Click on the Import option in the top left corner



Upload Certificate

Key  **Choose file** No file chosen

Certificate  **Choose file** No file chosen

☐ **Overwrite**
Overwrite allows to replace the existing key or certificate

Import

- In the Certificate option select the *Choose file* option and select your roots.pem and click on open
- Repeat the above for your device key and certificate

Create Your North Task

Having completed the pre-requisite steps it is now possible to create the north task to send data to GCP.

- Select the *North* option from the left-hand menu bar.
- Select GCP from the North Plugin list
- Name your North task and click on *Next*

The screenshot shows a configuration wizard with three steps: 1. Plugin & Name, 2. Review Configuration, and 3. Done. The 'Review Configuration' step is currently active. It displays a form with the following fields and values:

- Project ID:** fledge-demo
- The GCP Region:** us-central1
- Registry ID:** fledge
- Device ID:** demo
- Key Name:** fledge_private
- JWT Algorithm:** RS256
- Data Source:** readings

At the bottom of the form, there are two buttons: 'Back' and 'Next'.

- Configure your GCP plugin
 - **Project ID:** Enter the project ID you created in GCP
 - **The GCP Region:** Select the region in which you created your registry
 - **Registry ID:** The Registry ID you created should be entered here
 - **Device ID:** The Device ID you created should be entered here
 - **Key Name:** Enter the name of the device key you uploaded to the certificate store
 - **JWT Algorithm:** Select the algorithm that matches the key you created earlier
 - **Data Source:** Select the data to send to GCP, this may be readings or FogLAMP statistics
- Click on *Next*
- Enable your plugin and click on *Done*

8.2.4 Graphite

The *foglamp-north-graphite* plugin provides a means to send data from FogLAMP to the Carbon storage within Graphite to allow data to be graphed.

To create the connection to Graphite

- Select *North* from the left hand menu bar.
- Click on the + icon in the top left
- Choose *graphite* from the plugin selection list
- Name your task
- Click on *Next*
- Configure the plugin

1 Plugin & Name 2 Review Configuration 3 Done

Host: graphite.local

Port: 3000

Asset Root: foglamp

Source: readings

Back Next

- **Host:** The host where Graphite is running.
- **Port:** The carbon listening port of your Graphite Carbon engine.
- **Asset Root:** The root of the asset structure to use with Graphite.

- Click on *Next*
- Enable your north task and click on *Done*

8.2.5 HarperDB

The *foglamp-north-harperdb* plugin sends data from FogLAMP to the database. HarperDB is a geo-distributed database with hybrid SQL & NoSQL functionality in one powerful tool, accessed via a REST API. Each asset that is read by FogLAMP is written to a separate table within the specified HarperDB schema. The plugin will support both local installations and cloud installations of HarperDB.

The configuration of the *HarperDB* plugin requires a few simple configuration parameters to be set.

1 Plugin & Name 2 Review Configuration 3 Done

URL: http://localhost:9925/

Username:

Password: password

Schema: fledge

Source: readings

Back Next

- **URL:** The URL of the HarperDB database that will be used to store the data sent from FogLAMP. This may be either an HTTP or HTTPS URL

- **Username:** The username to use when authenticating with the HarperDB database.
- **Password:** The password of the user that will be used to store the data in HarperDB.
- **Schema:** The name of the schema in the HarperDB database in which the tables will be stored.
- **Source:** The source of the FogLAMP data to store in the HarperDB database; Readings or FogLAMP Statistics.

8.2.6 North HTTP

The *foglamp-north-http* plugin allows data to be sent from the north of one FogLAMP instance into the south of another FogLAMP instance. It allows hierarchies of FogLAMP instances to be built. The FogLAMP to which the data is sent must run the corresponding in order for data to flow between the two FogLAMP instances. The plugin supports both HTTP and HTTPS transport protocols and sends a JSON payload of reading data in the internal FogLAMP format.

The plugin may also be used to send data from FogLAMP to another system, the receiving system should implement a REST end point that will accept a POST request containing JSON data. The format of the JSON payload is described below. The required REST endpoint path is defined in the configuration of the plugin.

Filters may be applied to the connection in either the north task that loads this plugin or the receiving south service on the up stream FogLAMP.

A of this plugin exists also that performs the same function as this plugin, the pair are provided for purposes of comparison and the user may choose whichever they prefer to use.

To create a north task to send to another FogLAMP you should first create the that will receive the data. Then create a new north tasks by;

- Selecting *North* from the left hand menu bar.
- Click on the + icon in the top left
- Choose *http_north* from the plugin selection list
- Name your task
- Click on *Next*
- Configure the plugin

The screenshot displays the 'Review Configuration' step of the plugin setup process. At the top, a progress bar indicates the sequence: 1. Plugin & Name, 2. Review Configuration (current step), and 3. Done. The configuration form contains the following fields:

- URL:** A text input field containing 'http://localhost:6683/sensor-reading'.
- Source:** A dropdown menu with 'readings' selected.
- Verify SSL:** An unchecked checkbox.
- Apply Filter:** An unchecked checkbox.
- Filter Rule:** A text input field containing '.[]'.

At the bottom of the form, there are two buttons: 'Back' and 'Next'.

- **URL:** The URL of the receiving , the address and port should match the service in the up stream FogLAMP. The URL can specify either HTTP or HTTPS protocols.

- **Source:** The data to send, this may be either the reading data or the statistics data
 - **Verify SSL:** When HTTPS rather than HTTP is used this toggle allows for the verification of the certificate that is used. If a self signed certificate is used then this should not be enabled.
 - **Apply Filter:** This allows a simple jq format filter rule to be applied to the connection. This should not be confused with FogLAMP filters and exists for backward compatibility reason only.
 - **Filter Rule:** A jq filter rule to apply. Since the introduction of FogLAMP filters in the north task this has become deprecated and should not be used.
- Click *Next*
 - Enable your task and click *Done*

JSON Payload

The payload that is sent by this plugin is a simple JSON presentation of a set of reading values. A JSON array is sent with one or more reading objects contained within it. Each reading object consists of a timestamp, an asset name and a set of data points within that asset. The data points are represented as name value pair JSON properties within the reading property.

The fixed part of every reading contains the following

Name	Description
times-tamp	The timestamp as an ASCII string in ISO 8601 extended format. If no time zone information is given it is assumed to indicate the use of UTC.
asset	The name of the asset this reading represents.
read-ings	A JSON object that contains the data points for this asset.

The content of the *readings* object is a set of JSON properties, each of which represents a data value. The type of these values may be integer, floating point, string, a JSON object or an array of floating point numbers.

A property

```
"voltage" : 239.4
```

would represent a numeric data value for the item *voltage* within the asset. Whereas

```
"voltageUnit" : "volts"
```

Is string data for that same asset. Other data may be presented as arrays

```
"acceleration" : [ 0.4, 0.8, 1.0 ]
```

would represent acceleration with the three components of the vector, x, y, and z. This may also be represented as an object

```
"acceleration" : { "X" : 0.4, "Y" : 0.8, "Z" : 1.0 }
```

both are valid formats within FogLAMP.

An example payload with a single reading would be as shown below

```
[
  {
    "timestamp" : "2020-07-08 16:16:07.263657+00:00",
    "asset"      : "motor1",
    "readings"   : {
      "voltage"  : 239.4,
      "current"  : 1003,
      "rpm"      : 120147
    }
  }
]
```

8.2.7 North HTTP-C

The *foglamp-north-http-c* plugin allows data to be sent from the north of one FogLAMP instance into the south of another FogLAMP instance. It allows hierarchies of FogLAMP instances to be built. The FogLAMP to which the data is sent must run the corresponding in order for data to flow between the two FogLAMP instances. The plugin supports both HTTP and HTTPS transport protocols and sends a JSON payload of reading data in the internal FogLAMP format.

Additionally this plugin allows for two URL's to be configured, a primary URL and a secondary URL. If the connection to the primary URL fails then the plugin will switch over to using the secondary URL. It will switch back if the connection to the secondary fails or if when the north task completes and a new north task is later run.

The plugin may also be used to send data from FogLAMP to another system, the receiving system should implement a REST end point that will accept a POST request containing JSON data. The format of the JSON payload is described below. The required REST endpoint path is defined in the configuration of the plugin.

Filters may be applied to the connection in either the north task that loads this plugin or the receiving south service on the up stream FogLAMP.

A plugin exists also that performs the same function as this plugin, the pair are provided for purposes of comparison and the user may choose whichever they prefer to use.

To create a north task to send to another FogLAMP you should first create the that will receive the data. Then create a new north tasks by;

- Selecting *North* from the left hand menu bar.
- Click on the + icon in the top left
- Choose httpc from the plugin selection list
- Name your task
- Click on *Next*
- Configure the HTTP-C plugin

The screenshot shows the 'Review Configuration' step of the FogLAMP plugin configuration. The interface includes the following fields and options:

- URL:** `http://localhost:6683/sensor-reading`
- Secondary URL:** (empty)
- Proxy:** (empty)
- Source:** `readings`
- Headers:** A list with one item: `1: {}`
- Script:** A list with one item: `1:`
- Sleep Time Retry:** `1`
- Maximum Retry:** `3`
- Http Timeout (in seconds):** `10`
- Verify SSL:** ☐

At the bottom, there are 'Back' and 'Next' buttons.

- **URL:** The URL of the receiving , the address and port should match the service in the up stream FogLAMP. The URL can specify either HTTP or HTTPS protocols.
- **Secondary URL:** The URL to failover to if the connection to the primary URL fails. If failover is not required then leave this field empty.
- **Source:** The data to send, this may be either the reading data or the statistics data
- **Proxy:** The host and port of the proxy server to use. Leave empty is a proxy is not in use. This should be formatted as an address followed by a colon and then the port or a hostname followed by a colon and then the port. E.g. 192.168.0.42:8080. If the default port is used then the port may be omitted.
- **Headers:** An optional set of header fields to send in every request. The headers are defined as a JSON document with the name of each item in the document as header field name and the value the value of the header field.
- **Script:** An optional Python script that can be used to convert the payload format. If given the script should contain a method called *convert* that will be passed a single reading as a JSON DICT and must return the new payload as a string.
- **Sleep Time Retry:** A tuning parameter used to control how often a connection is retried to the up stream FogLAMP if it is not available. On every retry the time will be doubled.
- **Maximum Retry:** The maximum number of retries to make a connection to the up stream FogLAMP.

When this number is reached the current execution of the task is suspended until the next scheduled run.

- **Http Timeout (in seconds):** The timeout to set on the HTTP connection after which the connection will be closed. This can be used to tune the response of the system when communication links are unreliable.
- **Verify SSL:** When HTTPS rather the HTTP is used this toggle allows for the verification of the certificate that is used. If a self signed certificate is used then this should not be enabled.

- Click *Next*
- Enable your task and click *Done*

Header Fields

Header fields can be defined if required using the *Headers* configuration item. This is a JSON document that defines a set of key/value pairs for each header field. For example if a header field *token* was required with the value of *sfe93rjfk93rj* then the *Headers* JSON document would be as follows

```
{
  "token" : "sfe93rjfk93rj"
}
```

Multiple header fields may be set by specifying multiple key/value pairs in the JSON document.

JSON Payload

The payload that is sent by this plugin is a simple JSON presentation of a set of reading values. A JSON array is sent with one or more reading objects contained within it. Each reading object consists of a timestamp, an asset name and a set of data points within that asset. The data points are represented as name value pair JSON properties within the reading property.

The fixed part of every reading contains the following

Name	Description
ts	The timestamp as an ASCII string in ISO 8601 extended format. If no time zone information is given it is assumed to indicate the use of UTC. This timestamp is added by FogLAMP when it first reads the data.
user_ts	The timestamp as an ASCII string in ISO 8601 extended format. If no time zone information is given it is assumed to indicate the use of UTC. This timestamp is added by the device itself and can be used to reflect the timestamp the data refers to rather than the timestamp FogLAMP read the data.
asset	The name of the asset this reading represents.
readings	A JSON object that contains the data points for this asset.

The content of the *readings* object is a set of JSON properties, each of which represents a data value. The type of these values may be integer, floating point, string, a JSON object or an array of floating point numbers.

A property

```
"voltage" : 239.4
```

would represent a numeric data value for the item *voltage* within the asset. Whereas

```
"voltageUnit" : "volts"
```

Is string data for that same asset. Other data may be presented as arrays

```
"acceleration" : [ 0.4, 0.8, 1.0 ]
```

would represent acceleration with the three components of the vector, x, y, and z. This may also be represented as an object

```
"acceleration" : { "X" : 0.4, "Y" : 0.8, "Z" : 1.0 }
```

both are valid formats within FogLAMP.

An example payload with a single reading would be as shown below

```
[
  {
    "user_ts"    : "2020-07-08 16:16:07.263657+00:00",
    "ts"         : "2020-07-08 16:16:07.263657+00:00",
    "asset"      : "motor1",
    "readings"   : {
      "voltage"  : 239.4,
      "current"  : 1003,
      "rpm"      : 120147
    }
  }
]
```

Payload Script

If a script is given then it must provide a method called *convert*, that method is passed a single reading as a Python DICT and must return a formatted string payload for that reading.

As a simple example lets assume we want a JSON payload to be sent, but we want to use different keys to those in the default reading payload. We will replace *readings* with *data*, *user_ts* with *when* and *asset* with *device*. A simple Python script to do this would be as follows;

```
import json
def convert(reading):
    newReading = {
        "data" : reading["readings"],
        "when" : reading["user_ts"],
        "device" : reading["asset"],
    }
    return json.dumps(newReading)
```

An HTTP request would be sent with one reading per request and that reading would be formatted as a JSON payload of the format

```
{
  "data":
  {
    "sinusoid": 0.0,
    "sine10": 10.0
  },
  "when": "2022-02-16 15:12:55.196494+00:00",
  "device": "sinusoid"
}
```

Note that white space and newlines have been added to improve the readability of the payload.

The above example returns a JSON format payload, the return may however not be encoded as JSON, for example an XML payload

```
from dict2xml import dict2xml
def convert(reading):
    newReading = {
        "data" : reading["readings"],
        "when" : reading["user_ts"],
        "device" : reading["asset"],
    }
    payload = "<reading>" + dict2xml(newReading) + "</reading>"
    return payload
```

This return XML format data as follows

```
<reading>
  <data>
    <sine10>10.0</sine10>
    <sinusoid>0.0</sinusoid>
  </data>
  <device>sinusoid</device>
  <when>2022-02-16 15:12:55.196494+00:00</when>
</reading>
```

Note that white space and formatting have been added for ease of reading the XML data. You must also make sure you have installed the Python XML support as this is not normally installed with FogLAMP, To do this run

```
pip3 install dict2xml
```

from the command line of the FogLAMP machine.

8.2.8 InfluxDB Time Series Database

The *foglamp-north-influxdb* plugin is designed to send data from FogLAMP to the open source time series database.

The process of creating a North InfluxDB is similar to any other north setup

- Selecting the *North* option in the left-hand menu bar
- Click on the add icon in the top right corner.
- In the *North Plugin* list select the influxdb option.
- Click *Next*
- Configure your InfluxDB plugin

The screenshot displays the 'Review Configuration' step of the FogLAMP plugin setup. A progress bar at the top indicates the sequence: 1. Plugin & Name, 2. Review Configuration (active), and 3. Done. The configuration form includes the following fields:

- Host:** influxdb.local
- Port:** 8086
- Database:** foglamp
- Username:** (empty field)
- Password:** password
- Source:** readings (selected from a dropdown menu)

Navigation buttons at the bottom include a 'Back' button and a blue 'Next' button.

- **Host:** The hostname or IP address of the machine where your InfluxDB server is running.
 - **Port:** The port on which your InfluxDB server is listening.
 - **Database:** The database in your InfluxDB server into which to write data.
 - **Username:** The username if any to use to authenticate with your InfluxDB server.
 - **Password:** The password to use to authenticate with your InfluxDB server.
 - **Source:** The source of data to send, this may be either FogLAMP readings or the FogLAMP statistics
- Click *Next*
 - Enable your north task and click on *Done*

8.2.9 InfluxDB Cloud

The *foglamp-north-influxdbcloud* plugin is designed to send data from FogLAMP to the system for collection and analysis of data.

The process of creating a North InfluxDB Cloud connection is similar to any other north setup

- Selecting the *North* option in the left-hand menu bar
- Click on the add icon in the top right corner.
- In the *North Plugin* list select the *influxdbcloud* option.
- Click *Next*
- Configure your InfluxDB Cloud plugin

The screenshot shows the 'Review Configuration' step of the FogLAMP setup process. At the top, a progress bar indicates three steps: 1. Plugin & Name, 2. Review Configuration (current), and 3. Done. The configuration form includes the following fields:

- URL:** `https://eu-central-1-1.aws.cloud2.influxdata.com`
- InfluxDB token:** (empty text field)
- Organisation ID:** (empty text field)
- Bucket:** (empty text field)
- Measurement:** `foglamp`
- Source:** `readings` (dropdown menu)
- Filter Rule:** `.[]`
- Apply Filter:** ☐

At the bottom of the form, there are 'Back' and 'Next' buttons.

- **URL:** The URL of the InfluxDB instance you are using
- **InfluxDB token:** an authorization token that has been generated by the InfluxDB Cloud
- **Organisation ID:** Your organization ID from the InfluxDB Cloud. You can find this by looking at the URL you use after connecting the InfluxDB Cloud.
- **Bucket:** The bucket in InfluxDB CCloud where you wish to store your data.
- **Measurement:** The measurement to use for the data you send to InfluxDB Cloud
- **Source:** The source of data to send, this may be either FogLAMP readings or the FogLAMP statistics
- * **Apply Filter:** This allows a simple jq format filter rule to be applied to the connection. This should not be confused with FogLAMP filters and exists for backward compatibility reasons only.
- **Filter Rule:** A jq filter rule to apply. Since the introduction of FogLAMP filters in the north task this has become deprecated and should not be used.

- Click *Next*
- Enable your north task and click on *Done*

8.2.10 Kafka Producer

The *foglamp-north-kafka* plugin sends data from FogLAMP to the an Apache Kafka. FogLAMP acts as a Kafka producer, sending reading data to Kafka. This implementation is a simplified producer that sends all data on a single Kafka topic. Each message contains an asset name, timestamp and set of readings values as a JSON document.

The configuration of the *Kafka* plugin is very simple, consisting of four parameters that must be set.

1 Plugin & Name 2 Review Configuration 3 Done

Bootstrap Brokers

Kafka Topic

Send JSON

Data Source

- **Bootstrap Brokers:** A comma separate list of Kafka brokers to use to establish a connection to the Kafka system.
- **Kafka Topic:** The Kafka topic to which all data is sent.
- **Send JSON:** This controls how JSON data points should be sent to Kafka. These may be sent as strings or as JSON objects.
- **Data Source:** Which FogLAMP data to send to Kafka; Readings or FogLAMP Statistics.

8.2.11 PNG File Writer

The *foglamp-north-png* plugin is designed as a debugging aid for pipelines that make use of images. It will write copies of all image type data points to PNG file such that they can be verified. It is not expected that this plugin will be used for production systems, although it could be if there was a need to create image files.

The process of creating a North service or task to write images to files starts as with any other north setup by selecting the *North* option in the left-hand menu bar, then press the add icon in the top right corner. In the *North Plugin* list select the *png* option.

Set the name of the task or service, whether to run as a service and the repeat interval for tasks. Once complete click on *Next* and move on to the configuration of the plugin itself.

1 Plugin & Name 2 Review Configuration 3 Done

Directory

Prefix

Source

This second page allows for the setting of the configuration used for writing PNG files

- **Directory:** The directory into which the PNG files will be written.
- **Prefix:** A prefix that is added to the files that are created. The files that are created will use this prefix followed by the asset name, data point name and timestamp of the reading.
- **Source:** What data is being made available to the PNG. You may choose to make the reading data available or the FogLAMP statistics. Note however that the FogLAMP statistics do not contain image data.

Once you have completed your configuration click *Next* to move to the final page and then enable your north task and click *Done*.

8.2.12 Splunk Data Collector

The *foglamp-north-splunk* plugin is designed to send data from FogLAMP to the system for collecting and analysis of data.

The process of creating a North Splunk is similar to any other north setup

- Selecting the *North* option in the left-hand menu bar
- Click on the add icon in the top right corner.
- In the *North Plugin* list select the splunk option.
- Click *Next*
- Configure your Splunk plugin

The screenshot displays the 'Review Configuration' step of the Splunk Data Collector setup. A progress bar at the top indicates the current step (2) and the previous (1) and next (3) steps. The configuration form includes the following fields:

- URL:** A text input field containing 'http://splunk:8088/services/collector/event'.
- Source:** A dropdown menu with 'readings' selected.
- Splunk authorisation token:** A text input field containing '42b66064-ee25-407e-b0ac-c-3ded78d21b38'.
- Apply Filter:** An unchecked checkbox.
- Filter Rule:** A text input field containing '.[]'.

At the bottom of the form are two buttons: 'Back' and 'Next'.

- **URL:** The URL of the splunk collector for events
 - **Source:** The source of data to send, this may be either FogLAMP readings or the FogLAMP statistics
 - **Splunk authorisation token:** an authorisation token that has been issued by the splunk data collector
 - * **Apply Filter:** This allows a simple jq format filter rule to be applied to the connection. This should not be confused with FogLAMP filters and exists for backward compatibility reasons only.
 - **Filter Rule:** A jq filter rule to apply. Since the introduction of FogLAMP filters in the north task this has become deprecated and should not be used.
- Click *Next*

- Enable your north task and click on *Done*

8.2.13 ThingSpeak

The *foglamp-north-thingspeak* plugin provides a mechanism to , allowing an easy route to send data from an FogLAMP environment into MATLAB.

In order to send data to ThingSpeak you must first create a channel to receive it.

- Login to your account
- From the menu bar select the *Channels* menu and the *My Channels* option

The screenshot displays the ThingSpeak web interface. At the top, there is a navigation bar with 'Channels', 'Apps', and 'Support' menus. The main heading is 'My Channels'. Below this, there is a 'New Channel' button and a search bar. A table lists the channels, with one entry 'sinusoid' shown. The table has columns for 'Name', 'Created', and 'Updated'. Below the table name, there are buttons for 'Private', 'Public', 'Settings', 'Sharing', 'API Keys', and 'Data Import / Export'. To the right of the table, there is a 'Help' section with text explaining how to create a channel and sort the table. Below the help section is an 'Examples' section with a list of links to various devices: Arduino, Arduino MKR1000, ESP8266, Raspberry Pi, and Netduino Plus. At the bottom right, there is an 'Upgrade' section with the text 'Need to send more data faster?' and 'Need to use ThingSpeak for a commercial project?', followed by a green 'Upgrade' button.

Name	Created	Updated
sinusoid	2018-08-08	2019-07-26 15:04

Private Public Settings Sharing API Keys Data Import / Export

- Click on *New Channel* to create a new channel

[Channels](#)
[Apps](#)
[Support](#)

[Commercial Use](#)
[How to Buy](#)
[MR](#)

New Channel

Name

Description

Field 1 ☒

Field 2 ☐

Field 3 ☐

Field 4 ☐

Field 5 ☐

Field 6 ☐

Field 7 ☐

Field 8 ☐

Metadata

Tags

(Tags are comma separated)

Link to External Site

Link to GitHub

Elevation

Show Channel Location ☐

Latitude

Longitude

Show Video ☐

☒ YouTube
 ☐ Vimeo

Video URL

Show Status ☐

Help

Channels store all the data that a ThingSpeak application collects. Each channel includes eight fields that can hold any type of data, plus three fields for location data and one for status data. Once you collect data in a channel, you can use ThingSpeak apps to analyze and visualize it.

Channel Settings

- **Percentage complete:** Calculated based on data entered into the various fields of a channel. Enter the name, description, location, URL, video, and tags to complete your channel.
- **Channel Name:** Enter a unique name for the ThingSpeak channel.
- **Description:** Enter a description of the ThingSpeak channel.
- **Field#:** Check the box to enable the field, and enter a field name. Each ThingSpeak channel can have up to 8 fields.
- **Metadata:** Enter information about channel data, including JSON, XML, or CSV data.
- **Tags:** Enter keywords that identify the channel. Separate tags with commas.
- **Link to External Site:** If you have a website that contains information about your ThingSpeak channel, specify the URL.
- **Show Channel Location:**
 - **Latitude:** Specify the latitude position in decimal degrees. For example, the latitude of the city of London is 51.5072.
 - **Longitude:** Specify the longitude position in decimal degrees. For example, the longitude of the city of London is -0.1275.
 - **Elevation:** Specify the elevation position meters. For example, the elevation of the city of London is 35.052.
- **Video URL:** If you have a YouTube™ or Vimeo® video that displays your channel information, specify the full path of the video URL.
- **Link to GitHub:** If you store your ThingSpeak code on GitHub®, specify the GitHub repository URL.

Using the Channel

You can get data into a channel from a device, website, or another ThingsSpeak channel. You can then visualize data and transform it using ThingSpeak [Apps](#).

See [Get Started with ThingSpeak™](#) for an example of measuring dew point from a weather station that acquires data from an Arduino® device.

[Learn More](#)

- Enter the details for your channel, in particular name and the set of fields. These field names should match the asset names you are going to send from FogLAMP.
- When satisfied click on *Save Channel*
- You will need the channel ID and the API key for your channel. To get this for a channel, on the *My Channels* page click on the *API Keys* box for your channel

sinusoid
Channel ID: 556345
Author: markdianomic
Access: Private

Channels ▾ Apps ▾ Support ▾ Commercial Use How to Buy MR

Private View Public View Channel Settings Sharing API Keys Data Import / Export

Write API Key

Key

[Generate New Write API Key](#)

Read API Keys

Key

Note

[Save Note](#) [Delete API Key](#)

[Add New Read API Key](#)

Help

API keys enable you to write data to a channel or read data from a private channel. API keys are auto-generated when you create a new channel.

API Keys Settings

- **Write API Key:** Use this key to write data to a channel. If you feel your key has been compromised, click **Generate New Write API Key**.
- **Read API Keys:** Use this key to allow other people to view your private channel feeds and charts. Click **Generate New Read API Key** to generate an additional read key for the channel.
- **Note:** Use this field to enter information about channel read keys. For example, add notes to keep track of users with access to your channel.

API Requests

Write a Channel Feed

```
GET https://api.thingspeak.com/update?api_key=APIKEY&field1=0
```

Read a Channel Feed

```
GET https://api.thingspeak.com/channels/556345/feeds.json?api_key=APIKEY&results=2
```

Read a Channel Field

```
GET https://api.thingspeak.com/channels/556345/fields/1.json?api_key=APIKEY&results=2
```

Read Channel Status Updates

```
GET https://api.thingspeak.com/channels/556345/status.json?api_key=APIKEY
```

[Learn More](#)

Once you have created your channel on you may create your north task on FogLAMP to send data to this channel

- Select *North* from the left hand menu bar.
- Click on the + icon in the top left
- Choose ThingSpeak from the plugin selection list
- Name your task
- Click on *Next*
- Configure the plugin

The screenshot shows the 'Review Configuration' step of the FogLAMP configuration process. At the top, there are three steps: 1. Plugin & Name, 2. Review Configuration (current), and 3. Done. The configuration fields are as follows:

- URL:** `https://api.thingspeak.com/channels`
- API Key:** (Empty field)
- Source:** `readings` (dropdown menu)
- Fields:** A JSON document editor showing the following content:


```
1 {
2   "elements": [
3     {
4       "asset": "sinusoid",
5       "reading": "sinusoid"
6     }
7   ]
8 }
```
- Channel ID:** `0`

At the bottom, there are 'Back' and 'Next' buttons.

- **URL:** The URL of the ThingSpeak server, this can usually be left as the default.
- **API Key:** The write API key from the ThingSpeak channel you created
- **Source:** Controls if readings data or FogLAMP statistics are to be send to ThingSpeak
- **Fields:** Allows you to select what fields to send to ThingSpeak. It's a JSON document that contains a single array called elements. Each item of the array is a JSON object that has two properties, asset and reading. The asset should match the asset you wish to send and the reading the data point name.
- **Channel ID:** The channel ID of your ThingSpeak Channel

- Click on *Next*
- Enable your north task and click on *Done*

8.3 FogLAMP Filter Plugins

8.3.1 Asset Filter

The *foglamp-filter-asset* is a filter that allows for assets to be included, excluded or renamed in a stream. It may be used either in *South* services or *North* tasks and is driven by a set of rules that define for each named asset what action should be taken.

Asset filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.

- Select the *asset* plugin from the list of available plugins.
- Name your asset filter.
- Click *Next* and you will be presented with the following configuration page

1 Plugin Name 2 Review Configuration

Asset rules

```

1 {
2   "rules": [
3     {
4       "asset_name" : "temperature",
5       "action" : "rename",
6       "new_asset_name" : "abient"
7     }
8   ]
9 }

```

Enabled ☒

Previous Done

- Enter the *Asset rules*
- Enable the plugin and click *Done* to activate it

Asset Rules

The asset rules are an array of JSON objects which define the asset name to which the rule is applied and an action. Actions can be one of

- **include:** The asset should be forwarded to the output of the filter
- **exclude:** The asset should not be forwarded to the output of the filter
- **rename:** Change the name of the asset. In this case a third property is included in the rule object, "new_asset_name"

In addition a *defaultAction* may be included, however this is limited to *include* and *exclude*. Any asset that does not match a specific rule will have this default action applied to them. If the default action is not given it is treated as if a default action of *include* had been set.

A typical set of rules might be

```

{
  "rules": [
    {
      "asset_name": "Random1",

```

(continues on next page)

(continued from previous page)

```

        "action": "include"
    },
    {
        "asset_name": "Random2",
        "action": "rename",
        "new_asset_name": "Random92"
    },
    {
        "asset_name": "Random3",
        "action": "exclude"
    },
    {
        "asset_name": "Random4",
        "action": "rename",
        "new_asset_name": "Random94"
    },
    {
        "asset_name": "Random5",
        "action": "exclude"
    },
    {
        "asset_name": "Random6",
        "action": "rename",
        "new_asset_name": "Random96"
    },
    {
        "asset_name": "Random7",
        "action": "include"
    }
],
"defaultAction": "include"
}

```

8.3.2 Asset Split Filter

The *foglamp-filter-asset-split* plugin is a simple filter that allows an asset to split into multiple assets. It is useful when complex assets with multiple data points within the asset have been read from a south plugin but it is easier for either upstream processing or the north system that is to receive the data for that data to be lots of assets each with a single data point.

This filter is particularly useful if you have a north system that either only accepts single data points within an asset or you have assets where the set of data points is not constant. This can cause some destinations, such as the OSIsoft PI server to have issues with changing types. Splitting the asset into multiple assets results in more assets, but each asset only has a single value associated with it and hence it maps better to the PI Server tags.

When adding an asset split filter to either the south or north data pipelines, via the *Add Application* option of the user interface, a configuration page for the filter will be shown as below;

The screenshot shows a configuration window with a progress bar at the top. Step 1 is 'Plugin Name' and Step 2 is 'Review Configuration'. The configuration area has an 'Enabled' checkbox and a 'Prefix' text field. A 'Previous' button is on the left and a 'Done' button is on the right.

The options available for configuration are

- **Prefix:** A prefix to add to the asset name that is generated for each of the data points

Each asset that passes through the filter will have each of its data points removed and added to a new asset. the name of the new asset will be the prefix given configured as above, followed by the name of the original asset and the name of the data point within the asset. The data point within the new asset will be named in the same way as the asset is named.

For example, if we had a asset called *Motor* come into the filter with two data points, *Voltage* and *Current*, two assets would flow out of the filter. If we had set prefix to be *Pump001*, the assets that would flow out would be called *Pump001Motor.Voltage* and *Pump001Motor.Current*. Each would have a single value, with the same name as the asset, and would contain the Voltage value and current value form the original asset. The original asset would not be forwarded by the filter.

8.3.3 Change Filter

The *foglamp-filter-change* filter is used to only send information about an asset onward when a particular datapoint within that asset changes by more than a configured percentage. Data is sent for a period of time before and after the change in the monitored value. The amount of data to send before and after the change is configured in milliseconds, with a value for the pre-change time and one for the post-change time.

It is possible to define a rate at which readings should be sent regardless of the monitored value changing. This provides an average of the values of the period defined, e.g. send a 1 minute average of the values every minute.

This filter only operates on a single asset, all other assets are passed through the filter unaltered.

Change filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *change* plugin from the list of available plugins.
- Name your change filter.
- Click *Next* and you will be presented with the following configuration page

1 Plugin Name 2 Review Configuration

Asset	fan1
Trigger	current
Required Change %	5
Pre-trigger time (mS)	500
Post-trigger time (mS)	500
Reduced collection rate	2
Rate Units	per hour
Enabled	<input checked="" type="checkbox"/>

Previous Done

- Enter the configuration for your change filter
 - **Asset:** The asset to monitor and control with this filter. This asset is both the asset that is used to look for changes and also the only asset whose data is affected by the triggered or non-triggered state of this filter.
 - **Trigger:** The datapoint within the asset that is used to trigger the sending of data at full rate. This datapoint may be either a numeric value or a string. If it is a string then a change of value of the defined change percentage or greater will trigger the sending of data. If the value is a string then any change in value will trigger the sending of the data.
 - **Required Change %:** The percentage change required for a numeric value change to trigger the sending of data. If this value is set to 0 then any change in the trigger value will be enough to trigger the sending of data.
 - **Pre-trigger time:** The number of milliseconds worth of data before the change that triggers the sending of data will be sent.
 - **Post-trigger time:** The number of milliseconds after a change that triggered the sending of data will be sent. If there is a subsequent change while the data is being sent then this period will be reset and the sending of data will recommence.
 - **Reduced collection rate:** The rate at which to send averages if a change does not trigger full rate data. This is defined as a number of averages for a period defined in the rateUnit, e.g. 4 per hour.
 - **Rate Units:** The unit associated with the average rate above. This may be one of “per second”, “per minute”, “per hour” or “per day”.
- Enable the change filter and click on *Done* to activate your plugin

8.3.4 CSV Writer

The plugin collects the readings from south service into csv files and compresses them when limit set per file is exceeded. The files are collected in date-wise directories where a single directory contains all the files collected on that day. The directories are rotated when the limit set is exceeded. We may collect data continuously, periodically or using a CRON string.

The screenshot shows the configuration interface for the CSV Writer plugin. It consists of a list of configuration items on the left and their corresponding input fields on the right. The fields are as follows:

- Input assets name:** An empty text input field.
- Forward data:** A checkbox that is currently unchecked.
- Destination directory:** A text input field containing the value "FOGLAMP_DATA/readings-out".
- Subdirectory name:** A text input field containing the value "south-storage".
- File type:** A dropdown menu with "csv" selected.
- Sampling rate:** A text input field with a spinner, containing the value "8000".
- Cron mode:** A dropdown menu with "continuous" selected.
- Cron period starting time:** An empty text input field.
- Recording period duration:** A text input field containing the value "100ms".
- Event repetition time:** A text input field containing the value "16mins".
- Pre event time:** A text input field containing the value "1min".
- Event duration:** A text input field containing the value "1min".

- **‘inputAssets’:** type: string default: ‘’: The names of assets (comma separated) to be stored in the csv file. All the data points of these assets will become columns in the csv file. Other assets that are not included will be forwarded. If empty all asset names will be taken.
- **‘forwardData’:** type: boolean default: false: Forces data to be forwarded upstream, normally data is written to csv and not sent to storage.
- **‘destDir’:** type: string default: **‘FOGLAMP_DATA/readings-out’**: Destination directory inside \$FOGLAMP_DATA e.g. /usr/local/foglamp/data/readings-out if FOGLAMP_DATA/ is prefixed otherwise it will be created as specified. Default is FOGLAMP_DATA/readings-out. If the path given without FOGLAMP_DATA/ then directory will be created inside \$FOGLAMP_ROOT/services/<path>, path can be recursive.
- **‘filterName’(Subdirectory name):** type: string default: **‘south-storage’**: Name of the specific sub-directory under the “dest dir” where records are to be written. Useful when we have multiple instances of csv writer filter. Just for convention use the name source-destination to indicate that the filter is applied between source and destination. For example use filterName rms-database if the filter is applied between a rms filter and sqlite database.
- **‘fileType’:** type: string default: **‘csv’**: The file type (csv or pickle) in order to store readings.
- **‘samplingRate’:** type: integer default: **‘8000’**: The number of readings per second to be stored in the csv/pickle file.
- **‘cronMode’:** type: enumeration default: **continuous**: Cron style, either periodic, continuous, or table. In continuous mode files are continuously, in periodic mode the files are collected at every given interval of time (configurable). In table mode the files are collected according to a cron string similar to cron in Unix environments.

- **‘cronPeriodStart’**: type: string default: “”: The time at which the collection should start. If empty the collection will start immediately after the first reading. If a time stamp (a sample time stamp could be 2021-04-27 09:25:35.300875+00:00) is given then the plugin will start collection when the timestamp (will use user_ts of reading, if no user_ts then ts) of a reading that has arrived becomes greater than this value. Note cronPeriodStart is useful for periodic mode and won’t be used when cronMode is table. For periodic mode we can give cronPeriodStart either in the past or in the future.
- **‘cronPeriodDuration’**: type: string default: **100ms**: Amount of time records are written per file (in ms, sec[s], min[s], hr[s], day[s], week[s]). For example if cronPeriodDuration is 3 mins and sample rate is 8000, Then each file will have $3*60*8000=1440000$ records. Note: It won’t be used when cronMode is table. The limit will be picked from cron string. It also won’t be used for periodic mode.
- **‘eventRepetitionTime’**: type: string default: **16min**: Only for periodic mode. The time after which the collection starts again. (in ms, sec[s], min[s], hr[s], day[s], week[s]).
- **‘eventPreTime’**: type: string default: **1min**: Only for periodic mode. The amount of time to consider for collection before the desired event [eg., 1min].
- **‘eventDuration’**: type: string default: **1min**: Only for periodic mode. The duration for desired event [eg., 1min].

Post event time	<input type="text" value="1min"/>
Rotate after	<input type="text" value="14days"/>
Periodic collection spec	<input type="text"/>
Add timestamp to csv data	<input type="checkbox"/>
Enable compression	<input checked="" type="checkbox"/>
Compression type	<input type="text" value="bzip2"/>
Encryption password	<input type="text"/>
Enabled	<input checked="" type="checkbox"/>

- **‘eventPostTime’**: type: string default: **1min**: Only for periodic mode. The amount of time to consider for collection after the desired event [eg., 1min]. Note the cronPeriodDuration for periodic mode is the sum eventPreTime, eventDuration and eventPostTime.
- **‘rotateAfter’**: type: string default: **10min**: Total time after which rotation will occur. (eg., 4wks). For example if rotateAfter is 7 days. Then on 12:00:00 AM of ninth day then the first directory will be deleted.
- **‘cronTabSpec’**: type: string default: “”: This parameter controls the time at which collection takes place.

Note: It is only used when cronMode is table. It is a string which consists of seven parts separated by spaces. It takes the form ‘seconds(0-59) minute(0-59) hour(0-23) day-of-month (1-31) month(1-12/names) day-of-week(0-7 or names) duration(seconds[float]’

Examples:

- use string ‘0 0,15,30,45 * * * * 60’ if you want to collect at one minute worth of data zeroth, fifteenth, thirtieth, forty fifth minute of every hour.
- use string ‘0,15,30,35 * * * * * 5’ if you want to collect at five seconds worth of data zeroth, fifteenth, thirtieth, forty fifth second of every minute.

- **‘addTimestamp’:** type: boolean default: false: Add a timestamp to each csv entry.
- **‘enableCompress’:** type: boolean default: true: Compress files after they have reached their maximum size.
- **‘compressionType’:** enumeration [‘bzip2’, ‘gzip’, ‘7za’] default bzip2: Select compression type to be used when ‘enableCompress’ is true. if 7za is selected then the files will get encrypted.
- **‘encryptPw’:** type: password default: ‘’: The password used to to encrypt files if 7za compression is selected.
- **‘enable’:** type: boolean default: ‘false’: Enable / Disable plugin operation.

Execution

Part 1: Get some south service running

For starting a south service use any of the following commands.

1. Use csvplayback

Assuming you have a csv file named vibration.csv inside FOGlamp_ROOT/data/csv_data (Can give a pattern like vib. The plugin will search for all the files starting with vib and therefore find out the file named vibration.csv). The csv file has fixed number of columns per row. Also assuming the column names are present in the first line. The plugin will rename the file with suffix .tmp after playing. Here is the curl command for that.

```
res=$(curl -sX POST http://localhost:8081/foglamp/service -d @- << EOF
↪ | jq '.
{
  "name": "My_south",
  "type": "south",
  "plugin": "csvplayback",
  "enabled": false,
  "config": {
    "assetName": {"value": "My_csv_asset"},
    "csvDirName": {"value": "FOGLAMP_DATA/csv_data"},
    "csvFileName": {"value": "vib"},
    "headerMethod": {"value": "do_not_skip"},
    "variableCols": {"value": "false"},
    "columnMethod": {"value": "pick_from_file"},
    "rowIndexForColumnNames": {"value": "0"},
    "ingestMode": {"value": "burst"},
    "sampleRate": {"value": "8000"},
    "postProcessMethod": {"value": "rename"},
    "suffixName": {"value": ".tmp"}
  }
}
EOF
)

echo $res
```

2. Use dt9837 plugin

Assuming you have connected accelerometers to the DAQ, run the following command. This command uses 4 channel data.

```
curl -sX POST http://localhost:8081/foglamp/service -d '{"name": "My_south",
↪ "type": "south", "plugin": "dt9837", "enabled": "true", "config": {"range":
↪ {"value": "BiPolar 10 Volts"}, "lowChannel": {"value": "0"}, "highChannel":
↪ {"value": "3"}}}' | jq
```


Part 2: Add the filter & attach to service

```
curl -sX POST http://localhost:8081/foglamp/filter -d '{"name":"csv_writer_
↪continuous","plugin":"csv_writer","filter_config":{"samplingRate":"8000",
↪"enable":"true","enableCompress":"true","cronTabSpec":"","addTimestamp":
↪"true","filterName":"continuous","cronMode":"continuous",
↪"cronPeriodDuration":"5min","rotateAfter":"7days","forwardData":"true"}}' |jq
curl -sX PUT 'http://localhost:8081/foglamp/filter/My_south/pipeline?allow_
↪duplicates=true&append_filter=true' -d '{"pipeline":["csv_writer_continuous
↪"]}' |jq
```

Modes

Periodic

The following command will collect data after every 16 minutes and will collect 3 minutes (pre + post + event duration) worth of data in every file. The data will get rotated after 14 + 1 = 15 days. The collection will start at 2021-07-05 11:00:00.000000+00:00 (subtract the pre time of 1 minutes). If this time is of the past the plugin will calculate the time accordingly. Note times are considered in utc. The plugin will convert the time zone to utc.

```
# assign start time to a variable.
start_time="2021-07-05 11:01:00.000000+00:00"
curl -sX POST http://localhost:8081/foglamp/filter -d '{"name":"csv_writer_periodic",
↪"plugin":"csv_writer","filter_config":{"samplingRate":"8000","enable":"true",
↪"enableCompress":"true","cronTabSpec":"","addTimestamp":"true","filterName":
↪"periodic","cronPeriodStart":"'${start_time}'',"cronMode":"periodic",
↪"eventRepetitionTime":"16min","eventDuration":"1min","eventPreTime":"1min",
↪"eventPostTime":"1min","rotateAfter":"14days","forwardData":"true"}}' |jq
curl -sX PUT 'http://localhost:8081/foglamp/filter/My_south/pipeline?allow_
↪duplicates=true&append_filter=true' -d '{"pipeline":["csv_writer_periodic"]}' |jq
```

Some sample files will be as follows:

```
foglamp@foglamp:~/usr/local/foglamp/data/readings-out/periodic/2021-07-05.d$ ls
2021-07-05-11-00-00-0000.csv.bz2
2021-07-05-11-16-00-0000.csv.bz2
2021-07-05-11-32-00-0000.csv.bz2
..
..
..
```

Continuous

The following command collects files continuously. Each files has 5 minutes worth of data.

```
curl -sX POST http://localhost:8081/foglamp/filter -d '{"name":"csv_writer_continuous
↪","plugin":"csv_writer","filter_config":{"samplingRate":"8000","enable":"true",
↪"enableCompress":"true","cronTabSpec":"","addTimestamp":"true","filterName":
↪"continuous","cronMode":"continuous","cronPeriodDuration":"5min","rotateAfter":
↪"7days","forwardData":"true"}}' |jq
curl -sX PUT 'http://localhost:8081/foglamp/filter/My_south/pipeline?allow_
↪duplicates=true&append_filter=true' -d '{"pipeline":["csv_writer_continuous"]}' |jq
```

Some sample files will be as follows:

```
foglamp@foglamp:~/usr/local/foglamp/data/readings-out/continuous/2021-05-07.d$ ls
2021-05-07-10-00-00-0000.csv.bz2
2021-05-07-10-05-00-0000.csv.bz2
2021-05-07-10-10-00-0000.csv.bz2
..
..
..
```

Cron style collection

The following command collects 5 minutes of data in every two hours.

```
curl -sX POST http://localhost:8081/foglamp/filter -d '{"name":"csv_writer_
↳ discontinuous","plugin":"csv_writer","filter_config":{"samplingRate":"8000","enable
↳ ":"true","enableCompress":"true","cronTabSpec":"0 0 0,2,4,6,8,10,12,14,16,18,20,22,
↳ * * * 300","addTimestamp":"true","filterName":"discontinuous" ,"cronMode":"table",
↳ "rotateAfter":"4weeks", "forwardData":"true"}}' |jq
curl -sX PUT 'http://localhost:8081/foglamp/filter/My_south/pipeline?allow_
↳ duplicates=true&append_filter=true' -d '{"pipeline":["csv_writer_discontinuous"]}'
↳ |jq
```

The sample files will be like

```
foglamp@foglamp:~/usr/local/foglamp/data/readings-out/discontinuous/2021-05-07.d$ ls
2021-05-07-10-00-00-0000.csv.bz2
2021-05-07-12-00-00-0000.csv.bz2
2021-05-07-14-00-00-0000.csv.bz2
..
..
..
```

Cascading CSV writer filter

We can apply multiple instances of csv writer filter. Let's say we want to apply three filters. Then we need to keep forwardData of first two filters to be true. The third filter's forwardData may or may not be true.

Consider the following example

```
# continuous
curl -sX POST http://localhost:8081/foglamp/filter -d '{"name":"csv_writer_continuous
↳ ","plugin":"csv_writer","filter_config":{"samplingRate":"8000","enable":"true",
↳ "enableCompress":"true","cronTabSpec":"","addTimestamp":"true","filterName":
↳ "continuous" ,"cronMode":"continuous", "cronPeriodDuration":"5m","rotateAfter":
↳ "7days", "forwardData":"true"}}' |jq
curl -sX PUT 'http://localhost:8081/foglamp/filter/My_south/pipeline?allow_
↳ duplicates=true&append_filter=true' -d '{"pipeline":["csv_writer_continuous"]}' |jq

# discontinuous
curl -sX POST http://localhost:8081/foglamp/filter -d '{"name":"csv_writer_
↳ discontinuous","plugin":"csv_writer","filter_config":{"samplingRate":"8000","enable
↳ ":"true","enableCompress":"true","cronTabSpec":"0 0 0,2,4,6,8,10,12,14,16,18,20,22,
↳ * * * 300","addTimestamp":"true","filterName":"discontinuous" ,"cronMode":"table",
↳ "rotateAfter":"4weeks", "forwardData":"true"}}' |jq
```

(continues on next page)

(continued from previous page)

```

curl -sX PUT 'http://localhost:8081/foglamp/filter/My_south/pipeline?allow_
↳duplicates=true&append_filter=true' -d '{"pipeline":["csv_writer_discontinuous"]}' |jq
↳|jq

# periodic
# assigning the start time to a variable.
start_time="2021-07-05 11:01:00.000000+00:00"
curl -sX POST http://localhost:8081/foglamp/filter -d '{"name":"csv_writer_periodic",
↳"plugin":"csv_writer","filter_config":{"samplingRate":"8000","enable":"true",
↳"enableCompress":"true","cronTabSpec":"","addTimestamp":"true","filterName":
↳"periodic" ,"cronPeriodStart":"","start_time":"","cronMode":"periodic",
↳"eventRepetitionTime":"16min","eventDuration":"1min","eventPreTime":"1min",
↳"eventPostTime":"1min" ,"rotateAfter":"14days", "forwardData":"true"}}' |jq
curl -sX PUT 'http://localhost:8081/foglamp/filter/My_south/pipeline?allow_
↳duplicates=true&append_filter=true' -d '{"pipeline":["csv_writer_periodic"]}' |jq

```

If forwardData of first filter is false then only the first filter will collect data.

If forwardData of first filter is true and second filter is false only first and second filter will collect data.

The forwardData of third filter may or may not be true. It is advisable to switch it off to prevent ingesting into database.

The following table sums it up.

Table 2: **Two CSV filters cascaded together**

Filter 1 forwardData	Filter 2 forwardData	Behaviour
True	True	Both filters collect data and data is ingested into database.
True	False	Both filters collect data and data is NOT ingested into database.
False	True	Only filter 1 collects data and data is NOT ingested into database.
False	False	Only filter 1 collects data and data is NOT ingested into database.

Behaviour on restart and reconfigure

After restart the collection will resume normally which means collection will begin in the same directory as it was earlier. However it may happen that the plugin was writing a file and the file is uncompressed. This uncompressed file will be compressed when the plugin will restart.

Note that the plugin won't wait for compression as it would be offloaded to some other thread for compression.

If this is a csv file and is empty it will be deleted.

It may also happen the directory name is changed inside configuration of the plugin. Then collection will begin inside different directory without deleting existing files.

On reconfigure the plugin will behave similar to restart.

How data is rotated?

The plugin picks rotateAfter config parameter and converts into days. Since each day collection has got its own directory therefore when the number of directories exceed this number the first directory will get deleted and so on. (Assuming we already had transferred these files to somewhere else before rotation.)

Note: If rotateAfter is 1week, then limit calculated will be 8. (We are talking one more day to compensate the case when collection was started at let's say at 2 PM on first day. Had we taken 7 days then this is actually 6 days data.) Now at 12:00:00 AM at the ninth day the first directory will get deleted and so on.

Decryption

If you had selected 7z for compression then you will obtain encrypted files.

Use the following command to decrypt the file.

```
7za x -p<password> <file_name>

# example 7za x -ppassword123 vibration.7z
# assuming password123 is password and file name is vibration.7z.
```

For bzip2 and gzip compression use -d flag to uncompress the file.

```
bzip2 -d <file_name>
gzip -d <file_name>
```

8.3.5 Delta Filter

The *foglamp-filter-delta* is a filter that only forwards data that changes by more than a configurable percentage. It is used to remove duplicate data values from an asset stream. The definition of duplicate however allows for some noise in the reading value by requiring a delta percentage.

By defining a minimum rate it is possible to force readings to be sent at that defined rate when there is no change in the value of the reading. Rates may be defined as per second, per minute, per hour or per day.

Delta filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *delta* plugin from the list of available plugins.
- Name your delta filter.
- Click *Next* and you will be presented with the following configuration page

1 Plugin Name 2 Review Configuration

Tolerance % 0

Minimum Rate 0

Minimum Rate Units per second

Individual Tolerances 1 { }

Enabled ☐

Previous Done

- Configure the parameters of the delta filter
 - **Tolerance %**: The percentage tolerance when comparing reading data. Only values that differ by more than this percentage will be considered as different from each other.
 - **Minimum Rate**: The minimum rate at which readings should be sent. This is the rate at which readings will appear if there is no change in value.
 - **Minimum Rate Units**: The units in which minimum rate is define (per second, minute, hour or day)
 - **Individual Tolerances**: A JSON document that can be used to define specific tolerance values for an asset. This is defines as a set of name/value pairs for those assets that should use a tolerance percentage other than the global tolerances specified above. The following example would set the tolerance for the temperature asset to 15% and for the pressure asset to 5%. All other assets would use the tolerance specified in *Tolerance %*.

```
{
  "temperature" : 15,
  "pressure" : 5
}
```

- Enable the filter and click *Done* to complete the process of adding the new filter.

8.3.6 Down Sample Filter

The *foglamp-filter-downsample* filter is a mechanism to reduce the amount of data ingested, it allows the effective data rate to be reduced by a given factor, for example to have the data rate you select a down sample factor of 2, to get a third the rate you select a down sample factor of 3. There are a number of algorithms available to select the value to be sent.

- Sample - the first value in the sample is used as the value for the sample set.
- Mean - the average value in the down sampled set is sent as the down sampled value.
- Median - the mathematical median value is sent as the down sampled value. This is the number found by sorting the sample and choosing the mid point of the sample.
- Mode - the mathematical mode value is sent as the down sampled value. This is the number that appears most often in the sample.
- Minimum - the minimum value in the sample is sent forward.
- Maximum - the maximum value in the sample is used as the sample value.

Downsample filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *downsample* plugin from the list of available plugins.
- Name your downsample filter.
- Click *Next* and you will be presented with the following configuration page

Sine South Service

1 Plugin Name 2 Review Configuration

Down Sample Factor: 2

Down Sample Algorithm: Sample


Excluded Assets:

```
1 {  
2   "exclusions": []  
3 }
```

Enabled ☐

Previous Done

- Configure your downsample filter
 - **Down Sample Factor:** The number of incoming values in each sample set.
 - **Down Sample Algorithm:** The algorithm used to determine the value for the sample.



- **Excluded Assets:** A list of assets that are excluded from the down sampling process.
- Enable your filter and click *Done*

8.3.7 Edge ML Filter Plugin

The plugin takes a image saved by south plugin named webcam media, sends that to Edge ML cluster running somewhere else. The Edge ML cluster returns a response in the form of json which contains information about detected objects, their bounding boxes and confidence score. This information is appended to the readings generated from south service.

1 Plugin Name

2 Review Configuration

Name of entries asset: edgeML

Name of inferences output asset: edgeMLInference

Kubernetes deployment:

REST URL for ML model:

Forward inputs: ☒

Remove files after inference: ☐

Enable/disable plugin: ☒

Previous Done

- **‘assetName’:** type: string default: **‘edgeML’**: Name of asset to listen on; readings have path names of images to analyze. The plugin will pick these path names to read these images.
- **‘outAssetName’:** type: string default: **‘edgeMLInference’**: Name of asset to write ML inferences on.
- **‘deploymentName’:** type: string default: **“”**: Name of Kubernetes deployment for ML model

- **‘edgeMLUrl’:** type: string default: ‘’: REST URL for ML model which analyzes images; dynamically discovered if empty.
- **‘forwardData’:** type: boolean default: **‘true’**: Forward data as well as inferences.
- **‘rmFile’:** type: string default: **‘false’**: Remove source files after inference.
- **‘enable’:** type: boolean default: **‘true’**: Enable/ Disable the plugin.

Installation

To run the plugin you must follow these prerequisites.

1. **Run the south webcam media plugin.** To run the south webcam media plugin you can either

1. Copy some images inside some directory in FOGAMP_ROOT/data. Let’s say the directory name is pics. Run the following command.

```
curl -sX POST http://localhost:8081/foglamp/service -d '{"name":  
  ↪ "My_web_cam", "type": "south", "plugin": "webcam_media", "enabled  
  ↪ ": false, "config": {"assetName": {"value": "WebcamImages"}, "imageDir  
  ↪ ": {"value": "pics"}, "mediaType": {"value": "directory"}, "fpm": {"  
  ↪ "value": "10.0"}}}' |
```

2. Connect a camera to the machine and run the following command.

```
$ v4l2-ctl --list-formats-ext --device /dev/video0  
You will see something like  
'YUYV' (YUYV 4:2:2)  
  Size: Discrete 640x480  
    Interval: Discrete 0.033s (30.000 fps)  
  Size: Discrete 720x480  
    Interval: Discrete 0.033s (30.000 fps)  
  Size: Discrete 1280x720  
    Interval: Discrete 0.033s (30.000 fps)  
  Size: Discrete 1920x1080  
    Interval: Discrete 0.067s (15.000 fps)  
    Interval: Discrete 0.033s (30.000 fps)  
  Size: Discrete 2592x1944  
    Interval: Discrete 0.067s (15.000 fps)  
  Size: Discrete 0x0
```

Now we know that the id 0 is functional. If no output then try 1,2,3 and so on.

Finally launch the plugin using

```
curl -sX POST http://localhost:8081/foglamp/service -d '{"name":  
  ↪ "My_web_cam", "type": "south", "plugin": "webcam_media", "enabled  
  ↪ ": false, "config": {"assetName": {"value": "WebcamImages"}, "imageDir  
  ↪ ": {"value": "webcam"}, "mediaType": {"value": "camera"},  
  ↪ "cameraNumber": {"value": "0"}, "fpm": {"value": "10.0"}}}' | jq
```

2. **Start the Edge ML cluster.** For starting the Edge ML cluster you should follow this [README](#) file.
3. Add the filter Edge ML.


```
curl -sX POST http://localhost:8081/foglamp/filter -d '{"name":"edge_ml_
↪filter","plugin":"edgelm1","filter_config":{"deploymentName":"edgelm1-
↪deployment","assetName":"WebcamImages","outAssetName":"DetectionResults",
↪"enable":"true", "forwardData":"true", "rmFile":"false"}}}' |jq
curl -sX PUT 'http://localhost:8081/foglamp/filter/My_web_cam/pipeline?allow_
↪duplicates=true&append_filter=true' -d '{"pipeline":["edge_ml_filter"]}' |jq
↪|jq
```

4. Finally Enable the schedule.

```
curl -sX PUT http://localhost:8081/foglamp/schedule/enable -d '{"schedule_
↪name":"My_web_cam"}' |jq
```

8.3.8 Exponential Moving Average

The *foglamp-filter-ema* plugin implements an exponential moving average across a set of data. It also forms an example of how to write a filter plugin purely in Python. Filters written in Python have the same functionality and set of entry points as any other filter.

The `plugin_info` entry point that returns details of the plugin and the default configuration

```
def plugin_info():
    """ Returns information about the plugin
    Args:
    Returns:
        dict: plugin information
    Raises:
    """
    return {
        'name': 'ema',
        'version': '2.0.0',
        'mode': "none",
        'type': 'filter',
        'interface': '1.0',
        'config': _DEFAULT_CONFIG
    }
```

The `plugin_init` entry point that initialises the plugin

```
def plugin_init(config, ingest_ref, callback):
    """ Initialise the plugin
    Args:
        config: JSON configuration document for the Filter plugin configuration_
↪category
        ingest_ref:
        callback:
    Returns:
        data: JSON object to be used in future calls to the plugin

    ...
    return data
```

The `plugin_reconfigure` entry point that is called whenever the configuration is changed

```
def plugin_reconfigure(handle, new_config):
    """ Reconfigures the plugin
```

(continues on next page)

(continued from previous page)

```
Args:
    handle: handle returned by the plugin initialisation call
    new_config: JSON object representing the new configuration category for the_
↪category
Returns:
    new_handle: new handle to be used in the future calls
    """
    global rate, datapoint
    ...
    return new_handle
```

The `plugin_shutdown` entry point called to terminate the plugin

```
def plugin_shutdown(handle):
    """ Shutdowns the plugin doing required cleanup.
    Args:
        handle: handle returned by the plugin initialisation call
    Returns:
        plugin shutdown
    """
```

And the `plugin_ingest` call that is called to do the actual data processing

```
def plugin_ingest(handle, data):
    """ Modify readings data and pass it onward
    Args:
        handle: handle returned by the plugin initialisation call
        data: readings data
    """
```

Python filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *ema* plugin from the list of available plugins.
- Name your ema filter.
- Click *Next* and you will be presented with the following configuration page

Sine South Service

1 Plugin Name 2 Review Configuration

EMA datapoint

Rate

Enabled ☐

Previous Done

- Configure your ema filter
 - **EMA datapoint:** The name of the data point to create within the asset
 - **Rate:** The rate controls the rate of the average generated, in this case it is the percentage the current value contribute to the average value.
- Enable your plugin and click *Done*

8.3.9 Event Rate Filter

The *foglamp-filter-eventrate* is a filter plugin that has been explicitly designed to work with the notification server, mechanism and the north service. It can be used to reduce the rate a reading is sent northwards until an interesting event occurs. The filter will read data at full rate from the input side and buffer data internally, sending out averages for each value over a time frame determined by the filter configuration.

The user will provide two strings and a notification asset name that will be used to form a trigger for the filter. One trigger string will set the trigger and the other will clear it. When the trigger is set then the filter will no longer average the data over the configured time period, but will instead send the full bandwidth data out of the filter.

The trigger strings are values in the event data point of the notification asset that is named in the configuration. If the string given in the trigger is found within the event data point then the trigger is deemed to have fired. String matching is case sensitive, but strings given for trigger do not need to be the entire event reason, sub string searching is used to evaluate the trigger.

The filter also allows a pre-trigger time to be configured. In this case it will buffer this much data internally and when the trigger is initially set this pre-buffered data will be sent. The pre-buffered data is discarded if the trigger is not set and the data gets to the defined age for holding pre-trigger information.

Event rate filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *eventrate* plugin from the list of available plugins.
- Name your event rate filter.
- Click *Next* and you will be presented with the following configuration page

Sine South Service

1 Plugin Name 2 Review Configuration

Event asset: event

Trigger Reason:

Terminate on: Event

Stop Reason:

Full rate time (mS): 0

Pre-trigger time (mS): 1

Reduced collection rate: 0

Rate Units: per second

Exclusions:


```
1 {
2   "exclusions": [ ]
3 }
```

Enabled: ☐

Previous Done

- Configure your event rate filter
 - **Event asset:** The asset used to trigger the full rate sending of data. This is the asset that is inserted by the plugin.
 - **Trigger Reason:** A trigger reason to set the trigger for full rate data
 - **Terminate on:** A switch to control if the end condition is a trigger or time based
 - **Step Reason:** An untrigger reason to clear the trigger for full rate data, if left blank this will simply be the trigger filter evaluating to false
 - **Full rate time (ms):** A full rate time after which the reduce rate is again active
 - **Pre-trigger time (mS):** An optional pre-trigger time expressed in milliseconds
 - **Reduced collection rate:** The nominal data rate to send data out. This defines the period over which is outgoing data item is averaged.

- **Rate Units:** This defines the units used for the above rate. This can be per second, per minute, per hour or per day.



- **Exclusions:** A set of asset names that are excluded from the rate limit processing and always sent at full rate
- Enable your plugin and click *Done*

8.3.10 Expression Filter

The *foglamp-filter-expression* allows an arbitrary mathematical expression to be applied to data values. The expression filter allows user to augment the data at the edge to include values calculate from one or more asset to be added and acted upon both within the FogLAMP system itself, but also forwarded on to the up stream systems. Calculations can range from very simply manipulates of a single value to convert ranges, e.g. a linear scale to a logarithmic scale, or can combine multiple values to create composite value. E.g. create a power reading from voltage and current or work out a value that is normalized for speed.

Expression filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *expression* plugin from the list of available plugins.
- Name your expression filter.
- Click *Next* and you will be presented with the following configuration page

- Configure the expression filter
 - **Datapoint Name:** The name of the new data point into which the new value will be stored.
 - **Expression to apply:** This is the expression that will be evaluated for each asset reading. The expression will use the data points within the reading as symbols within the asset. See [Expressions](#) below.

- Enable the plugin and click *Done* to activate your filter

Expressions

The *foglamp-filter-expression* plugin makes use of the library to do run time expression evaluation. This library provides a rich mathematical operator set, the most useful of these in the context of this plugin are;

- Logical operators (and, nand, nor, not, or, xor, xnor, mand, mor)
- Mathematical operators (+, -, *, /, %, ^)
- Functions (min, max, avg, sum, abs, ceil, floor, round, roundn, exp, log, log10, logn, pow, root, sqrt, clamp, inrange, swap)
- Trigonometry (sin, cos, tan, acos, asin, atan, atan2, cosh, cot, csc, sec, sinh, tanh, d2r, r2d, d2g, g2d, hyp)

Within the expression the data points of the asset become symbols that may be used; therefore if an asset contains values “voltage” and “current” the expression will contain those as symbols and an expression of the form

```
voltage * current
```

can be used to determine the power in Watts.

When the filter is used in an environment in which more than one asset is passing through the filter then symbols are created of the form <asset name>.<data point>. As an example if you have one asset called “electrical” that has data points of “voltage” and “current” and another asset called “speed” that has a data point called “rpm” then you can write an expression to obtain the power per 1000 RPM’s of the motor as follows;

```
(electrical.voltage * electrical.current) / (speed.rpm / 1000)
```

8.3.11 Fast Fourier Transform Filter

The *foglamp-filter-fft* filter is designed to accept some periodic data such as a sample electrical waveform, audio data or vibration data and perform a Fast Fourier Transform on that data to supply frequency data about that waveform.

Data is added as a new asset which is named as the sampled asset with “FFT” append. This FFT asset contains a set of data points that each represent the a band of frequencies, or as a frequency spectrum in a single array data point. The band information that is returned by the filter can be chosen by the user. The options available to represent each band are;

- the average in the band,
- the peak
- the RMS
- or the sum of the band.

The bands are created by dividing the frequency space into a number of equal ranges after first applying a low and high frequency filter to discard a percentage of the low and high frequency results. The bands are not created if the user instead opts to return the frequency spectrum.

If the low Pass filter is set to 15% and the high Pass filter is set to 10%, with the number of bands set to 5, the lower 15% of results are discarded and the upper 10% are discarded. The remaining 75% of readings are then divided into 5 equal bands, each of which representing 15% of the original result space. The results within each of the 15% bands are then averaged to produce a result for the frequency band.

FFT filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.

- Select the *fft* plugin from the list of available plugins.
- Name your FFT filter.
- Click *Next* and you will be presented with the following configuration page

Waveform South Service

1 Plugin Name 2 Review Configuration

Asset to analysis: wave

Result Data: average

Frequency Bands: 10

Band Prefix: Band

No. of samples per FFT: 8194

Low Frequency Reject %: 0

High Frequency Reject %: 0

Enabled: ☒

Previous Done

- Configure your FFT filter
 - **Asset to analysis:** The name of the asset that will be used as the input to the FFT algorithm.

Result Data

Frequency Bands

Band Prefix

average (selected)

peak

sum

rms

spectrum

- **Result Data:** The data that should be returned for each band. This may be one of average, sum, peak, rms or spectrum. Selecting average will return the average amplitude within the band, sum returns the sum of all amplitudes within the frequency band, peak the greatest amplitude and rms the root mean square of the amplitudes within the band. Setting the output type to be spectrum will result in the full FFT spectrum data being written. Spectrum data however can not be sent to all north destinations as it is not supported natively on all the systems FogLAMP can send data to.

- **Frequency Bands:** The number of frequency bands to divide the resultant FFT output into
- **Band Prefix:** The prefix to add to the data point names for each band in the output
- **No. of Samples per FFT:** The number of input samples to use. This must be a power of 2.
- **Low Frequency Reject %:** A percentage of low frequencies to discard, effectively reducing the range of frequencies to examine
- **High Frequency Reject %:** A percentage of high frequencies to discard, effectively reducing the range of frequencies to examine

8.3.12 FFT2 Filter

The *foglamp-filter-fft2* is a filter that applies Fourier Transform to signal data to convert it to frequency domain.

Data is returned in a single reading with a set of datapoints that each represent the average amplitude for a band of frequencies. These bands are created by dividing the frequency space into a number of equal ranges after first applying a low and high frequency filter to discard a percentage of the low and high frequency results.

E.g. if the lowPass filter is set to 15% and the highPass filter is set to 10%, with the number of bands set to 5, the lower 15% of results are discarded and the upper 10% are discarded. The remaining 75% of readings are then divided into 5 equal bands, each of which representing 15% of the original result space. The results within each of the 15% bands are then averaged to produce a result for the frequency band.

It is also possible to get the top 'k' dominant frequencies in the input signal. Also one could analyze a base frequency and its harmonics.

FFT2 filter is added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *fft2* plugin from the list of available plugins.
- Name your FFT filter instance
- Click *Next* and you will be presented with the configuration page as in the image below. Configure as required.
- Enable the filter and click *Done* to activate it

s South Service

1 Plugin Name 2 Review Configuration

Asset to analyse: vibration

Sampling rate (Hz): 8000

No. of samples to use per FFT operation: 8192

Low Frequency Reject %: 0

High Frequency Reject %: 0

Asset for generated output: fft-output

Enable summarization into bands: ☐

Number of frequency Bands: 3

Summarization method: average

Band Prefix: Band

Output spectrum data: ☐

Output dominant frequencies: ☐

Number of dominant frequencies: 3

Analyze a frequency and its harmonics: ☐

Source of base frequency value: manual

Base frequency (Hz): 1

Asset name for frequency of interest: Tachometer

Number of harmonics to analyze: 3

Forward raw signal data: ☐

Enabled: ☐

Previous Done

- **Asset to analyse:** The asset to apply the FFT filter to
- **Sampling rate (Hz):** Sampling rate of the input signal (in Hz)
- **No. of samples to use per FFT operation:** The number of input samples to use. This must be a power of 2.
- **Low Frequency Reject %:** A percentage of low frequencies to discard, effectively reducing the range of frequencies to examine
- **High Frequency Reject %:** A percentage of high frequencies to discard, effectively reducing the range of frequencies to examine
- **Asset for generated output:** The asset name to use for reading containing FFT filter output
- **Enable summarization into bands:** Whether to summarize the FFT output into bands after applying configured lowPass and highPass
- **Number of frequency Bands:** The number of frequency bands to split the resultant FFT output into
- **Summarization method:** The data that should be returned for each band. This should be one of average, sum, peak or rms. Selecting 'average' will return the average amplitude within the band, 'sum' returns the sum of all amplitudes within the frequency band, 'peak' returns the highest amplitude and 'rms' returns the root mean square of the amplitudes within the band.

- **Band Prefix:** Prefix for band datapoints
- **Output spectrum data:** Output reading would result in the full FFT spectrum data being added to output. Spectrum data however can not be sent to all north destinations as it is not supported natively on all the systems FogLAMP can send data to.
- **Output dominant frequencies:** Selecting this option results in adding “dominant frequencies” information to output reading.
- **Number of dominant frequencies:** The number of dominant frequency components to put into output of this filter.
- **Analyze a frequency and its harmonics:** Selecting this option allows analysis of a base frequency and its ‘k’ harmonics.
- **Source of base frequency value:** Selects whether the base frequency would be entered manually or it should be picked from an asset value
- **Base frequency (Hz):** Base frequency value (in Hz)
- **Asset name for frequency of interest:** If base frequency is to be picked from an asset value, that asset’s name
- **Number of harmonics to analyze:** Number of harmonics of the base frequency to be analyzed
- **Forward raw signal data:** Selects whether raw signal data should be forwarded towards storage
- **Enabled:** Whether to enable this filter

8.3.13 Flir Validity Filter

The *foglamp-filter-Flir-Validity* plugin is a simple filter that filters out unused boxes and spot temperatures in the Flir temperature data stream. The filter also allows the naming of the boxes such that the data points added to the asset will use these names rather than the default box1, box2 etc.

Adding the filter to a south plugin you will receive a configuration screen as below

AX8 South Service

1
2

Plugin Name
Review Configuration

Area Labels

1
2
3
4
5
6
7
8
9
10
11
12

{
"areas": [
"1",
"2",
"3",
"4",
"5",
"6",
"7",
"8",
"9",
"10",
]
}

Enabled
☒

Previous
Done

The JSON document *Area Labels* can be used to set the labels to use for each of the boxes and replace the min1, min2 etc. The value of this configuration option is a JSON document that has a single element called *areas* which is a JSON array. Each element in that area is the name to assign to the particular box. The default values would set the name of box1 to simply be 1, box2 to 2 etc.

If we assume we are monitoring a lathe with the camera and taking the temperature of the motor, the bearing and cutting bit using the boxes 1, 2, and 3 in the camera. We wish to rename the first box to be called *Motor*, the second box to be called *Bearing* and the third to be called *Tool*, setting an *areas* array as follows would achieve this.

```
{
  "areas" : [
    "Motor",
    "Bearing",
    "Tool",
    "4",
    "5",
    "6",
    "7",
  ]
}
```

(continues on next page)

(continued from previous page)


```
        "8",  
        "9",  
        "10"  
    ]  
}
```

Note that we do not change the boxes 4 to 10 as these are not in use and have not been defined within the area interface. Using the above configuration setting for areas will result in asset names of *minMotor*, *maxMotor* and *averageMotor* being generated for the motor temperature. Similarly the bearing temperatures would be *minBearing*, *maxBearing* and *averageBearing*. The tool would have asset names of *minTool*, *maxTool* and *averageTool*.

8.3.14 Greyscale Filter

The *foglamp-filter-greyscale* plugin is a filter that is designed to covert 24bit RGB colour images to greyscale images. any data points that contain 24 bit colour images will have those images replaced with greyscale versions of the same image. All other data points will be passed unaltered by the plugin. The plugin using the simple NTSC scaling algorithm to convert the pixels from 24 bit pixels to greyscale. This approximates the response of the three colours within the human eye.

When adding a greyscale filter to either the south service or north task, via the *Add Application* option of the user interface, a configuration page for the filter will be shown as below;



The only options are to enable and disable the filter and to define the bit depth of the greyscale image to create. The filter supports both 8 bit per pixel and 16 bit per pixel greyscale output.

8.3.15 Log Filter

The *foglamp-filter-log* plugin is a simple filter that converts data to a logarithmic scale.

When adding a scale filter to either the south service or north task, via the *Add Application* option of the user interface, a configuration page for the filter will be shown as below;

The screenshot shows a configuration interface for the *foglamp-filter-log* plugin. At the top, a progress bar indicates two steps: '1 Plugin Name' and '2 Review Configuration'. The 'Review Configuration' step is currently active. Below the progress bar, there is a configuration box containing two settings: 'Asset filter' with an empty text input field, and 'Enabled' with a checked checkbox. At the bottom of the configuration box, there are two buttons: 'Previous' and 'Done'.

The *Asset Filter* entry is a regular expression that can be used to limit the assets that the filter will effect. To change all assets leave this entry blank.

8.3.16 Metadata Filter

The *foglamp-filter-metadata* filter allows data to be added to assets within FogLAMP. Metadata takes the form of fixed data points that are added to an asset used to add context to the data. Examples of metadata might be unit of measurement information, location information or identifiers for the piece of equipment to which the measurement relates.

A metadata filter may be added to either a south service or a north task. In a south service it will be adding data for just those assets that originate in that service, in which case it probably relates to a single machine that is being monitored and would add metadata related to that machine. In a north task it causes metadata to be added to all assets that the FogLAMP is sending to the up stream system, in which case the metadata would probably related to that particular FogLAMP instance. Adding metadata in the north is particularly useful when a hierarchy of FogLAMP systems is used and an audit trail is required with the data or the individual FogLAMP systems related to some physical location information such as building, floor and/or site.

To add a metadata filter

- Click on the Applications add icon for your service or task.
- Select the *metadata* plugin from the list of available plugins.
- Name your metadata filter.
- Click *Next* and you will be presented with the following configuration page

1 Plugin Name

2 Review Configuration

Metadata to add

```

1 {
2   "floor": "Third",
3   "location": "AirIntake",
4   "unit": "Celsius",
5   "serialNo": "A73953-42492-3229"
6 }

```

Enabled ☒

- Enter your metadata in the JSON array shown. You may add multiple items in a single filter by separating them with commas. Each item takes the format of a JSON key/value pair and will be added as data points within the asset.
- Enable the filter and click on *Done* to activate it

Example Metadata

Assume we are reading the temperature of air entering a paint booth. We might want to add the location of the paint booth, the booth number, the location of the sensor in the booth and the unit of measurement. We would add the following configuration value

```

{
  "value": {
    "floor": "Third",
    "booth": 1,
    "units": "C",
    "location": "AirIntake"
  }
}

```

In above example the filter would add “floor”, “booth”, “units” and “location” data points to all the readings processed by it. Given an input to the filter of

```
{ "temperature" : 23.4 }
```

The resultant reading that would be passed onward would become

```

{ "temperature" : 23.5, "booth" : 1, "units" : "C", "floor" : "Third", "location" :
  ↪ "AirIntake" }

```

This is an example of how metadata might be added in a south service. Turning to the north now, assume we have a configuration whereby we have several sites in an organization and each site has several building. We want to monitor data about the buildings and install a FogLAMP instance in each building to collect building data. We also install a

FogLAMP instance in each site to collect the data from each individual FogLAMP instance per building, this allows us to then send the site data to the head office without having to allow each building FogLAMP to have access to the corporate network. Only the site FogLAMP needs that access. We want to label the data to say which building it came from and also which site. We can do this by adding metadata at each stage.

To the north task of a building FogLAMP, for example the “Pearson” building, we add the following metadata

```
{
  "value" : {
    "building": "Pearson"
  }
}
```

Likewise to the “Lawrence” building FogLAMP instance we add the following to the north task

```
{
  "value" : {
    "building": "Lawrence"
  }
}
```

These buildings are both in the “London” site and will send their data to the site FogLAMP instance. In this instance we have a north task that sends the data to the corporate headquarters, in this north task we add

```
{
  "value" : {
    "site": "London"
  }
}
```

If we assume we measure the power flow into each building in terms of current, and for the Pearson building we have a value of 117A at 11:02:15 and for the Lawrence building we have a value of 71.4A at 11:02:23, when the data is received at the corporate system we would see readings of

```
{ "current" : 117, "site" : "London", "building" : "Pearson" }
{ "current" : 71.4, "site" : "London", "building" : "Lawrence" }
```

By adding the data like this it gives us more flexibility, if for example we want to change the way site names are reported, or we acquire a second site in London, we only have to change the metadata in one place.

8.3.17 Image Mirror Filter Plugin

The *foglamp-filter-mirror* plugin allows the user to specify a vertical or horizontal mirror operation. All image data points will then be mirrored as configured in the plugin.

When adding a mirror filter to either the south or north, via the *Add Application* option of the user interface, a configuration page for the filter will be shown as below;

The screenshot displays the configuration interface for the *foglamp-filter-mirror* plugin. At the top, a progress bar indicates two steps: '1 Plugin Name' and '2 Review Configuration'. The main configuration area is titled 'Mirror' and includes an 'Enabled' toggle. A dropdown menu is open, showing 'Vertically' as the selected option, with 'Horizontally' as an alternative. A help icon (?) is visible in the top right corner of the configuration area. At the bottom, there are 'Previous' and 'Done' buttons.

Select the desired mirroring operation, *Vertically* or *Horizontally* to be applied to the image datapoints within the readings.

Click on the *Enabled* option and then click on *Done* to add the filter.

8.3.18 OMF Hint Filter

The *foglamp-filter-omfhint* filter allows hints to be added to assets within FogLAMP that will be used by the plugin. These hints allow for individual configuration of specific assets within the OMF plugin.

A OMF hint filter may be added to either a south service or a north task. In a south service it will be adding data for just those assets that originate in that service. In a north task it causes OMF hints to be added to all assets that the FogLAMP is sending to the upstream system, it would normally only be used in a north that was using the OMF plugin, however it could be used in a north that is sending data to another FogLAMP that then forwards to OMF.

To add an OMF hints filter:

- Click on the Applications add icon for your service or task.
- Select the *omfhint* plugin from the list of available plugins.
- Name your OMF hint filter.
- Click *Next* and you will be presented with the following configuration page

1 Plugin Name 2 Review Configuration

OMF Hint

```

1 {
2   "asset": {
3     "number": "float64"
4   }
5 }

```

Enabled ☐

Previous Done

- Enter your OMF Hints in the JSON editor shown. You may add multiple hints for multiple assets in a single filter instance. See [OMF Hint data](#).
- Enable the filter and click on *Done* to activate it.

OMF Hint data

OMF Hints comprise of an asset name which the hint applies and a JSON document that is the hint. A hint is a name/value pair: the name is the hint type and the value is the value of that hint.

The asset name may be expressed as a regular expression, in which case the hint is applied to all assets that match that regular expression.

The following hint types are currently supported by :

- *integer*: The format to use for integers, the value is a string and may be any of the PI Server supported formats; int64, int32, int16, uint64, uint32 or uint16
- *number*: The format to use for numbers, the value is a string and may be any of the PI Server supported formats; float64, float32 or float16
- *typeName*: Specify a particular type name that should be used by the plugin when it generates a type for the asset. The value of the hint is the name of the type to create.
- *tagName*: Specify a particular tag name that should be used by the plugin when it generates a tag for the asset. The value of the hint is the name of the tag to create.

- *type*: Specify a pre-existing type that should be used for the asset. In this case the value of the hint is the type to use. The type must already exist within your PI Server and must be compatible with the values within the asset.
- *datapoint*: Specifies that this hint applies to a single datapoint within the asset. The value is a JSON object that contains the name of the datapoint and one or more hints.
- *AFLocation*: Specifies a location in the OSIsoft Asset Framework for the asset. This hint is fully documented in the plugin page.

The following example shows a simple hint to set the number format to use for all numeric data within the asset names *supply*.

```
{
  "supply": {
    "number": "float32"
  }
}
```

To apply a hint to all assets, the single hint definition can be used with a regular expression.

```
{
  ".*": {
    "number": "float32"
  }
}
```

Regular expressions may also be used to select subsets of assets, in the following case only assets with the prefix OPCUA will have the hint applied.

```
{
  "OPCUA.*": {
    "number": "float32"
  }
}
```

To apply a hint to a particular data point the hint would be as follows

```
{
  "supply": {
    "datapoint": {
      "name": "frequency"
      "integer": "uint16"
    }
  }
}
```

This example sets the datapoint *frequency* within the *supply* asset to be stored in the PI server as a uint16.

Datapoint hints can be combined with asset hints

```
{
  "supply": {
    "number": "float32",
    "datapoint": {
      "name": "frequency"
      "integer": "uint16"
    }
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

In this case all numeric data except for *frequency* will be stored as float32 and *frequency* will be stored as uint16.

8.3.19 Python 2.7 Filter

The *foglamp-filter-python27* filter allows snippets of Python to be easily written that can be used as filters in FogLAMP. A similar filter exists that uses Python 3.5 syntax, the *filter*. A Python code snippet will be called with sets of asset readings as they are read or processed in a filter pipeline. The data appears in the Python code as a JSON document passed as a Python Dict type.

The user should provide a Python function whose name matches the name given to the plugin when added to the filter pipeline of the south service or north task, e.g. if you name your filter *myPython* then you should have a function named *myPython* in the code you enter. This function is send a set of readings to process and should return a set of processed readings. The returned set of readings may be empty if the filter removes all data.

A general code syntax for the function that should be provided is;

```
def myPython(readings):
    for elem in list(readings):
        ...
    return readings
```

Each element that is processed has a number of attributes that may be accessed

Attribute	Description
asset_code	The name of the asset the reading data relates to.
timestamp	The data and time FogLAMP first read this data
user_timestamp	The data and time the data for the data itself, this may differ from the timestamp above
readings	The set of readings for the asset, this is itself an object that contains a number of key/value pairs that are the data points for this reading.

In order to access an data point within the readings, for example one named *temperature*, it is a simple case of extracting the value of with *temperature* as its key.

```
def myPython(readings):
    for elem in list(readings):
        reading = elem['readings']
        temp = reading['temperature']
        ...
    return readings
```

It is possible to write your Python code such that it does not know the data point names in advance, in which case you are able to iterate over the names as follows;

```
def myPython(readings):
    for elem in list(readings):
        reading = elem['readings']
        for attribute in reading:
            value = reading[attribute]
            ...
    return readings
```

A second function may be provided by the Python plugin code to accept configuration from the plugin that can be used to modify the behavior of the Python code without the need to change the code. The configuration is a JSON document which is again passed as a Python Dict to the `set_filter_config` function in the user provided Python code. This function should be of the form

```
def set_filter_config(configuration):
    config = json.loads(configuration['config'])
    value = config['key']
    ...
    return True
```

Python27 filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *python27* plugin from the list of available plugins.
- Name your python27 filter, this should be the same name as the Python function you will provide.
- Click *Next* and you will be presented with the following configuration page

The screenshot shows the 'Review Configuration' step of the plugin configuration process. At the top, a progress bar indicates two steps: '1 Plugin Name' and '2 Review Configuration'. The main area is divided into two sections: 'Python script' and 'Configuration'. The 'Python script' section contains a code editor with the following Python code:

```
1 # generate exponential moving average
2
3 import json
4
5 # exponential moving average rate default value: include 7%
  of current value
6 rate = 0.07
7 # latest ema value
8 latest = None
9
10 # get configuration if provided.
11 # set this JSON string in configuration:
12 # {"rate":0.07}
```

Below the code editor, the filename 'ema.py' is displayed. To the right of the code editor is a file upload button labeled 'Choose Files' and a status 'No file chosen'. The 'Configuration' section contains a JSON editor with the following content:

```
1 {"rate" : 0.75}
```

At the bottom left, there is an 'Enabled' checkbox. At the bottom right, there are two buttons: 'Previous' and 'Done'.

- Enter the configuration for your python27 filter

- **Python script:** This is the script that will be executed. Initially you are unable to type in this area and must load your initial script from a file using the *Choose Files* button below the text area. Once a file has been chosen and loaded you are able to update the Python code in this page.

Note: Any changes made to the script in this screen will **not** be written back to the original file it was loaded from.

- **Configuration:** You may enter a JSON document here that will be passed to the *set_filter_config* function of your Python code.

- Enable the python27 filter and click on *Done* to activate your plugin

Example

The following example uses Python to create an exponential moving average plugin. It adds a data point called *ema* to every asset. It assumes a single data point exists within the asset, but it does not assume the name of that data point. A rate can be set for the EMA using the configuration of the plugin.

```
# generate exponential moving average

import json

# exponential moving average rate default value: include 7% of current value
rate = 0.07
# latest ema value
latest = None

# get configuration if provided.
# set this JSON string in configuration:
# {"rate":0.07}
def set_filter_config(configuration):
    global rate
    config = json.loads(configuration['config'])
    if ('rate' in config):
        rate = config['rate']
    return True

# Process a reading
def doit(reading):
    global rate, latest

    for attribute in list(reading):
        if not latest:
            latest = reading[attribute]
        else:
            latest = reading[attribute] * rate + latest * (1 - rate)
            reading[b'ema'] = latest

# process one or more readings
def ema(readings):
    for elem in list(readings):
        doit(elem['reading'])
    return readings
```

Examining the content of the Python, a few things to note are;

- The filter is given the name `ema`. This name defines the default method which will be executed, namely `ema()`.
- The function `ema` is passed 1 or more readings to process. It splits these into individual readings, and calls the function `doit` to perform the actual work.
- The function `doit` walks through each attribute in that reading, updates a global variable `latest` with the latest value of the `ema`. It then adds an `ema` attribute to the reading.
- The function `ema` returns the modified readings list which then is passed to the next filter in the pipeline.
- `set_filter_config()` is called whenever the user changes the JSON configuration in the plugin. This function will alter the global variable `rate` that is used within the function `doit`.

8.3.20 Python 3.5 Filter

The *foglamp-filter-python35* filter allows snippets of Python to be easily written that can be used as filters in FogLAMP. A similar filter exists that uses Python 2.7 syntax, the filter, however it is recommended that the *python35* filter is used by preference as it provides a more compatible interface to the rest of the FogLAMP components.

The purpose of the filter is to allow the user the flexibility to add their own processing into a pipeline that is not supported by any of the existing filters offered by FogLAMP. The philosophy of FogLAMP is to provide processing by adding a set of filters that act as a pipeline between a data source and a data sink. The data source may be the south plugin in the south service or the buffered readings data from the storage service in the case of a north service or task. The data sink is either the buffer within the storage service in the case of a south service or the north plugin that sends the data to the upstream system in the case of a north service or task. Each of the filters provides a small, focused operation on the data as it traverses the pipeline, data is passed from one filter in the pipeline to another.

The functionality provided by this filter gives the user the opportunity to write Python code that can manipulate the reading data as it flows, however that modification should follow the same guidelines and principles as followed in the filters that come supplied as part of the FogLAMP distribution or are contributed by other FogLAMP users. The overriding principles however are

- Do not duplicate existing functionality provided by existing filters.
- Keep the operations small and focused. It is better to have multiple filters each with a specific purpose than to create large, complex Python scripts.
- Do not buffer large quantities of data, this will effect the footprint of the service and also slow the data pipeline.

The Python code snippet that the user provides within this filter will be called with sets of asset readings as they are read or processed in a filter pipeline. The data appears in the Python code as a JSON document passed as a Python Dict type.

The user should provide a Python function whose name matches the name given to the plugin when added to the filter pipeline of the south service or north task, e.g. if you name your filter `myPython` then you should have a function named `myPython` in the code you enter. This function is passed a set of readings to process and should return a set of processed readings. The returned set of readings may be empty if the filter removes all data.

A general code syntax for the function that should be provided is as follows;

```
def myPython(readings):  
    for elem in list(readings):  
        ...  
    return readings
```

The script iterates over the set of readings it is passed from the previous filter or the south plugin and returns the result of processing that data. Multiple readings are passed for two reasons, one is to improve the efficiency of processing and the second is because a south service or filter may ingest or process multiple readings in a single operation.

Each element that is processed has a number of attributes that may be accessed

Attribute	Description
asset_code	The name of the asset the reading data relates to.
timestamp	The data and time FogLAMP first read this data
user_timestamp	The data and time the data for the data itself, this may differ from the timestamp above
readings	The set of readings for the asset, this is itself an object that contains a number of key/value pairs that are the data points for this reading.

In order to access an data point within the readings, for example one named *temperature*, it is a simple case of extracting the value of with *temperature* as its key.

```
def myPython(readings):
    for elem in list(readings):
        reading = elem['readings']
        temp = reading['temperature']
        ...
    return readings
```

It is possible to write your Python code such that it does not know the data point names in advance, in which case you are able to iterate over the names as follows;

```
def myPython(readings):
    for elem in list(readings):
        reading = elem['readings']
        for attribute in reading:
            value = reading[attribute]
            ...
    return readings
```

The Python script is not limited to returning the same number of readings it receives, additional readings may be added into the pipeline or readings may be removed. If the filter removes all the readings it was sent it must still return either an empty list or it may return the *None* object.

A second function may be provided by the Python plugin code to accept configuration from the plugin that can be used to modify the behavior of the Python code without the need to change the code. The configuration is a JSON document which is again passed as a Python Dict to the `set_filter_config` function in the user provided Python code. This function should be of the form

```
def set_filter_config(configuration):
    config = json.loads(configuration['config'])
    value = config['key']
    ...
    return True
```

This function is called each time the configuration of the filter is changed. The function is responsible for taking whatever actions are required to change the behavior of the Python script. The most common approach taken with the configuration function is to record the configuration information in global variables for reference by the Python script. This however is contrary to the recommendations for writing Python scripts that are embedded within plugins.

There is little choice but to use globals in this case, however precautions should be taken than minimize the risk of sharing common global variables between instances.

- Do not use common names or names that are not descriptive. E.g. avoid simply calling the variable *config*.
- Do not use multiple variables, there are other options that can be used.
 - Use a single Python DICT as reference individuals items within the DICT

- Create a Python class and use a global instance of the class

Adding Python35 Filters

Python35 filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *python35* plugin from the list of available plugins.
- Name your python35 filter, this should be the same name as the Python function you will provide.
- Click *Next* and you will be presented with the following configuration page

The screenshot shows the 'Review Configuration' page for a Python35 filter. At the top, there are two numbered steps: 1. Plugin Name and 2. Review Configuration. The main content area is divided into two sections: 'Python script' and 'Configuration'.

Python script

```

1  # generate exponential moving average
2
3  import json
4
5  # exponential moving average rate default value: include 7%
6  # of current value
7  rate = 0.07
8  # latest ema value
9  latest = None
10
11 # get configuration if provided.
12 # set this JSON string in configuration:
13 #     {"rate":0.07}

```

Below the script is a file selection button labeled 'Choose Files' and 'No file chosen'. The filename 'ema.py' is displayed below the button.

Configuration

```

1  {"rate" : 0.75}

```

At the bottom left, there is an 'Enabled' checkbox which is currently checked. At the bottom right, there are two buttons: 'Previous' and 'Done'.

- Enter the configuration for your python35 filter
 - **Python script:** This is the script that will be executed. Initially you are unable to type in this area and must load your initial script from a file using the *Choose Files* button below the text area. Once a file has been chosen and loaded you are able to update the Python code in this page.

Note: Any changes made to the script in this screen will **not** be written back to the original file it was

loaded from.

- **Configuration:** You may enter a JSON document here that will be passed to the `set_filter_config` function of your Python code.

- Enable the python35 filter and click on *Done* to activate your plugin

Example

The following example uses Python to create an exponential moving average plugin. It adds a data point called *ema* to every asset. It assumes a single data point exists within the asset, but it does not assume the name of that data point. A rate can be set for the EMA using the configuration of the plugin.

```
# generate exponential moving average

import json

# exponential moving average rate default value: include 7% of current value
rate = 0.07
# latest ema value
latest = None

# get configuration if provided.
# set this JSON string in configuration:
# {"rate":0.07}
def set_filter_config(configuration):
    global rate
    config = json.loads(configuration['config'])
    if ('rate' in config):
        rate = config['rate']
    return True

# Process a reading
def doit(reading):
    global rate, latest

    for attribute in list(reading):
        if not latest:
            latest = reading[attribute]
        else:
            latest = reading[attribute] * rate + latest * (1 - rate)
            reading[b'ema'] = latest

# process one or more readings
def ema(readings):
    for elem in list(readings):
        doit(elem['reading'])
    return readings
```

Examining the content of the Python, a few things to note are;

- The filter is given the name *ema*. This name defines the default method which will be executed, namely *ema()*.
- The function *ema* is passed 1 or more readings to process. It splits these into individual readings, and calls the function *doit* to perform the actual work.
- The function *doit* walks through each attribute in that reading, updates a global variable *latest* with the latest value of the *ema*. It then adds an *ema* attribute to the reading.

- The function `ema` returns the modified readings list which then is passed to the next filter in the pipeline.
- `set_filter_config()` is called whenever the user changes the JSON configuration in the plugin. This function will alter the global variable `rate` that is used within the function `doit`.

Scripting Guidelines

The user has the full range of Python functionality available to them within the script code they provide to this filter, however caution should be exercised as it is possible to adversely impact the functionality and performance of the FogLAMP system by misusing Python features to the detriment of FogLAMP's own features.

The overriding guidance given above should always be observed

- Do not duplicate existing functionality provided by existing filters.
- Keep the operations small and focused. It is better to have multiple filters each with a specific purpose than to create large, complex Python scripts.
- Do not buffer large quantities of data, this will effect the footprint of the service and also slow the data pipeline.

Importing Python Packages

The user is free to import whatever packages they wish in a Python script, this includes the likes of the numpy packages and other that are limited to a single instance within a Python interpreter.

Do not import packages that you do not use or are not required. This adds an extra overhead to the filter and can impact the performance of FogLAMP. Only import packages you actually need.

Python does not provide a mechanism to remove a package that has previously been imported, therefore if you import a package in your script and then update your script to no longer import the package, the package will still be in memory from the previous import. This is because we reload updated scripts without closing down as restarting the Python interpreter. This is part of the sharing of the interpreter that is needed in order to allow packages such as numpy and scipy to be used. This can lead to misleading behavior as when the service gets restarted the package will not be loaded and the script may break because it makes use of the package still.

If you remove a package import from your script and you want to be completely satisfied that the script will still run without it, then you must restart the service in which you are using the plugin. This can be done by disabling and then re-enabling the service.

Use of Global Variables

You may use global variables within your script and these globals will retain their value between invocations of the processing function. You may use global variables as a method to keep information between executions and perform such operations as trend analysis based on data seen in previous calls to the filter function.

All Python code within a single service shares the same Python interpreter and hence they also share the same set of global variables. This means you must be careful as to how you name global variables and also if you need to have multiple instances of the same filter in a single pipeline you must be aware that the global variables will be shared between them. If your filter uses global variables it is normally not recommended to have multiple instances of them in the same pipeline.

It is tempting to use this sharing of global variables as a method to share information between filters, this is not recommended as should not be used. There are several reasons for this

- It provides data coupling between filters, each filter should be independent of each other filter.
- One of the filters sharing global variables may be disabled by the user with unexpected consequences.

- Filter order may be changed, resulting in data that is expected by a later filter in the chain not being available.
- Intervening filters may add or remove readings resulting in the data in the global variables not referring to the same reading, or set of readings that it was intended to reference.

If you no wish one filter to pass data onto a later filter in the pipeline this is best done by adding data to the reading, as an extra data point. This data point can then be removed by the later filter. An example of this is the way FogLAMP adds OMF hints to readings that are processed and removed by the OMF north plugin.

For example let us assume we have calculated some value *delta* that we wish to pass to a later filter, we can add this as a data point to our reading which we will call *_hintDelta*.

```
def myPython(readings):
    for elem in list(readings):
        reading = elem['readings']
        ...
        reading['_hintDelta'] = delta
        ...
    return readings
```

This is far better than using a global as it is attached to the reading to which it refers and will remain attached to that reading until it is removed. It also means that it is independent of the number of readings that are processed per call, and resilient to readings being added or removed from the stream.

The name chosen for this data point in the example above has no significance, however it is good practice to choose a name that is unlikely to occur in the data normally and portrays the usage or meaning of the data.

File IO Operations

It is possible to make use of file operations within a Python35 filter function, however it is not recommended for production use for the following reasons;

- Pipelines may be moved to other hosts where files may not be accessible.
- Permissions may change dependent upon how FogLAMP systems are deployed in the various different scenarios.
- Edge devices may also not have large, high performance storage available, resulting in performance issues for FogLAMP or failure due to lack of space.
- FogLAMP is designed to be managed solely via the FogLAMP API and applications that use the API. There is no facility within that API to manage arbitrary files within the filesystem.

It is common to make use of files during development of a script to write information to in order to aid development and debugging, however this should be removed, along with associated imports of packages required to perform the file IO, when a filter is put into production.

Threads within Python

It is tempting to use threads within Python to perform background activity or to allow processing of data sets in parallel, however there is an issue with threading in Python, the Python Global Interpreter Lock or GIL. The GIL prevents two Python statements from being executed within the same interpreter by two threads simultaneously. Because we use a single interpreter for all Python code running in each service within FogLAMP, if a Python thread is created that performs CPU intensive work within it, we block all other Python code from running within that FogLAMP service.

We therefore avoid using Python threads within FogLAMP as a means to run CPU intensive tasks, only using Python threads to perform IO intensive tasks, using the asyncio mechanism of Python 3.5.3 or later. In older versions of FogLAMP we used multiple interpreters, one per filter, in order to workaround this issue, however that had the side

effect that a number of popular Python packages, such as *numpy*, *pandas* and *scipy*, could not be used as they can not support multiple interpreters within the same address space. It was decided that the need to use these packages was greater than the need to support multiple interpreters and hence we have a single interpreter per service in order to allow the use of these packages.

Interaction with External Systems

Interaction with external systems, using network connections or any form of blocking communication should be avoided in a filter. Any blocking operation will cause data to be blocked in the pipeline and risks either large queues of data accumulating in the case of asynchronous south plugins or data begin missed in the case of polled plugins.


Scripting Errors

If an error occurs in the plugin or Python script, including script coding errors and Python exception, details will be logged to the error log and data will not flow through the pipeline to the next filter or into the storage service.

Warnings raised will also be logged to the error log but will not cause data to cease flowing through the pipeline.

To view the error log you may examine the file directly on your host machine, for example `/var/log/syslog` on a Ubuntu host, however it is also possible to view the error logs specific to FogLAMP from the FogLAMP user interface. Select the *System* option under *Logs* in the left hand menu pane. You may then filter the logs for a specific service to see only those logs that refer to the service which uses the filter you are interested in.

Alternatively if you open the dialog for the service in the *South* or *North* menu items you will see two icons displayed in the bottom left corner of the dialog that lets you alter the configuration of the service.



The left most icon, with the ? in a circle, allows you to view the documentation for the plugin, the right most icon, which looks like a page of text with a corner folded over, will open the log view page filtered to view the service.

Error Messages & Warnings

The following are some errors you may see within the log with some description of the cause and remedy for the error.

Unable to obtain a reference to the asset tracker. Changes will not be tracked The service is unable to obtain the required reference to the asset tracker within FogLAMP. Data will continue to flow through the pipeline, but there will not be any trace of the assets that have been modified by this plugin within the pipeline.

The return type of the python35 filter function should be a list of readings. The python script has returned an incorrect data type. The return value of the script should be a list of readings

Badly formed reading in list returned by the Python script One or more of the readings in the list returned by the Python script is an incorrectly formed reading object.

Each element returned by the script must be a Python DICT The list returned by the Python script contains an element that is not a DICT and therefore can not be a valid reading.

Badly formed reading in list returned by the Python script: Reading has no asset code element One or more of the readings that is returned in the list from the script is missing the *asset_code* key. This item is the name of the asset to which the reading refers.

Badly formed reading in list returned by the Python script: Reading is missing the reading element which should contain the data One or more of the readings that is returned in the list from the script is missing the *reading* DICT that contains the actual data.

Badly formed reading in list returned by the Python script: The reading element in the python Reading is of an incorrect type, i One or more of the readings that is returned in the list from the script has an item with a key of *reading* which is not a Python Dict. This item should always be a DICT and contains the data values as key/value pairs.

Badly formed reading in list returned by the Python script: Unable to parse the asset code value. Asset codes should be a string One or more of the readings that is returned in the list from the script has an item with a key of *asset_code* whose value is not a string.

8.3.21 Rate Filter

The *foglamp-filter-rate* plugin that can be used to reduce the rate a reading is stored until an interesting event occurs. The filter will read data at full rate from the input side and buffer data internally, sending out averages for each value over a time frame determined by the filter configuration.

The user can provide either one or two simple expressions that will be evaluated to form a trigger for the filter. One expressions will set the trigger and the other will clear it. When the trigger is set then the filter will no longer average the data over the configured time period, but will instead send the full bandwidth data out of the filter. If the second expression, the one that clears the full rate sending of data is omitted then the full rate is cleared as soon as the trigger expression returns false. Alternatively the filter can be configured to clear the sending of full rate data after a fixed time.

The filter also allows a pre-trigger time to be configured. In this case it will buffer this much data internally and when the trigger is initially set this pre-buffered data will be sent. The pre-buffered data is discarded if the trigger is not set and the data gets to the defined age for holding pre-trigger information.

Rate filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *rate* plugin from the list of available plugins.

- Name your rate filter.
- Click *Next* and you will be presented with the following configuration page

1 Plugin Name

2 Review Configuration

Trigger expression:

Terminate on:

End Expression:

Full rate time (mS):

Pre-trigger time (mS):

Reduced collection rate:

Rate Units:

Exclusions:

```
1 {
2   "exclusions": [ "speed" ]
3 }
```

Enabled: ☒

- Configure your rate filter
 - **Trigger Expression:** An expression to set the trigger for full rate data
 - **Terminate ON:** The mechanism to stop full rate forwarding, this may be another expression or a time window


Terminate on

Expression

Time

- **End Expression:** An expression to clear the trigger for full rate data, if left blank this will simply be the trigger filter evaluating to false
- **Full rate time (ms):** The time window, in milliseconds to forward data at the full rate

- **Pre-trigger time (ms):** An optional pre-trigger time expressed in milliseconds
- **Reduced collection rate:** The nominal data rate to send data out. This defines the period over which is outgoing data item is averaged.
- **Rate Units:** The units that the reduced collection rate is expressed in; per second, minute, hour or day



- **Exclusions:** A set of asset names that are excluded from the rate limit processing and always sent at full rate

- Enable your filter and click *Done*

For example if the filter is working with a SensorTag and it reads the tag data at 10ms intervals but we only wish to send 1 second averages under normal circumstances. However if the X axis acceleration exceed 1.5g then we want to send full bandwidth data until the X axis acceleration drops to less than 0.2g, and we also want to see the data for the 1 second before the acceleration hit this peak the configuration might be:

- **Nominal Data Rate:** 1, data rate unit “per second”
- **Trigger set expression:** $X > 1.5$
- **Trigger clear expression:** $X < 0.2$
- **Pre-trigger time (mS):** 1000

The trigger expression uses the same expression mechanism, as the , and plugins

Expression may contain any of the following. . .

- Mathematical operators (+, -, *, /, %, ^)
- Functions (min, max, avg, sum, abs, ceil, floor, round, roundn, exp, log, log10, logn, pow, root, sqrt, clamp, inrange, swap)
- Trigonometry (sin, cos, tan, acos, asin, atan, atan2, cosh, cot, csc, sec, sinh, tanh, d2r, r2d, d2g, g2d, hyp)
- Equalities & Inequalities (=, ==, <>, !=, <, <=, >, >=)
- Logical operators (and, nand, nor, not, or, xor, xnor, mand, mor)

Note: This plugin is designed to work with streams with a single asset in the stream, there is no mechanism in the expression syntax to support multiple asset names.

8.3.22 Rename Filter

The *foglamp-filter-rename* filter that can be used to modify the name of an asset, datapoint or both. It may be used either in *South* services or *North* services or *North* tasks.

To add a Rename filter

- Click on the Applications add icon for your service or task.
- Select the *rename* plugin from the list of available plugins.
- Name your Rename filter.
- Click *Next* and you will be presented with the following configuration page
- Configure the plugin

- **Operation:** Search and replace operation be performed on asset name, datapoint name or both
- **Find:** A regular expression to match for the given operation
- **Replace With:** A substitution string to replace the matched text with
- Enable the filter and click on *Done* to activate it

Example

The simplest following example perform on given below reading object

```
{
  "readings": {
    "sinusoid": -0.978147601,
    "a": {
      "sinusoid": "2.0"
    }
  },
  "asset": "sinusoid",
  "id": "a1bedea3-8d80-47e8-b256-63370ccfce5b",
  "ts": "2021-06-28 14:03:22.106562+00:00",
}
```

(continues on next page)

(continued from previous page)

```

    "user_ts": "2021-06-28 14:03:22.106435+00:00"
  }

```

1. To replace an asset apply a configuration would be as follows

- Operation : asset
- Find : sinusoid
- Replace With : sin

Output

```

{
  "readings": {
    "sinusoid": -0.978147601,
    "a": {
      "sinusoid": 2.0
    }
  },
  "asset": "sin",
  "id": "albedea3-8d80-47e8-b256-63370ccfce5b",
  "ts": "2021-06-28 14:03:22.106562+00:00",
  "user_ts": "2021-06-28 14:03:22.106435+00:00"
}

```

2. To replace a datapoint apply a configuration would be as follows

- Operation : datapoint
- Find : sinusoid
- Replace With : sin

Output

```

{
  "readings": {
    "sin": -0.978147601,
    "a": {
      "sin": 2.0
    }
  },
  "asset": "sinusoid",
  "id": "albedea3-8d80-47e8-b256-63370ccfce5b",
  "ts": "2021-06-28 14:03:22.106562+00:00",
  "user_ts": "2021-06-28 14:03:22.106435+00:00"
}

```

3. To replace both asset and datapoint apply a configuration would be as follows

- Operation : both
- Find : sinusoid
- Replace With : sin

Output

```
{
  "readings": {
    "sin": -0.978147601,
    "a": {
      "sin": 2.0
    }
  },
  "asset": "sin",
  "id": "albedea3-8d80-47e8-b256-63370ccfce5b",
  "ts": "2021-06-28 14:03:22.106562+00:00",
  "user_ts": "2021-06-28 14:03:22.106435+00:00"
}
```

8.3.23 Replace Filter

The *foglamp-filter-replace* is a filter that allows an be used to replace all occurrence of a set of characters with a single replacement character. This can be used to change reserved characters in the names of assets and datapoints.

The screenshot displays the configuration interface for the *foglamp-filter-replace* plugin. It is divided into two main steps: **1. Plugin Name** and **2. Review Configuration**. The configuration form contains the following fields:

- Replace:** A text input field containing the regular expression `[\:\?]`. A help icon (?) is visible to the right of the field.
- With:** A text input field containing the replacement character `-`.
- Enabled:** A checkbox that is currently unchecked.

At the bottom of the interface, there are two buttons: **Previous** (disabled) and **Done** (active).

- **Replace:** The set of reserved characters to be replaced.
- **With:** The character to replace each occurrence of the above characters with

8.3.24 Root Mean Squared (RMS) Filter

The *foglamp-filter-rms* filter is designed to accept some periodic data such as a sample electrical waveform, audio data or vibration data and perform a Root Mean Squared, *RMS* operation on that data to supply power of the waveform. The filter can also return the *peak to peak* amplitude of the waveform over the sampled period and the *crest* factor of the waveform.

Note: peak values may be less than individual values of the input if the asset value does not fall to or below zero. Where a data value swings between negative and positive values then the peak value will be greater than the maximum value in the data stream. For example if the minimum value of a data point in the sample set is 0.3 and the maximum

is 3.4 then the peak value will be 3.1. If the maximum value is 2.4 and the minimum is zero then the peak will be 2.4. If the maximum value is 1.7 and the minimum is -0.5 then the peak value will be 2.2.

RMS, also known as the quadratic mean, is defined as the square root of the mean square (the arithmetic mean of the squares of a set of numbers).

Peak to peak, is the difference between the smallest value in the sampled data and the highest, this give the maximum amplitude variation during the period sampled.

Crest factor is a parameter of a waveform, showing the ratio of peak values to the effective value. In other words, crest factor indicates how extreme the peaks are in a waveform. Crest factor 1 indicates no peaks, such as direct current or a square wave. Higher crest factors indicate peaks, for example sound waves tend to have high crest factors.

The user may also choose to include or not the raw data that is used to calculate the RMS values via a switch in the configuration.

Where a data stream has multiple assets within it the RMS filter may be limited to work only on those assets whose name matches a regular expression given in the configuration of the filter. The default for this expression is `.*`, i.e. all assets are processed.

RMS filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *rms* plugin from the list of available plugins.
- Name your RMS filter.
- Click *Next* and you will be presented with the following configuration page

Waveform South Service

1 Plugin Name 2 Review Configuration

Sample size: 10

RMS Asset name: %a RMS

Include Peak Values: ☐

Include Crest Values: ☐

Include Raw Data: ☐

Asset filter: .*

Enabled: ☐

Previous Done

- Configure your RMS filter
 - **Sample size:** The number of data samples to perform a calculation over.

- **RMS Asset name:** The asset name to use to output the RMS values. “%a” will be replaced with the original asset name.
- **Include Peak Values:** A switch to include peak to peak measurements for the same data set as the RMS measurement.
- **Include Crest Values:** A switch to include crest measurements for the same data set as the RMS measurement.
- **Include Raw Data:** A switch to include the raw input data in the output.
- **Asset Filter:** A regular expression to limit the asset names on which this filter operations. Useful when multiple assets appear in the input data stream as it allows data which is not part of the periodic function that is being examined to be excluded.

8.3.25 Image Rotation Filter Plugin

The *foglamp-filter-rotate* plugin allows the user to specify an angle of rotation of either 90, 180 or 270 degrees clockwise. All image data points will then be rotated by that angle.

When adding a rotate filter to either the south or north, via the *Add Application* option of the user interface, a configuration page for the filter will be shown as below;

Select the desired angle of clockwise rotation to be applied to the image datapoints within the readings.

Click on the *Enabled* option and then click on *Done* to add the filter.

8.3.26 Scale Filter

The *foglamp-filter-scale* plugin is a simple filter that allows a scale factor and an offset to be applied to numerical data. Its primary uses are for adjusting values to match different measurement scales, for example converting temperatures from Centigrade to Fahrenheit or when a sensor reports a value in non-base units, e.g. 1/10th of a degree.

When adding a scale filter to either the south service or north task, via the *Add Application* option of the user interface, a configuration page for the filter will be shown as below;

The configuration options supported by the scale filter are detailed in the table below

Setting	Description
Scale Factor	The scale factor to multiply the numeric values by
Constant Offset	A constant to add to all numeric values after applying the scale
Asset filter	This is useful when applying the filter in the north, it allows the filter to be applied only to those assets that match the regular expression given. If left blank then the filter is applied to all assets/

8.3.27 Scale Set Filter

The *foglamp-filter-scale-set* plugin is a filter that allows a scale factor and an offset to be applied to numerical data where an asset has multiple data points. It is very similar to the filter, which allows a single scale and offset to be applied to all assets and data points. It's primary uses are for adjusting values to match different measurement scales, for example converting temperatures from Centigrade to Fahrenheit or when a sensor reports a value in non-base units, e.g. 1/10th of a degree.

Scale set filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *scale-set* plugin from the list of available plugins.
- Name your scale-set filter.
- Click *Next* and you will be presented with the following configuration page

- Enter the configuration for your change filter
 - **Scale factors:** A JSON document that defines a set of factors to apply. It is an array of JSON objects that define the scale factor and offset, a regular expression that is matched against the asset name and another that matches the data point name within the asset.

Name	Description
asset	A regular expression to match against the asset name. The scale factor is only applied to assets whose name matches this regular expression.
data-point	A regular expression to match against the data point name within a matching asset. The scale factor is only applied to assets whose name matches this regular expression.
scale	The scale factor to apply to the numeric data.
off-set	The offset to add to the matching numeric data.

- Enable the scale-set filter and click on *Done* to activate your plugin

Example

In the following example we have an asset whose name is *environment* which contains two data points; *temperature* and *humidity*. We wish to allow two different scale factors and offsets to these two data points whilst not affecting assets of any other name in the data stream. We can accomplish this by using the following JSON document in the plugin configuration;

```
{
  "factors" : [
    {
```

(continues on next page)

(continued from previous page)

```

    "asset"      : "environment",
    "datapoint"  : "temperature",
    "scale"      : 1.8,
    "offset"     : 32
  },
  {
    "asset"      : "environment",
    "datapoint"  : "humidity",
    "scale"      : 0.1,
    "offset"     : 0
  }
]
}

```

If instead we had multiple assets that contain *temperature* and *humidity* we can accomplish the same transformation on all these assets, whilst not affecting any other assets, by changing the *asset* regular expression to something that matches more asset names;

```

{
  "factors" : [
    {
      "asset"      : ".*",
      "datapoint"  : "temperature",
      "scale"      : 1.8,
      "offset"     : 32
    },
    {
      "asset"      : ".*",
      "datapoint"  : "humidity",
      "scale"      : 0.1,
      "offset"     : 0
    }
  ]
}

```


8.3.28 Sigma Data Cleansing Filter

The *foglamp-filter-sigmacleanse* filter is designed to cleanse data in a stream by removing outliers from the data stream. The method used to remove these outliers is to build an average and standard deviation for the data over time and remove any data that differs by more than a certain factor of the standard deviation from that average.

The plugin is designed to be used in situations when a sensor or item of equipment produces occasional anomalous results, these will be removed from the data passed onward within the system to provide a cleaner data stream. Care should be taken however that these values that are removed do represent sensor anomalies and are not the result of problems with the condition that is being monitored. If a sensor produces a high percentage of anomalous results then it should be considered for replacement.

In order to monitor the anomalous rates the plugin can optionally produce an hourly statistics report that will show the number of readings that have been forwarded as good and the number that have been discarded.

The method used to determine if a value is anomalous is based on the premise that data from a given sensor will follow a normal distribution from the mean value that is sampled over time. The probability of a value being valid reduces as the value differs more greatly from the mean value. This gives rise to the classical bell shaped distribution of values as shown below.



It can be seen from the diagram above how the probability drops as the values move away from the mean, the sigma values here are the standard deviations observed for good data samples. Outlier values that are discarded do not contribute to the calculation of the standard deviation.

To add a sigma cleansing filter to your service:

- Click on the Applications add icon for your service or task.
- Select the *sigmacleanse* plugin from the list of available plugins.
- Name your cleansing filter.
- Click *Next* and you will be presented with the following configuration page

Progress: 1 Plugin Name ————— 2 Review Configuration

Sample Size (hours)	<input type="text" value="1"/>
Sigma	<input type="text" value="3"/>
Statistics Asset	<input type="text"/>
Enabled	<input type="checkbox"/>

Previous Done

- Configure your sigma cleanse filter
 - **Sample Size:** The number of hours over which an initial mean and standard deviation is built before any cleansing commences

- **Sigma:** The factor to apply to the standard deviation, the default is 3. Any value that differs from the mean by more than 3 * sigma will be removed.
- **Statistics Asset:** If this is not empty a statistics asset will be added every hour that details the number of readings that have been forwarded by the filter and the number removed. The name is that asset matches the value added here.
- Enable your filter and click *Done*

8.3.29 Simple Python Filter

The *foglamp-filter-simple-python* plugin allows very simple Python code to be used as a filter. A user may effectively write expressions in Python and have them execute in a filter.

The data is available within your Python code as a variable, named *reading* for each asset. You may access each data point within the asset by indexing the reading with the data point name. For example if your asset has two data points, *voltage* and *current*, then you would access these two values as

```
voltage = reading[b'voltage']
current = reading[b'current']
```

Using this type of filter it is possible to modify values of data points within an asset, remove data points in an asset or add new data points to an asset. It is **not** possible to remove assets or add new assets. The filter uses a Python 3 run time environment, therefore Python 3 syntax should be used.

The following examples show how to filter the readings data,

- Change datapoint value

```
reading[b'point_1'] = reading[b'point_1'] * 2 + 15
```

- Create a new datapoint while filtering

```
reading[b'temp_fahr'] = reading[b'temperature'] * 9 / 5 + 32
```

- Generate an exponential moving average (ema)

In this case we need to parse some data while filtering current data set the filter receives in input. A global 'user_data' empty dictionary is available to the Python interpreter and key values can be easily added. This illustrates the ability to maintain state within your filter.

```
global user_data
if not user_data:
    user_data['latest'] = None
for attribute in list(reading):
    if not user_data['latest']:
        user_data['latest'] = reading[attribute]
    user_data['latest'] = reading[attribute] * 0.07 + user_data['latest'] *
    ↪ (1 - 0.07)
    reading[b'ema'] = user_data['latest']
```

Simple Python filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *simple-python* plugin from the list of available plugins.
- Name your python filter.
- Click *Next* and you will be presented with the following configuration page

Sine South Service

1 Plugin Name 2 Review Configuration

Python code

```
1 reading[b'sinusoid'] = reading[b'sinusoid'] * 2 + 15
```

Enabled ☒

Previous Done

- Configure your simple Python filter
 - **Python Code:** Enter the code required for your filter.
- Enable your filter and click *Done*

8.3.30 Specgram Filter

The *foglamp-filter-specgram* is a filter that generates spectrogram images from incoming signal data based on defined configuration.

Specgram filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *specgram* plugin from the list of available plugins.
- Name your specgram filter instance
- Click *Next* and you will be presented with the configuration page as in the image below. Configure as required.
- Enable the filter and click *Done* to activate it

s South Service

Enabled	<input type="checkbox"/>
Log level	INFO
Asset from which vibration data is read	vibration
Output asset indicating small shoot discharge length	smallShootDischarge
Comma separated TS and Data channel(s)	user_ts,Channel0,Channel1,Channel2,Channel3
Sampling rate (per sec)	8000
Forward data	<input type="checkbox"/>
Vibration data evaluation interval	1hr
Delete spectrograms images older than these many days	30
Destination directory	/usr/local/foglamp/data/specgrams
Periodic event UTC timestamp	
Event repetition time	16min
Pre event time	1min
Event duration	1min
Post event time	1min

- **Enabled:** Whether this filter is to be enabled
- **Log level:** Logging level for debug/error information
- **Asset from which vibration data is read:** Asset name containing the signal data to create spectrograms for
- **Output asset indicating small shoot discharge length:** The asset name under which small shoot discharge width and timing is reported
- **Comma separated TS and Data channel(s):** comma separated list of columns for timestamp (mandatory) and data channels to be processed
- **Sampling rate (per sec):** The sampling rate of the incoming signal data
- **Forward data:** Whether to forward original vibration data down the filter chain
- **Vibration data evaluation interval:** Time interval after which to generate spectrogram images
- **Delete spectrogram images older than these many days:** Number of days after which older spectrogram images may be discarded
- **Destination directory:** Destination directory to store spectrogram images organized into date-wise directories
- **Periodic event UTC timestamp:** UTC Timestamp (potentially in the past) when a periodic event of interest happened/would happen
- **Event repetition time:** The time after which the desired event repeats
- **Pre event time:** The amount of time to consider for collection before the desired event [eg., 1min]
- **Event duration:** The duration for desired event [eg., 1min]
- **Post event time:** The amount of time to consider for collection after the desired event [eg., 1min]

8.3.31 Statistics Filter

The *foglamp-filter-statistics* filter is designed to accept data from one or more asset and produces statistics over specified time intervals, for example produce the mean, standard deviation and variance for 100 milliseconds samples of the data. The statistics that can be produced are;

- mean - the average of all the values in the time period calculated by adding up all the values and dividing by the number of values.
- mode - the number that appears most often in the time period.
- median - the median is found by sorting all the values in the time period and then choosing the middle number in this sorted set
- minimum - the minimum value that appears within the time period
- maximum - the maximum value that appears within the time period
- standard deviation - the standard deviation measures the spread of the numbers above and below the mean value
- variance - the variance is the average of the squared differences from the mean value calculated over the time period

Statistics filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *statistics* plugin from the list of available plugins.
- Name your statistics filter.
- Click *Next* and you will be presented with the following configuration page

Sine South Service

1 Plugin Name 2 Review Configuration

Sample Size (mS) 100

Mean	<input checked="" type="checkbox"/>
Mode	<input checked="" type="checkbox"/>
Median	<input checked="" type="checkbox"/>
Minimum	<input checked="" type="checkbox"/>
Maximum	<input checked="" type="checkbox"/>
Standard Deviation	<input checked="" type="checkbox"/>
Variance	<input checked="" type="checkbox"/>
Enabled	<input checked="" type="checkbox"/>

Previous Done

- Configure your statistics filter
 - **Mean:** A toggle that controls inclusion of the mean value
 - **Mode:** A toggle that controls inclusion of the mode value
 - **Median:** A toggle that controls inclusion of the median value
 - **Minimum:** A toggle that controls inclusion of the minimum value
 - **Maximum:** A toggle that controls inclusion of the maximum value
 - **Standard Deviation:** A toggle that controls inclusion of the standard deviation value
 - **Variance:** A toggle that controls inclusion of the variance value
- Enable your filter and click *Done*

8.3.32 Threshold Filter

The *foglamp-filter-threshold* plugin is a filter that is used to control the forwarding of data within FogLAMP. Its use is to only allow data to be stored or forwarded if a condition about that data is true. This can save storage or network bandwidth by eliminating data that is of no interest.

The filter uses an expression, that is entered by the user, to evaluate if data should be forwarded, if that expression evaluates to true then the data is forwarded, in the case of a south service this would be to the FogLAMP storage. In the case of a north task this would be to the upstream system.

Note: If the threshold filter is part of a chain of filters and the data is not forwarded by the threshold filter, i.e. the expression evaluates to false, then the following filters will not receive the data.

If an asset in the case of a south service, or data stream in the case of a north task, has other data points or assets that are not part of the expression, then they too are subject to the threshold. If the expression evaluates to false then no assets will be forwarded on that stream. This allows a single value to control the forwarding of data.

Another example use might be to have two north streams, one that uses a high cost, link to send data when some condition that requires close monitoring occurs and the other that is used to send data by a lower cost mechanism when normal operating conditions apply.

E.g. We have a temperature critical process, when the temperature is above 80 degrees it must be closely monitored. We use a high cost link to send data north wards in this case. We would have a north task setup that has the threshold filter with the condition:

```
temperature >= 80
```

We then have a second, lower cost link with a north task using the threshold filter with the condition:

```
temperature < 80
```

This way all data is sent once, but data is sent in an expedited fashion if the temperature is above the 80 degree threshold.

Threshold filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *threshold* plugin from the list of available plugins.
- Name your threshold filter.
- Click *Next* and you will be presented with the following configuration page

- Enter the expression to control forwarding in the box labeled *Expression*
- Enable the filter and click on *Done* to activate it

Expressions

The *foglamp-filter-threshold* plugin makes use of the library to do run time expression evaluation. This library provides a rich mathematical operator set, the most useful of these in the context of this plugin are;

- Comparison operators (=, ==, <>, !=, <, <=, >, >=)
- Logical operators (and, nand, nor, not, or, xor, xnor, mand, mor)
- Mathematical operators (+, -, *, /, %, ^)
- Functions (min, max, avg, sum, abs, ceil, floor, round, roundn, exp, log, log10, logn, pow, root, sqrt, clamp, inrange, swap)
- Trigonometry (sin, cos, tan, acos, asin, atan, atan2, cosh, cot, csc, sec, sinh, tanh, d2r, r2d, d2g, g2d, hyp)

Within the expression the data points of the asset become symbols that may be used; therefore if an asset contains values “voltage” and “current” the expression will contain those as symbols and an expression of the form

```
voltage * current > 1000
```

can be used to determine if power (voltage * current) is greater than 1kW.

8.3.33 Vibration Features Filter

The *foglamp-filter-vibration_features* filter collects readings for configured observation interval, then calculates statistics on these readings and puts these statistics into a new reading.

- mean - the average of all the values in the time period calculated by adding up all the values and dividing by the number of values.
- median - the median is found by sorting all the values in the time period and then choosing the middle number in this sorted set
- standard deviation - the standard deviation measures the spread of the numbers above and below the mean value
- variance - the variance is the average of the squared differences from the mean value calculated over the time period

- RMS - the root mean squared of the waveform
- kurtosis - is a measure of the combined sizes of the two tails. It measures the amount of probability in the tails.

Vibration feature filters are added in the same way as any other filters.

- Click on the Applications add icon for your service or task.
- Select the *vibration_features* plugin from the list of available plugins.
- Name your vibration feature filter.
- Click *Next* and you will be presented with the following configuration page

The screenshot shows a configuration window titled "Sine South Service". At the top, a progress bar indicates two steps: "1 Plugin Name" and "2 Review Configuration". The "Review Configuration" step is active. Below the progress bar, there is a form with three fields: "Asset name" (containing "VibrationFeature"), "Observation interval (ms)" (containing "10"), and "Enabled" (a checked checkbox). At the bottom of the form, there are two buttons: "Previous" and "Done".

- Configure your vibration filter
 - **Asset name:** The name of the asset to create. This is the asset that will hold the vibration feature data.
 - **Observation interval (ms):** The interval over which the statistics are compiled.
- Enable your filter and click *Done*

8.4 FogLAMP Notification Rule Plugins

8.4.1 Threshold Rule

The threshold rule is used to detect the value of a data point within an asset going above or below a set threshold.

The configuration of the rule allows the threshold value to be set, the operation and the datapoint used to trigger the rule.

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Asset name

Datapoint name

Condition

Trigger value

Evaluation data

Window evaluation

Time window

Previous Next

- **Asset name:** The name of the asset that is tested by the rule.
- **Datapoint Name:** The name of the datapoint in the asset used for the test.
- **Condition:** The condition that is being tested, this may be one of >, >=, <= or <.
- **Trigger value:** The value used for the test.
- **Evaluation data:** Select if the data evaluate is a single value or a window of values.
- **Window evaluation:** Only valid if evaluation data is set to Window. This determines if the value used in the rule evaluation is the average, minimum or maximum over the duration of the window.
- **Time window:** Only valid if evaluation data is set to Window. This determines the time span of the window.

8.4.2 Moving Average Rule

The *foglamp-rule-average* plugin is a notification rule that is used to detect when a value moves outside of the determined average by more than a specified percentage. The plugin only monitors a single asset, but will monitor all data points within that asset. It will trigger if any of the data points within the asset differ by more than the configured percentage, an average is maintained for each data point separately.

During the configuration of a notification use the screen presented to choose the average plugin as the rule.

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Rule Plugin

- Average
- OutOfBound
- SimpleExpression
- Threshold

Trigger if the current value deviates from the moving average by more than a defined percentage

[available plugins](#)

Previous Next

The next screen you are presented with provides the configuration options for the rule.

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Asset temperature

Deviation % 10

Direction Both

Average Simple Moving Average

EMA Factor 10

Previous Next

The *Asset* entry field is used to define the single asset that the plugin should monitor.

The *Deviation %* defines how far away from the observed average the current value should be in order to be considered as triggering the rule.

Deviation %	<div>Above Average Below Average ✓ Both</div>
Direction	

The *Direction* entry is used to define if the rule should trigger when the current value is above average, below average or in both cases.

Average	<div>✓ Simple Moving Average Exponential Moving Average 10</div>
EMA Factor	

The *Average* entry is used to determine what type of average is used for the calculation. The average calculated may be either a simple moving average or an exponential moving average. If an exponential moving average is chosen then a second configuration parameter, *EMA Factor*, allows the setting of the factor used to calculate that average.

Exponential moving averages give more weight to the recent values compared to historical values. The smaller the EMA factor the more weight recent values carry. A value of 1 for *EMA Factor* will only consider the most recent value.

Note: The Average rule is not applicable to all data, only simple numeric values are considered and those values should not deviate with an average of 0 or close to 0 if good results are required. Data points that deviate wildly are also not suitable for this plugin.

8.4.3 Delta Rule

The *foglamp-rule-delta* plugin is a notification rule that triggers when a data point value changes. When a datapoint that is monitored changes the plugin will trigger. An alias value is given for the triggered datapoint and is included in the reason message when the plugin triggers.

During the configuration of a notification use the screen presented to choose the delta plugin as the rule.

The screenshot shows the 'Rule' configuration step in a four-step wizard. The steps are: 1. Notification Instance, 2. Rule (current), 3. Delivery Channel, and 4. Done. The 'Rule Plugin' section contains a dropdown menu with the following options: Average, Delta, OutOfBound, and Simple-Expression. Below the dropdown is a link labeled 'available_plugins'. At the bottom of the screen are 'Previous' and 'Next' buttons.

The next screen you are presented with provides the configuration options for the rule.

The screenshot shows the 'Rule' configuration step, specifically the 'JSON Configuration' section. The 'Asset' field is empty. The 'JSON Configuration' field contains the following JSON code:

```
1 {  
2   "datapoint_name": "alias_name",  
3   "sinusoid": "cosinus"  
4 }
```

At the bottom of the screen are 'Previous' and 'Next' buttons.

- **Asset:** define the single asset that the plugin should monitor.
- **Datapoints:** the datapoints monitor and the alias for the datapoint that is used in the trigger reason.

8.4.4 Periodic Rule

The *foglamp-rule-periodic* plugin is a notification rule that is used to trigger at regular intervals when data is being received for an asset.

The screenshot displays the configuration interface for the *foglamp-rule-periodic* plugin. At the top, a progress bar indicates the current step in the configuration process, with four steps labeled: 1. Notification Instance, 2. Rule, 3. Delivery Channel, and 4. Done. Step 2, 'Rule', is the active step. Below the progress bar, a configuration form is shown. It contains two input fields: 'Asset' and 'Interval'. The 'Interval' field is currently set to 3600. At the bottom of the form, there are two buttons: 'Previous' and 'Next'.

- **Asset:** the single asset that the plugin should monitor.
- **Interval:** the minimum time interval, in seconds between the rule triggering.

Whenever the plugin receives data for the specified it will check to see how long it has been since it last triggered. If it is *interval* seconds or longer since it triggered then the rule will be triggered. If no data is being received for the specified asset then the rule will not trigger.

This rule is commonly used to perform periodic operations, such as data sampling on an asset, the action associated with the rule could be to enable and disable an instance of the to control sending this data to the north.

A similar rule, exists that will trigger if data is not seen for a given asset within a given time frame.

8.4.5 Expression Rule

The *foglamp-rule-simple-expression* is a notification rule plugin that evaluates a user defined function to determine if a notification has triggered or not. The rule will work with a single asset, but does allow access to all the data points within the asset.

During the configuration of a notification use the screen presented to choose the average plugin as the rule.

The next screen you are presented with provides the configuration options for the rule.

The *Asset* entry field is used to define the single asset that the plugin should monitor.

The *Expression to apply* defines the expression that will be evaluated each time the rule is checked. This should be a boolean expression that returns true when the rule is considered to have triggered. Each data point within the asset will become a symbol in the expression, therefore if your asset contains a data point called voltage, the symbol voltage can be used in the expression to obtain the current voltage reading. As an example to create an under voltage notification if the voltage falls below 48 volts, the expression to use would be;

```
voltage < 48
```

The trigger expression uses the same expression mechanism, as the , and plugins

Expression may contain any of the following. . .

- Mathematical operators (+, -, *, /, %, ^)
- Functions (min, max, avg, sum, abs, ceil, floor, round, roundn, exp, log, log10, logn, pow, root, sqrt, clamp, inrange, swap)
- Trigonometry (sin, cos, tan, acos, asin, atan, atan2, cosh, cot, csc, sec, sinh, tanh, d2r, r2d, d2g, g2d, hyp)
- Equalities & Inequalities (=, ==, <>, !=, <, <=, >, >=)

- Logical operators (and, nand, nor, not, or, xor, xnor, mand, mor)

8.4.6 Simple-Sigma Rule

The *foglamp-rule-simple-sigma* is a notification rule plugin that uses the principle of normal distribution to trigger a notification if a value is found to be an , a value that is outside of the normal distribution. The normal distribution is discovered by taking the mean of all the values over time and calculating the standard deviation, or sigma, from that mean. Until the rule has built up a reasonable sample of data on which to calculate the mean and standard deviation the rule will not trigger. This reasonable sample is defined as a time period, in hours, for which the rule will simply sample the data to determine the mean and sigma values. During this period the rule will not trigger. The default time period for this is 1 hour, however it may be overridden.

Once a mean and standard deviation have been determined the rule will mode into a mode in which it will trigger. Whilst in triggering mode the rule will still refine the mean and standard deviation values. If a value is found in trigger mode that is more than a certain number of standard deviations from the mean, then the rule will trigger. The number of standard deviations is the sigma factor and defaults to 3.0, however the user can configure this to be more or less than 3.0.

To use the Simple-Sigma plugin create your notification rule as normal, when selecting the rule to use select the *Simple-Sigma* rule and click on next. You will be presented with a dialog as below

The screenshot shows a configuration dialog for the Simple-Sigma rule. At the top, a progress bar indicates four steps: 1. Notification Instance, 2. Rule (active), 3. Delivery Channel, and 4. Done. Below the progress bar, the configuration fields are: Asset name (Current), SigmaRule Factor (3.0), and Sample Size (hours) (1). At the bottom, there are 'Previous' and 'Next' buttons.

Configure the Simple-Sigma rule

- **Asset name:** The asset name to monitor with the rule
- **Sigma Factor:** The factor to use for determining range, a factor of 3.0 will trigger when a value is more the 3.0 * Sigma from the current mean
- **Sample Size:** The number of hours to build a mean and standard deviation before the rule will trigger.

Click on *Next* and complete the configuration of your notification.

8.4.7 Watchdog Rule

The *foglamp-rule-watchdog* is a notification rule plugin that is to detect the absence of data. The plugin is simply configured with the name of an asset to track and a period to test over. If no new readings are observed within that specified period of time then the rule will trigger.

During configuration of the notification choose the WatchDog rule from the set of available rules.

The screenshot shows a four-step configuration wizard. Step 2, 'Rule', is active. A dropdown menu for 'Rule Plugin' is open, showing options: OutOfBound, Simple-Expression, Threshold, and WatchDog (which is selected). To the right of the dropdown, a description reads: 'Generate a notification based on the last time of received data'. Below the dropdown is a link that says 'available plugins'. At the bottom of the wizard are 'Previous' and 'Next' buttons.

The next step in the process is to enter the configuration options for the watchdog rule.

The screenshot shows the same four-step configuration wizard, now at step 2 'Rule'. The 'Rule Plugin' dropdown is closed. Two input fields are visible: 'Asset name' with the value 'sinusoid' and 'Evaluation interval in ms' with the value '5000'. A help icon (?) is in the top right corner of the configuration box. 'Previous' and 'Next' buttons are at the bottom.

- **Asset name:** The name of the asset to track.
- **Evaluation interval in ms:** The period of time to monitor for new readings of the asset. This is expressed in milliseconds.

As soon as the configured time period expires, if no readings for the asset have been observed then the rule will trigger.

It is important to consider the tuning of the south service buffering when setting up watchdog plugins. The watchdog is based on data entering the FogLAMP buffer, hence if the south service buffers data for a period that is the same or close to the watchdog period then false triggering of the watchdog may occur. Ensure the south service maximum latency is less than the watchdog interval for reliable behavior.

8.5 FogLAMP Notification Delivery Plugins

8.5.1 Amazon Alexa Notification

The *foglamp-notify-alexa* notification delivery plugin sends notifications via Amazon Alexa devices using the Alexa *NotifyMe* skill.

When you receive a notification Alexa will make a noise to say you have a new notification and the green light on your Alexa device will light to say you have waiting notifications. To hear your notifications simply say “Alexa, read my notifications”

To enable notifications on an Alexa device

- You must enable the NotifyMe skill on your Amazon Alexa device.
- Link this skill to your Amazon account
- NotifyMe will send you an access code that is required to configure this plugin.

Once you have created your notification rule and move on to the delivery mechanism

- Select the alexa plugin from the list of plugins
- Click *Next*

The screenshot shows a configuration wizard for the Amazon Alexa notification plugin. At the top, a progress bar indicates four steps: 1. Notification Instance, 2. Rule, 3. Delivery Channel, and 4. Done. The 'Rule' step is currently active. Below the progress bar, there is a form with three fields: 'Access Code' (empty), 'Title' (containing 'The level in tank 15 is below 10%'), and 'Enabled' (checked). At the bottom, there are 'Previous' and 'Next' buttons.

- Configure the plugin
 - **Access Code:** Paste the access code you received from the *NotifyMe* application here
 - **Title:** This is the title that the Alexa device will read to you
- Enable the plugin and click *Next*
- Complete your notification setup

When you notification triggers the Alexa device will read the title text to you followed by either “Notification has triggered” or “Notification has cleared”.

8.5.2 Asset Notification

The *foglamp-notify-asset* notification delivery plugin is unusual in that it does not notify an external system, instead it creates a new asset which is then processed like any other asset within FogLAMP. This plugin is useful to inform up stream systems that a event has occurred and allow them to take action or merely as a way to have a record of a condition occurring which may not require any further actions.

Once you have created your notification rule and move on to the delivery mechanism

- Select the asset plugin from the list of plugins
- Click *Next*

- Now configure the asset delivery plugin
 - **Asset:** The name of the asset to create.
 - **Description:** A textual description to add to the asset
- Enable the plugin and click *Next*
- Complete your notification setup

The asset that will be created when the notification triggers will contain

- The timestamp of the trigger event
- Three data points
 - **rule:** The name of the notification that triggered this asset creation
 - **description:** The textual description entered in the configuration of the delivery plugin
 - **event:** This will be one of *triggered* or *cleared*. If the notification type was not set to be *toggled* then the *cleared* event will not appear. If *toggled* was set as the notification type then there will be a *triggered* value in the asset created when the rule triggered and a *cleared* value in the asset generated when the rule moved from the triggered to untriggered state.

8.5.3 Configuration Update

The *foglamp-notify-config* plugin is designed to allow a notification to alter the configuration of one of the configuration items within the local FogLAMP.

The plugin can be used to trigger changes to the way data is collected, for example by altering the *readingsPerSec* item in a south server Advanced category. It is not limited to this however and could equally be used to effect some configuration of a filter, for example to change a scale factor or threshold. It may also change configuration of notification rule or delivery plugins.

Once you have created your notification rule and moved on to the delivery mechanism

- Select the *config* plugin from the list of plugins
- Click *Next*

The screenshot shows a configuration form for a notification rule. At the top, a progress bar indicates four steps: 1. Notification Instance, 2. Rule, 3. Delivery Channel, and 4. Done. The current step is 'Rule'. The form contains the following fields:

Field	Value
Category	SineAdvanced
Item	readingsPerSecond
Trigger Value	20
Cleared Value	1
Enabled	<input checked="" type="checkbox"/>

At the bottom of the form, there are two buttons: 'Previous' and 'Next'.

- Configure the delivery plugin
 - **Category:** The name of the configuration category to be updated.
 - **Item:** The name of the item within the configuration category to be updated.
 - **Trigger Value:** The value to set the item to when an notification is triggered.
 - **Clear Value:** The value to set the item to when the notification is cleared. Note you must set the notification type to *toggled* if you wish to use a *Clear Value*.
- Enable the plugin and click *Next*
- Complete your notification setup

8.5.4 Control Dispatcher Notification

The *foglamp-notify-control* notification delivery plugin is a mechanism by which a notification can be used to send set point control writes and operations via the control dispatcher service.

Once you have created your notification rule and move on to the delivery mechanism

- Select the control plugin from the list of plugins
- Click *Next*

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Trigger Value

```
1 {  
2   "write": {  
3     "name": "value"  
4   }  
5 }
```

Cleared Value

```
1 {  
2   "write": {  
3     "name": "value"  
4   }  
5 }
```

Enabled ☐

Previous Next

- Configure the plugin
 - **Trigger Value:** The control payload to send to the dispatcher service. These are set when the notification rule triggers.
 - **Cleared Value:** The control payload to send to the dispatcher service. These are set when the notification rule clears.
- Enable the plugin and click *Next*
- Complete your notification setup

Trigger Values

The *Trigger Value* and *Cleared Value* are JSON documents that are sent to the dispatcher. The format of these is a JSON document that describes the control operation to perform. The document contains a definition of the recipient of the control input, this may be a south service, an asset, an automation script or all south services. If the recipient is a service then the document contains a *service* key, the value of which is the name of the south service that should receive the control operation. To send the control request to the south service responsible for the ingest of a particular asset, then an *asset* key would be given. If sending to an automation script then a *script* key would be given. If known of the *service*, *asset* or *script* keys are given then the request will be sent to all south services that support control.

The document also contains the control request that should be made, either a *write* or an *operation* and the values to write or the name and parameters of the operation to perform.

The example below shows a JSON document that would cause the two values *temperature* and *rate* to be written to the south service called *oven001*.

```
{
  "service" : "oven001",
  "write": {
    "temperature" : "110",
    "rate"       : "245"
  }
}
```

In this example the values are constants defined in the plugin configuration. It is possible however to use values that are in the data that triggered the notification.

As an example of this assume we are controlling the speed of a fan based on the temperature of an item of equipment. We have a south service that is reading the temperature of the equipment, let's assume this is in an asset called *equipment* which has a data point called *temperature*. We add a filter using the *foglamp-filter-expression* filter to calculate a desired fan speed. The expression we will use in this example is *desiredSpeed = temperature * 100*. This will cause the asset to have a second data point called *desiredSpeed*.

We create a notification that is triggered if the *desiredSpeed* is greater than 0. The delivery mechanism will be this plugin, *foglamp-notify-setpoint*. We want to set two values in the south plugin *speed* to set the speed of the fan and *run* which controls if the fan is on or off. We set the *Trigger Value* to the following

```
{
  "service" : "fan001",
  "write": {
    "speed" : "$equipment.desiredSpeed$",
    "run"   : "1"
  }
}
```

In this case the *speed* value will be substituted by the value of the *desiredSpeed* data point of the *equipment* asset that triggered the notification to be sent.

If then fan is controlled by the same south service that is ingesting the data into the asset *equipment*, then we could use the *asset* destination key rather than name the south service explicitly.

```
{
  "asset" : "equipment",
  "write": {
    "speed" : "$equipment.desiredSpeed$",
    "run"   : "1"
  }
}
```

Another option for controlling the destination of a control request is to broadcast it to all south services. In this example we will assume we want to trigger a shutdown operation across all the devices we monitor.

```
{
  "operation" : {
    "shutdown" : { }
  }
}
```

Here we are not giving *asset*, *script* or *service* keys, therefore the control request will be broadcast. Also we have used an *operation* rather than a *write* request. The operation name is *shutdown* and we have assumed it takes no arguments.

8.5.5 Custom Asset Notification

The *foglamp-notify-customasset* notification delivery plugin is a plugin that creates an event asset in the FogLAMP readings. This event asset can be customised via the configuration and may include data from the asset that triggered the notification.

The asset created will contain a number of data points

- *description*: A fixed description that is set in the plugin configuration.
- *event*: The event that caused the asset to be created. This may be one of *triggered* or *cleared*.
- *rule*: The name of the notification rule that triggered the notification.
- *store*: The data that triggered the notification.

Once you have created your notification rule and move on to the delivery mechanism

- Select the *customasset* plugin from the list of plugins
- Click *Next*

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

CustomAsset

Description

JSON Configuration

```

1 {
2   "sinusoid": [
3     {
4       "sinusoid": "cosinusoid"
5     }
6   ]
7 }

```

Enabled ☐

Enable authentication ☐

Username

Password

- Configure the plugin
 - **Custom Asset:** The name of the assert to create
 - **Description:** The content to add in the *description* data point within the asset
 - **JSON Configuration:** This is a description of how to map the asset that triggered the notification to the data in the *store* data point of the event asset.
 - **Enable authentication:** Enable the authentication of the plugin to the FogLAMP API
 - **Username:** The user name to use when connecting to the FogLAMP API
 - **Password:** The password to use when connecting to the FogLAMP API
- Enable the plugin and click *Next*
- Complete your notification setup

Store Configuration

The content of the *store* data point would normally contain data from the reading that triggered the notification. It will be written as a JSON data point type and will always contain the timestamp of the reading that triggered this notification.

The configuration consists of one or more asset names as a key, the value is an array of objects that defines the data point names to extract from the triggering asset and an alias to use in the store. The example below would include the *sinusoid* asset and the data point within *sinusoid*, also called *sinusoid*. However it would write this value using an alias of *cosinusoid*.

```
{
  "sinusoid": [
    {
      "sinusoid": "cosinusoid"
    }
  ]
}
```

This would result in an asset with a *store* data point that would be as follows

```
{"sinusoid":{"cosinusoid":0.994521895,"timestamp":"2022-09-08 11:31:29.323666 +0000"}}
```

8.5.6 Email Notifications

The *foglamp-notify-email* delivery notification plugin allows notifications to be delivered as email messages. The plugin uses an SMTP server to send email and requires access to this to be configured as part of configuring the notification delivery method.

During the creation of your notification select the email notification plugin from the list of available notification mechanisms. You will be prompted with a configuration dialog in which to enter details of your SMTP server and of the email you wish to send.

The screenshot shows a configuration dialog for an email notification plugin. At the top, a progress bar indicates four steps: 1. Notification Instance, 2. Rule, 3. Delivery Channel (current step), and 4. Done. The dialog contains the following fields:

- To address:** alert.subscriber@dianomic.com
- To:** Notification alert subscriber
- Subject:** Fledge alert notification
- From address:** dianomic.alerts@gmail.com
- From name:** Notification alert
- SMTP Server:** smtp.gmail.com
- SMTP Port:** 587
- SSL/TLS:** ☒
- Username:** dianomic.alerts@gmail.com
- Password:** pass
- Enabled:** ☐

At the bottom, there are 'Previous' and 'Next' buttons.

- **To address:** The email address to which the notification will be sent
- **To:** A textual name for the recipient of the email
- **Subject:** A Subject to put in the email message
- **From address:** A from address to use for the email message
- **From name:** A from name to include in the email
- **SMTP Server:** The address of the SMTP server to which to send messages

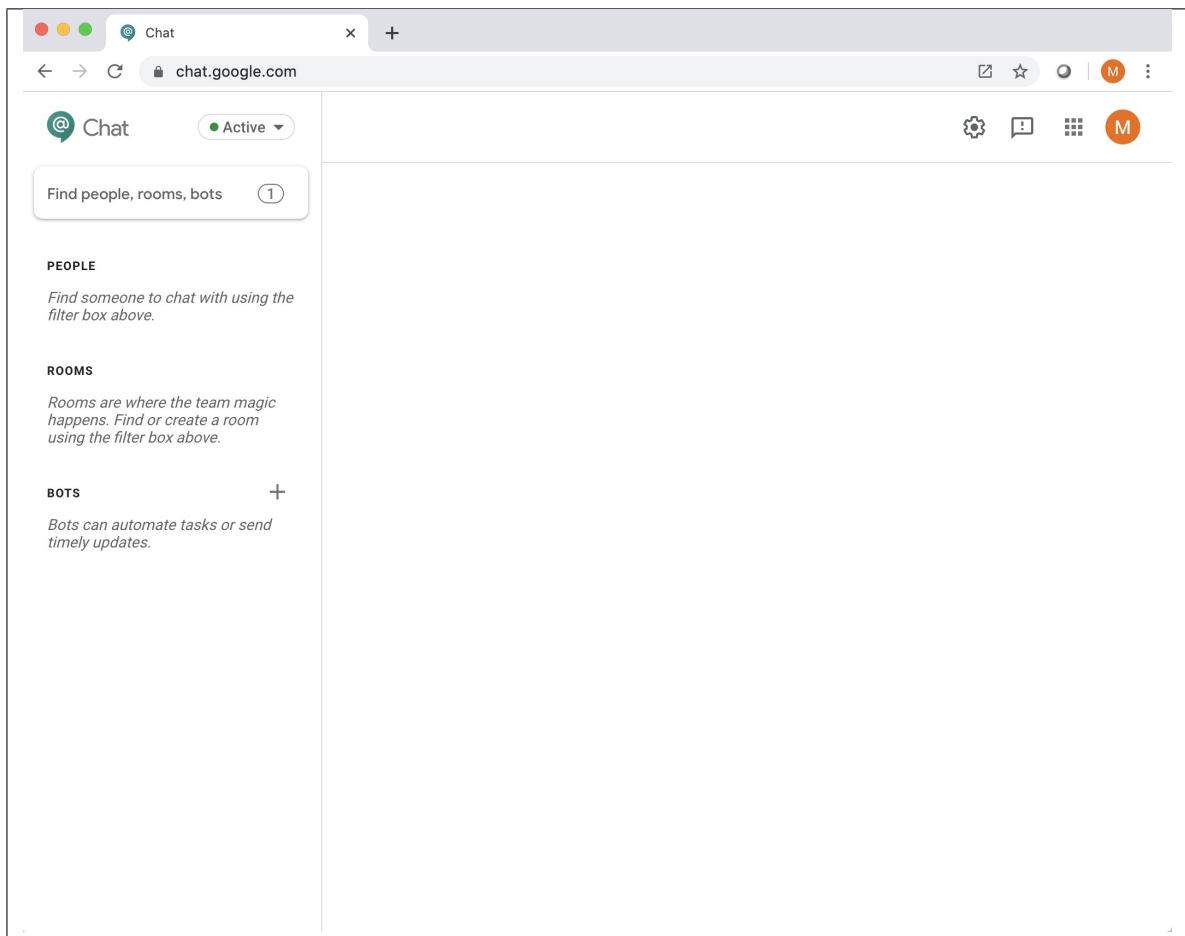
- **SMTP Port:** The port of your SMTP server
- **SSL/TLS:** A toggle to control if SSL/TLS encryption should be used when communicating with the SMTP server
- **Username:** A username to use to authenticate with the SMTP server
- **Password:** A password to use to authenticate with the SMTP server.

8.5.7 Google Chat

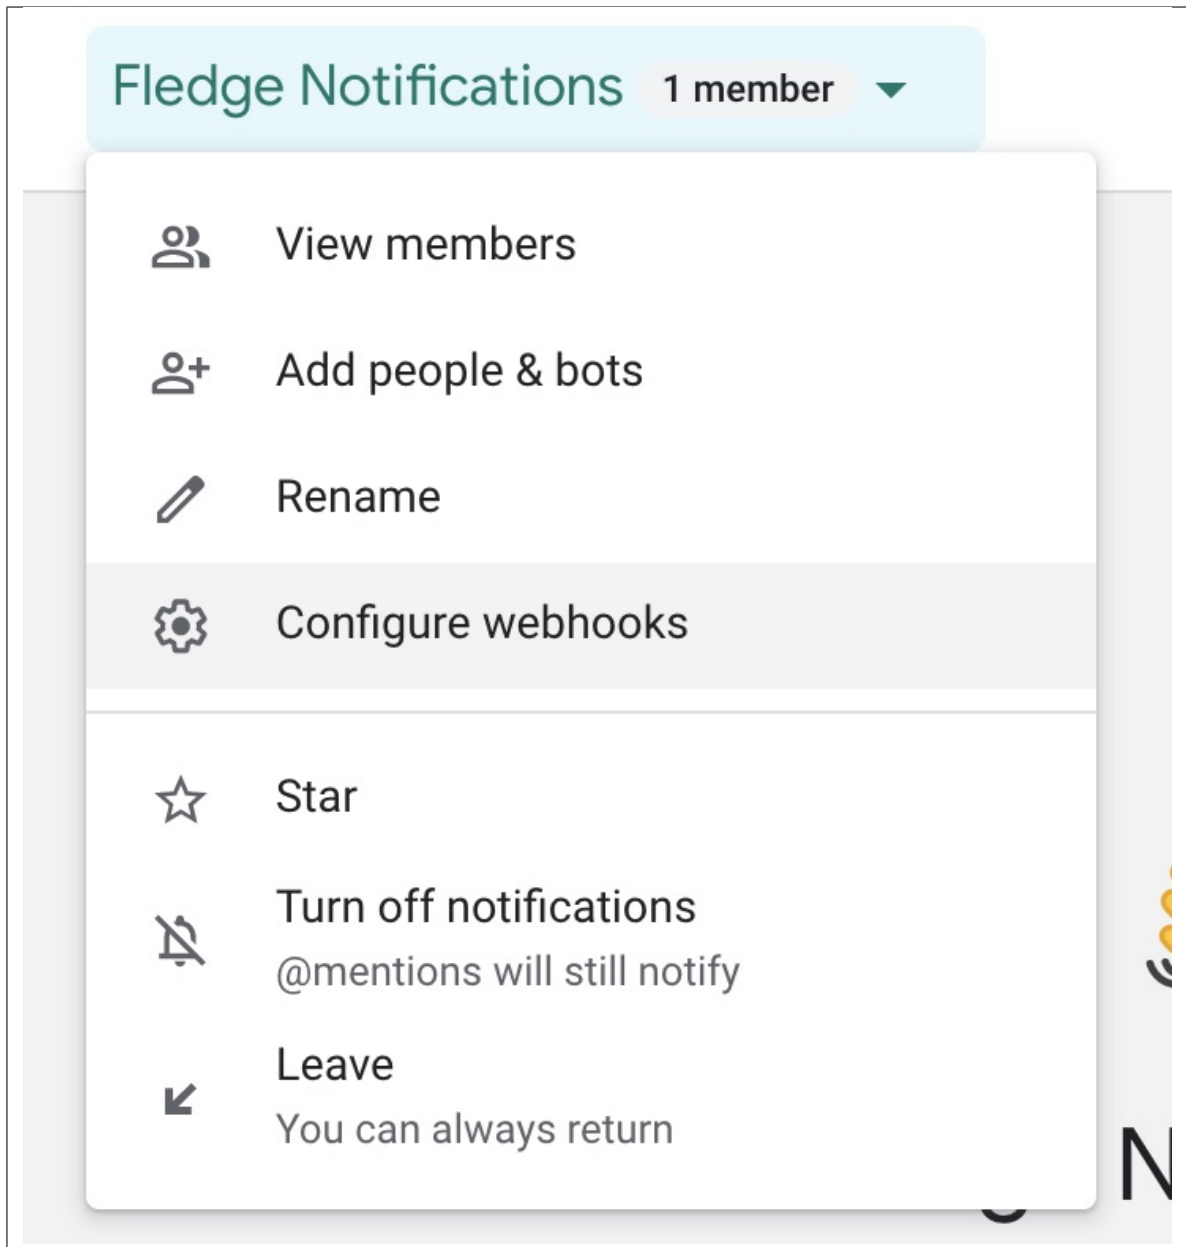
The *foglamp-notify-google-hangouts* plugin allows notifications to be delivered to the Google chat platform. The notification are delivered into a specific chat room within the application, in order to allow access to the chat room you must create a webhook for sending data to that chatroom.

To create a webhook

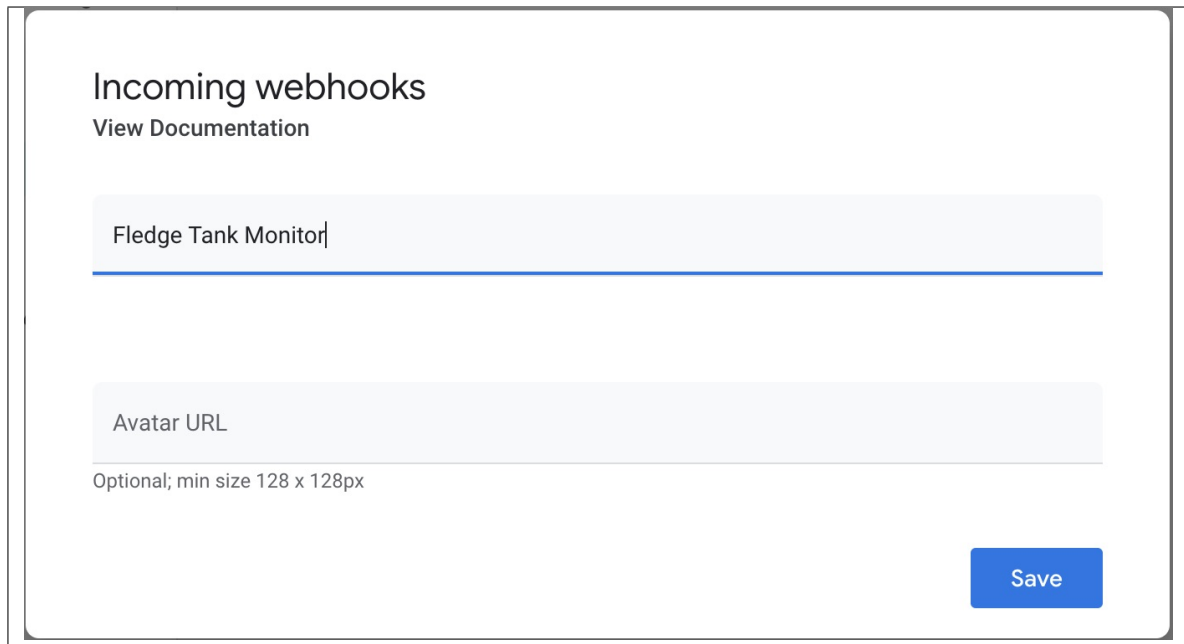
- Go to the page in your browser



- Select the chat room you wish to use or create a new chat room
- In the menu at the top of the screen select *Configure webhooks*



- Enter a name for your webhook and optional avatar and click *Save*



Incoming webhooks

[View Documentation](#)

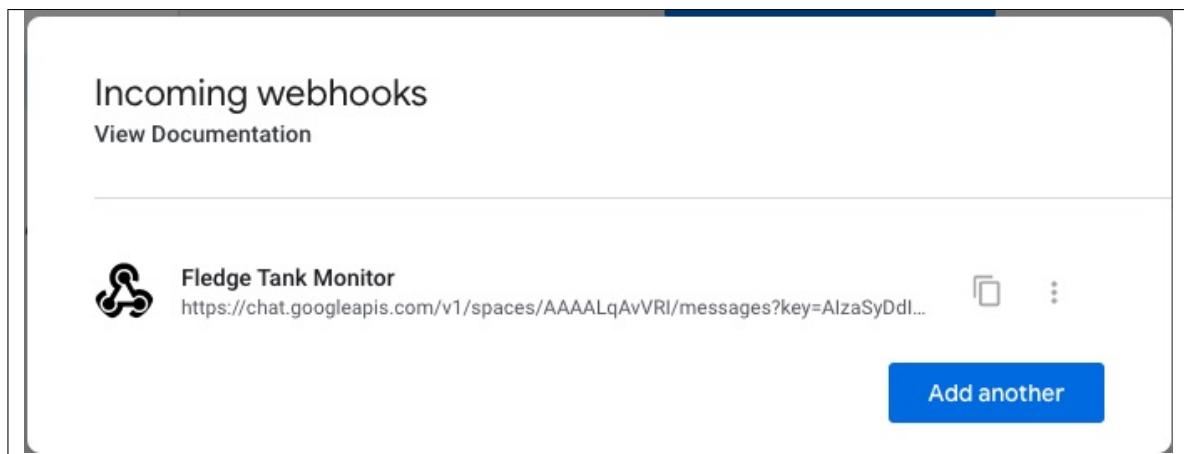
Fledge Tank Monitor

Avatar URL

Optional; min size 128 x 128px


Save

- Copy the URL that appears under your webhook name, you can use the copy icon next to the URL to place it in the clipboard



Incoming webhooks

[View Documentation](#)

 **Fledge Tank Monitor**
https://chat.googleapis.com/v1/spaces/AAAAALqAvVRI/messages?key=AlzaSyDdl...

Copy icon and menu icon

Add another

- Close the webhooks window by clicking outside the window

Once you have created your notification rule and move on to the delivery mechanism

- Select the Hangouts plugin from the list of plugins
- Click *Next*

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Google Hangout Webhook URL

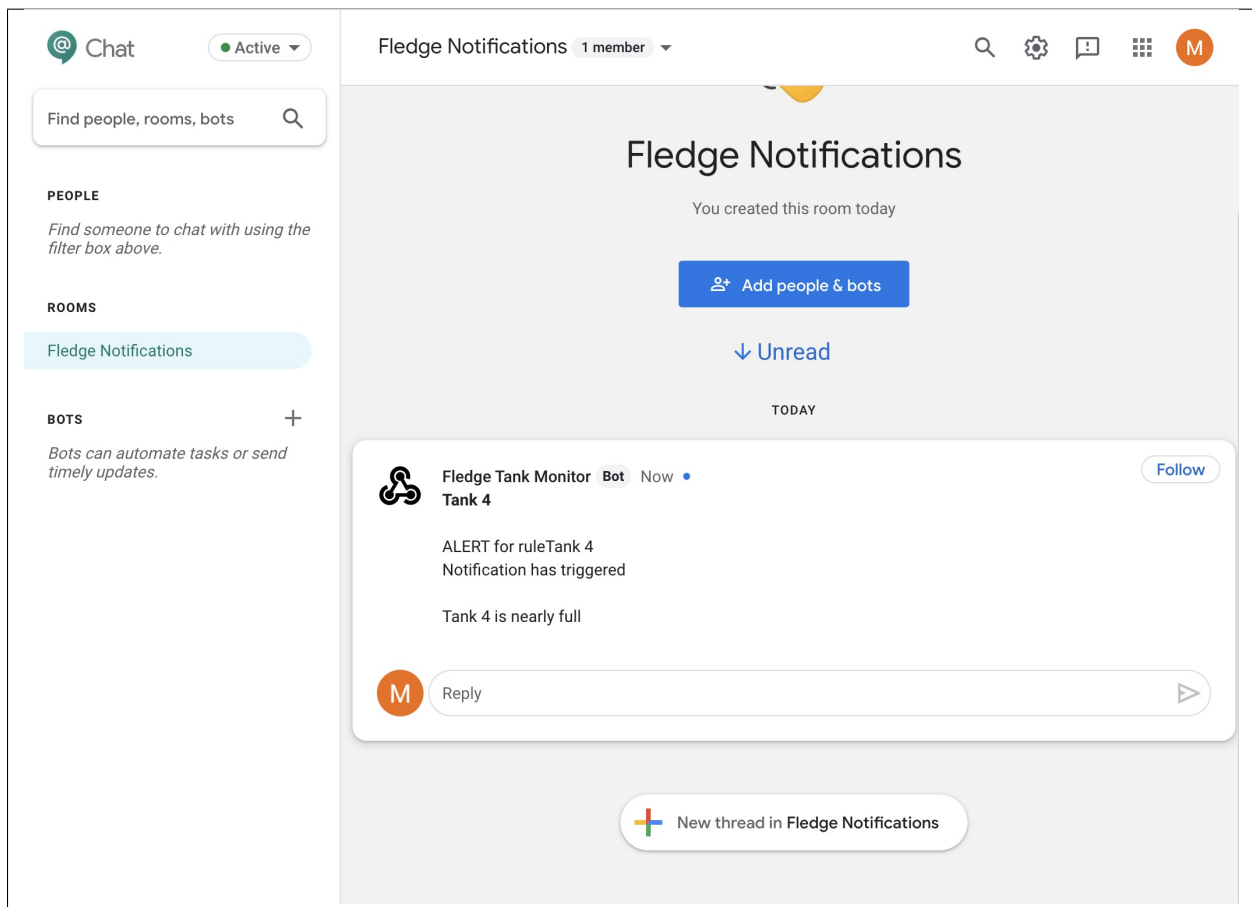
Message Text

Enabled ☒

[Previous](#) [Next](#)

- Now configure the asset delivery plugin
 - **Google Hangout Webhook URL:** Paste the URL obtain above here
 - **Message Text:** Enter the message text you wish to send
- Enable the plugin and click *Next*
- Complete your notification setup

A message will be sent to this chat room whenever a notification is triggered.

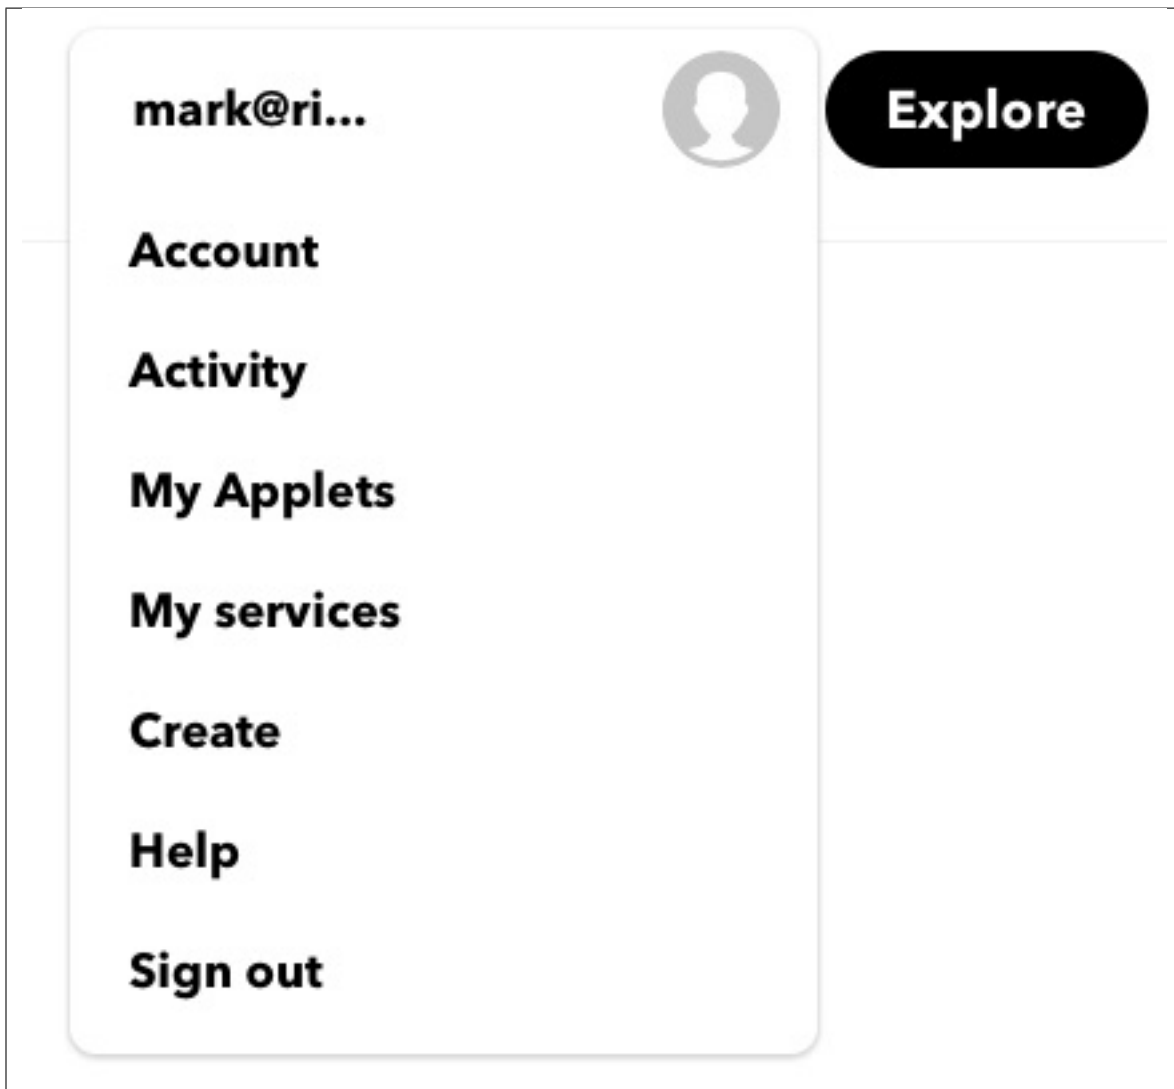


8.5.8 IFTTT Delivery Plugin

The *foglamp-notify-ifttt* is a notification delivery plugin designed to trigger an action on the *If This Than That* IoT platform. IFTTT allows the user to setup a webhook that can be used to trigger processing on the platform. The webhook could be sending an IFTTT notification to a destination not support by any FogLAMP plugin to controlling a device that is controllable via IFTTT.

In order to use the IFTTT webhook you must obtain a key from IFTTT by visiting your IFTTT account

- Select the “My Applets” page from your account pull down menu



- Select “New Applet”
- Click on the blue “+ this” logo
- Choose the service Webhooks
- Click on the blue box “Receive a web request”
- Enter an “Event Name”, this may be of your choosing and will be put in the configuration entry ‘Trigger’ for the FogLAMP plugin

- Click on the “+ that” logo
- Select the action you wish to invoke

Once you have setup your webhook on IFTTT you can now proceed to setup the FogLAMP delivery notification plugin. Create you notification, choose and configure your notification rule. Select the IFTTT delivery plugin and click on *Next*. You will be presented with the IFTTT plugin configuration page.

The screenshot displays the IFTTT plugin configuration interface. At the top, a progress bar indicates the current step is 'Rule' (step 2 of 4). The configuration area contains three fields: 'IFTTT Trigger' set to 'button_press', 'IFTTT Key' set to a long alphanumeric string, and 'Enabled' checked. Navigation buttons 'Previous' and 'Next' are located at the bottom.

There are two important items to be configured

- **IFTTT Trigger:** This is the *Maker Event* that you used in IFTTT when defining the action that the webhook should trigger.
- **IFTTT Key:** This is the webhook key you obtain from the IFTTT platform.

Enable the delivery and click on *Next* to move to the final stage of completing your notification.

8.5.9 Jira Ticket Creation

The *foglamp-notify-jira* delivery notification plugin allows notifications to be used to create tickets within Jira. The tickets are created within a specified project with a summary, description and other information supplied by FogLAMP.

To obtain an API token from Jira


- Visit the page
- Select *Create API token*
- Enter a name for your application, this must be unique for each FogLAMP Jira application you create
- Click on Create

Once you have created your notification rule and move on to the delivery mechanism

- Select the *jira* plugin from the list of plugins
- Click *Next*

- Configure the delivery plugin
 - **Hostname:** The hostname where your Jira instance is installed. This may be a local instance or a cloud instance.
 - **Project:** The project into which you are creating the Jira tickets. The project name should be the one that appears as projectKey in the URL bar when browsing the Jira boards.
 - **User:** Your Jira user name, this is the name of the account you used to create the API token
 - **API Token:** The API token you created above
 - **Summary:** The text to add into the ticket summary, this may include text substitutions (see below).
 - **Description:** The text to add into the ticket description, this may include text substitution (see below).
 - **Type:** The issue type to create. This must be the name of one of the types that is valid for your Jira project.
 - **Additional Fields:** This is a JSON document that contains a number of key/value pairs, each of these pairs is a field name and content to add to the ticket. Text substitutions may be applied here also.
- Enable the plugin and click *Next*
- Complete your notification setup

When the notification rule triggers you a Jira ticket will be created.



Text Substitution

Text markers may be used to substitution text with the fields in the Jira ticket. The markers supported are

- **%MESSAGE%**: this is replaced with the message generated in the notification system
- **%REASON%**: this is replaced with the reason for the notification, it may be the string *triggered* or *cleared*.
- **%TIMESTAMP%**: this is replaced with the timestamp of the reading data that caused the notification to trigger.

8.5.10 JSON Configuration Update

The *foglamp-notify-jsonconfig* plugin is designed to allow a notification to alter the configuration of one of the JSON configuration items within the local FogLAMP.

The plugin can be used to trigger changes to the way data is collected by altering individual items within a complex JSON configuration items. The delivery plugin allows you to set a value when the notification is raised and a different value when it is cleared.

Once you have created your notification rule and moved on to the delivery mechanism

- Select the *config* plugin from the list of plugins
- Click *Next*

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Category

Item

JSON Path

Property

Trigger Value

Cleared Value

Enabled ☐

Previous Next

- Configure the delivery plugin
 - **Category:** The name of the configuration category to be updated.
 - **Item:** The name of the item within the configuration category to be updated.
 - **JSON Path:** The JSON path of the object that contains the item to be modified.
 - **Property:** The name of the JSON property to modify.
 - **Trigger Value:** The value to set the item to when an notification is triggered.
 - **Clear Value:** The value to set the item to when the notification is cleared. Note you must set the notification type to *toggled* if you wish to use a *Clear Value*.
- Enable the plugin and click *Next*
- Complete your notification setup

JSON Path

A subset of the full JSON Path expressions are supported in this plugins. Each path element is proceeded by a / character and may be one of

- **Literals:** A literal object name within the JSON document. E.g. `/a/b/c`
- **An Array Index:** An absolute index within an array. E.g. `a[2]`
- **A conditional test:** A property value to match within an array or object. `a[prop==value]`

To match the object under the *registers* element within the *map* element an expression would be of the form

```
/map/registers
```

To match the first element in the array called *assets* under the *exclusions* object the expression would be

```
/exclusions/assets[0]
```


To match the object that contains a property called *id* whose values in *QTE123* within the *connections* object the expression would be

```
/connections[id=="QTE123"]
```

8.5.11 Management Poll Notification

The *foglamp-notify-management* notification delivery plugin is designed to trigger the FogMan agent microservice of the current FogLAMP to poll its FogMan to retrieve any configuration updates for this FogLAMP.

Once you have created your notification rule and move on to the delivery mechanism

- Select the management plugin from the list of plugins
- Click *Next*
- There is no specific configuration for this plugin
- Enable the plugin and click *Next*
- Complete your notification setup

Plugin Uses

The plugin is designed for an environment whereby the updates of configuration of the FogLAMP are coordinated with the state of the equipment that is being monitored by the FogLAMP. This might be because you may wish to prevent updates from occurring during critical periods of operation or maybe because the FogLAMP is monitoring the network connectivity and you wish to synchronize updates with network availability.

8.5.12 MQTT Notification

The *foglamp-notify-mqtt* notification delivery plugin sends notifications via an MQTT broker. The MQTT topic and the payloads to send when the notification triggers or is cleared are configurable.

Once you have created your notification rule and move on to the delivery mechanism

- Select the mqtt plugin from the list of plugins
- Click *Next*

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

MQTT Broker

MQTT Topic

Trigger Payload

Clear Payload

Enabled ☒

[Previous](#) [Next](#)

- Configure the plugin
 - **MQTT Broker:** The URL of your MQTT broker.
 - **Topic:** The MQTT topic on which to publish the messages.
 - **Trigger Payload:** The payload to send when the notification triggers
 - **Clear Payload:** The payload to send when the notification clears
- Enable the plugin and click *Next*
- Complete your notification setup

8.5.13 Conditional Forwarding

The *foglamp-notify-north* plugin is designed to allow conditional forwarding of data to an existing north application from within FogLAMP.

The scenario the plugin addresses is the need to send data to a system north of FogLAMP when a condition occurs. The sending is done via a standard FogLAMP north task and can use any plugin such as OMF, GCP, InfluxDB, etc. The condition used to send this data is monitored using the notification server, when the rule in the notification triggers we send data from the FogLAMP storage service to the specified north task.

The data that is sent is based on the time the notification triggered and two configuration parameters, pre-trigger and post-trigger times. The pre-trigger setting control how long before the event the data is sent and the post-trigger for how long after the event data is sent.

The data that is sent may be anything that is buffered in the Foglamp storage service. A list of assets to send may be configured as part of the plugin configuration.

Once you have created your notification rule and move on to the delivery mechanism

- Select the North plugin from the list of plugins
- Click *Next*

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

North task name

Assets to send

```

1 {
2   "assets": []
3 }

```

Pre-trigger time

Post-trigger time

Block Size

Enabled ☐

- Configure the delivery plugin
 - **North task name:** This is the name of a north task to use for the sending of the data. The north task should have already been created but should be disabled.
 - **Assets to send:** A JSON structure that contains the list of assets that should be sent via the north task. This list is a simple JSON array of asset names.
 - **Pre-trigger time:** The length of time in seconds before the notification triggers for which data should be sent.
 - **Post-trigger time:** The length of time in seconds after the notification triggers for which data should be sent.
 - **Block size:** The size of the data block sent to the north service, this is a tuning parameter to throttle the data sent, under most circumstances it may be left as the default.
- Enable the plugin and click *Next*
- Complete your notification setup

8.5.14 Operation Notification

The *foglamp-notify-operation* notification delivery plugin is a mechanism by which a notification can be used to send a request to a south services to perform an operation.

Once you have created your notification rule and move on to the delivery mechanism

- Select the operation plugin from the list of plugins
- Click *Next*

- Configure the plugin
 - **Service:** The name of the south service you wish to control
 - **Trigger Value:** The operation payload to send to the south service when the rule triggers. This is the name of the operation to perform and a set of name, value pairs which are the optional parameters to pass that operations.
 - **Cleared Value:** The operation payload to send to the south service when the rule clears. This is the name of the operation to perform and a set of name, value pairs which are the optional parameters to pass that operations.
- Enable the plugin and click *Next*
- Complete your notification setup

8.5.15 Python 3 Script

The *foglamp-notify-python35* notification delivery plugin allows a user supplied Python script to be executed when a notification is triggered or cleared. The script should be written in Python 3 syntax.

A Python script should be provided in the form of a function, the name of that function should match the name of the file the code is loaded from. E.g if you have a script to run which you have saved in a file called `alert_light.py` it should contain a function `alert_light`. ~that function is called with a message which is defined in notification itself as a simple string.

A second function may be provided by the Python plugin code to accept configuration from the plugin that can be used to modify the behavior of the Python code without the need to change the code. The configuration is a JSON document which is again passed as a Python Dict to the `set_filter_config` function in the user provided Python code. This function should be of the form

```
def set_filter_config(configuration):
    config = json.loads(configuration['config'])
    value = config['key']
    ...
    return True
```

Once you have created your notification rule and move on to the delivery mechanism

- Select the python35 plugin from the list of plugins
- Click *Next*

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Python script

```
1 from time import sleep
2 from enviropat import leds
3
4 def flash_leds(message):
5     for count in range(4):
6         leds.on()
7         sleep(0.5)
8         leds.off()
9         sleep(0.5)
10
```

flash_leds.py

Choose Files flash_leds.py

Configuration

```
1 {}
```

Enabled ☐

Previous Next

- Configure the plugin
 - **Python Script:** This is the script that will be executed. Initially you are unable to type in this area and must load your initial script from a file using the *Choose Files* button below the text area. Once a file has been chosen and loaded you are able to update the Python code in this page.

Note: Any changes made to the script in this screen will be written back to the original file it was loaded from.

- **Configuration:** You may enter a JSON document here that will be passed to the *set_filter_config* function of your Python code.
- Enable the plugin and click *Next*
- Complete your notification setup

Example Script

The following is an example script that flashes the LEDs on the Enviro pHAT board on a Raspberry Pi

```
from time import sleep
from envirophat import leds
def flash_leds(message):
    for count in range(4):
        leds.on()
        sleep(0.5)
        leds.off()
        sleep(0.5)
```

This code imports some Python libraries and then in a loop will turn the leds on and then off 4 times.

Note: This example will take 4 seconds to execute, unless multiple threads have been turned on for notification delivery this will block any other notifications from being delivered during that time.

8.5.16 Set Point Control Notification

The *foglamp-notify-setpoint* notification delivery plugin is a mechanism by which a notification can be used to send set point control writes into south services which support set point control

Once you have created your notification rule and move on to the delivery mechanism

- Select the setpoint plugin from the list of plugins
- Click *Next*

The screenshot shows the 'Delivery Channel' configuration step (3) of a notification rule. The interface includes a progress bar at the top with steps: 1 Notification Instance, 2 Rule, 3 Delivery Channel, and 4 Done. The main configuration area has the following fields:

- Service:** A text input field.
- Trigger Value:** A JSON template editor showing:


```
1 {
2   "values": {
3     "name": "value"
4   }
5 }
```
- Cleared Value:** A JSON template editor showing:


```
1 {
2   "values": {
3     "name": "value"
4   }
5 }
```
- Enabled:** A checkbox that is currently unchecked.

- Configure the plugin
 - **Service:** The name of the south service you wish to control
 - **Trigger Value:** The set point control payload to send to the south service. This is a list of name, value pairs to be set within the service. These are set when the notification rule triggers.
 - **Cleared Value:** The set point control payload to send to the south service. This is a list of name, value pairs to be set within the service. These are set when the notification rule clears.
- Enable the plugin and click *Next*
- Complete your notification setup

Trigger Values

The *Trigger Value* and *Cleared Value* are JSON documents that are sent to the set point entry point of the south service. The format of these is a set of name and value pairs that represent the data to write via the south service. A simple example would be as below

```
{
  "values": {
    "temperature" : "11",
    "rate"        : "245"
  }
}
```

In this example we are setting two variables in the south service, one named *temperature* and the other named *rate*. In this example the values are constants defined in the plugin configuration. It is possible however to use values that are in the data that triggered the notification.

As an example of this assume we are controlling the speed of a fan based on the temperature of an item of equipment. We have a south service that is reading the temperature of the equipment, let's assume this is in an asset called *equipment* which has a data point called *temperature*. We add a filter using the *foglamp-filter-expression* filter to calculate a desired fan speed. The expression we will use in this example is *desiredSpeed = temperature * 100*. This will cause the asset to have a second data point called *desiredSpeed*.

We create a notification that is triggered if the *desiredSpeed* is greater than 0. The delivery mechanism will be this plugin, *foglamp-notify-setpoint*. We want to set two values in the south plugin *speed* to set the speed of the fan and *run* which controls if the fan is on or off. We set the *Trigger Value* to the following

```
{
  "values" : {
    "speed" : "$equipment.desiredSpeed$",
    "run"   : "1"
  }
}
```

In this case the *speed* value will be substituted by the value of the *desiredSpeed* data point of the *equipment* asset that triggered the notification to be sent.

8.5.17 Slack Messages

The *foglamp-notify-slack* delivery notification plugin allows notifications to be delivered as instant messages on the Slack messaging platform. The plugin uses a Slack webhook to post the message.

To obtain a webhook URL from Slack

- Visit the page
- Select *Create New App*
- Enter a name for your application, this must be unique for each FogLAMP slack application you create
- Select your Slack workspace in which to deliver your notification. If not already logged in you may need to login to your workspace
- Click on Create
- Select *Incoming Webhooks*
- Activate your webhook
- Add your webhook to the workspace
- Select the channel or individual to send the notification to
- Authorize your webhook
- Copy the Webhook URL which you will need when configuring the plugin

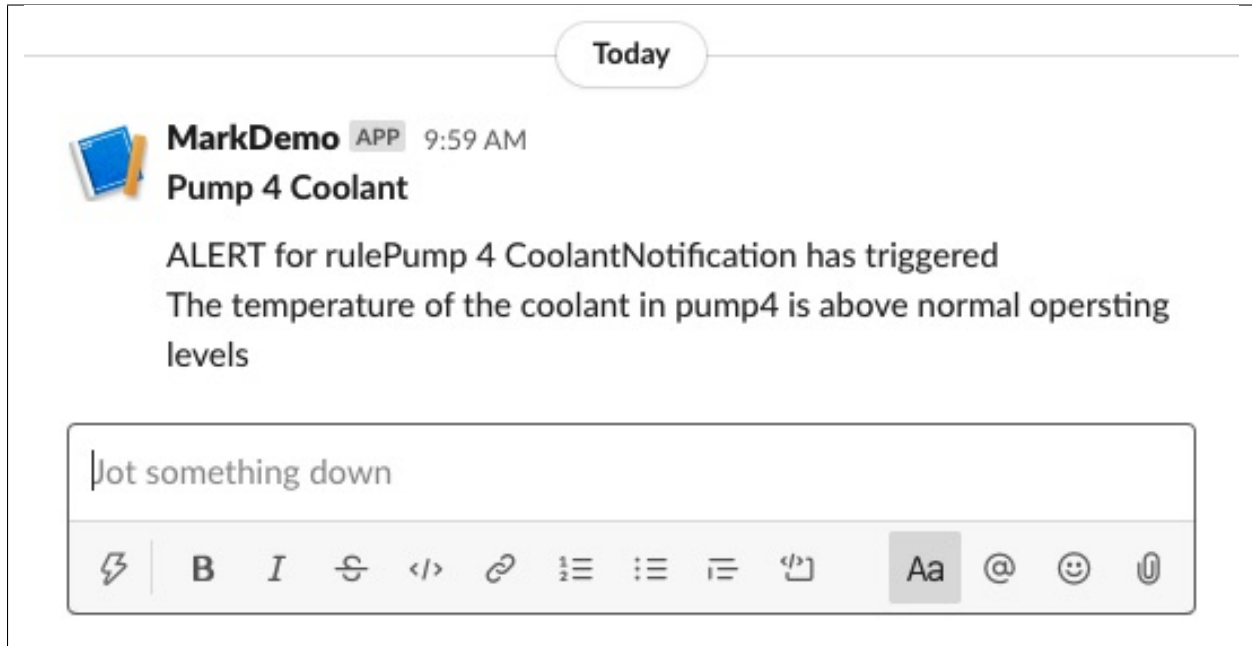
Once you have created your notification rule and move on to the delivery mechanism

- Select the slack plugin from the list of plugins
- Click *Next*


The screenshot shows a configuration window for a notification rule. At the top, a progress bar with four steps is visible: 1. Notification Instance, 2. Rule, 3. Delivery Channel, and 4. Done. The 'Rule' step is currently active. The configuration form contains three fields: 'Slack Webhook URL' with a text input containing a long URL, 'Message Text' with a text input containing 'The temperature of the coolant in pump4 is above normal operating levels', and 'Enabled' with a checked checkbox. At the bottom of the form are two buttons: 'Previous' and 'Next'.

- Configure the delivery plugin
 - **Slack Webhook URL:** Paste the URL you obtain above from the page
 - **Message Text:** Static text that will appear in the slack message you receive when the rule triggers
- Enable the plugin and click *Next*
- Complete your notification setup

When the notification rule triggers you will receive messages in you Slack client on your desk top

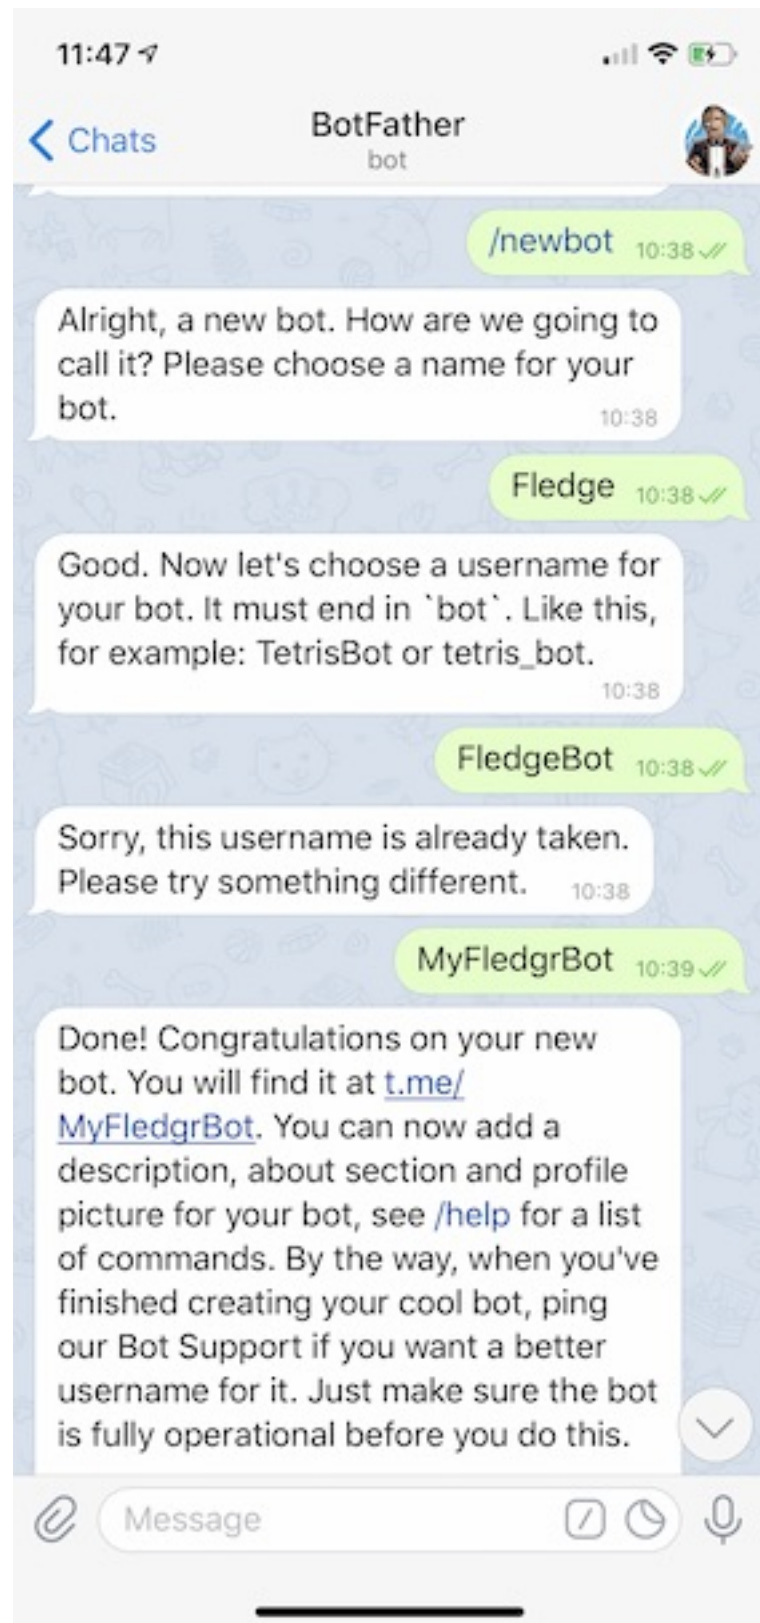


and/or your mobile devices



8.5.18 Telegram Messages

The *foglamp-notify-telegram* delivery notification plugin allows notifications to be delivered as instant messages on the Telegram messaging platform. The plugin uses Telegram BOT API, to use this you must create a BOT and obtain a token.



To obtain a Telegram BOT token

- Use the Telegram application to send a message to *botfather*.

- In your message send the text /start
- Then send the message /newbot
- Follow the instructions to name your BOT

- Copy your BOT token.

You now need to get a chat id

- In the Telegram application send a message to you chat BOT
- Run the following command at the your shell command line or use a web browser to go to the URL <https://api.telegram.org/bot<YourBOTToken>/getUpdates>

```
wget https://api.telegram.org/bot<YourBOTToken>/getUpdates
```

Examine the contents of the getUpdates file or the output from the web browser

- Extract the id from the “chat” JSON object

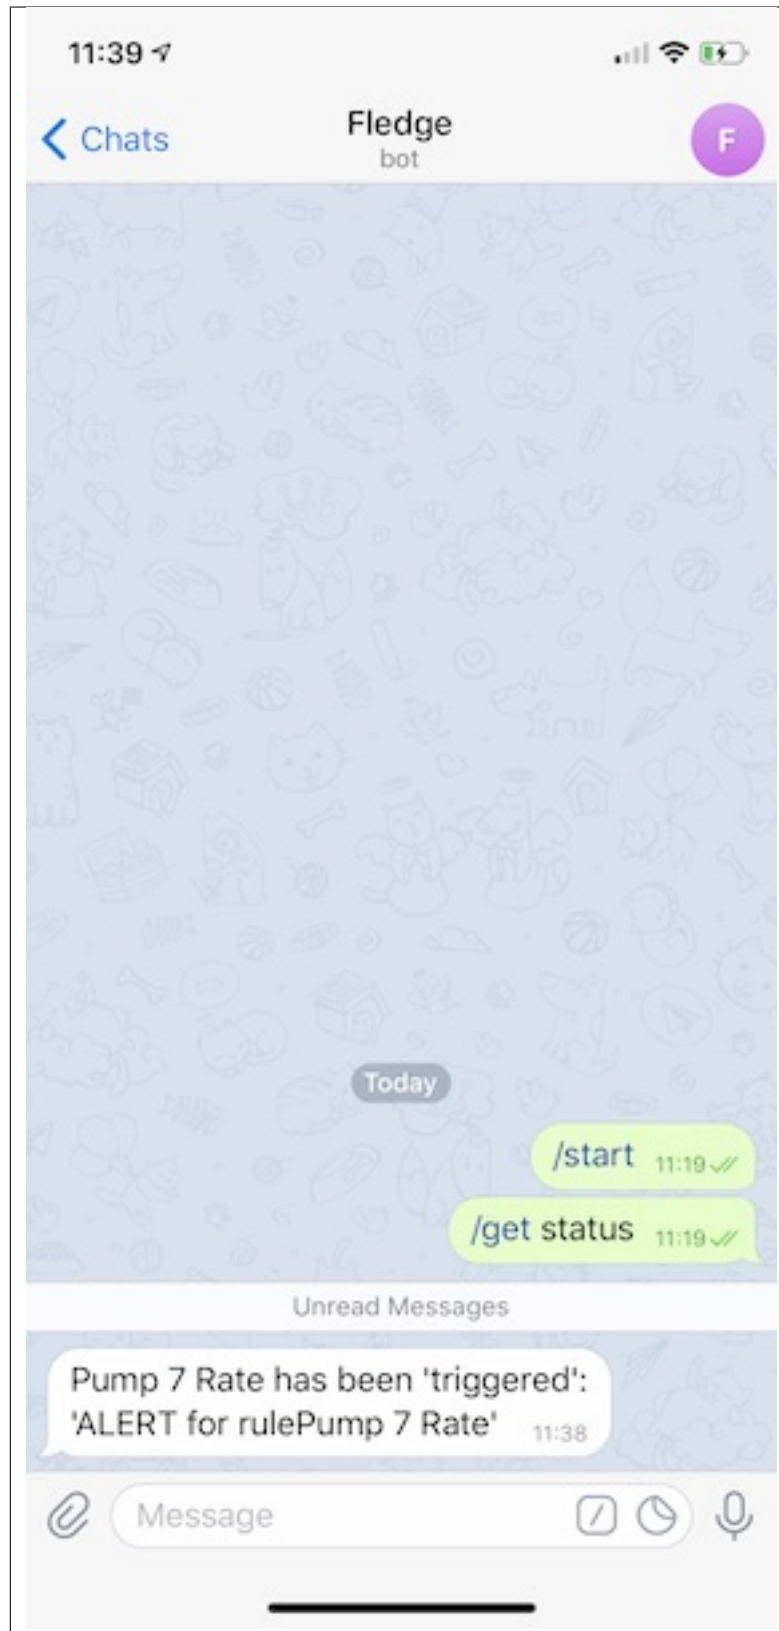
```
{
  "ok": true,
  "result": [
    {
      "update_id": 562812724,
      "message": {
        "message_id": 1,
        "from": {
          "id": 1166366214,
          "is_bot": false,
          "first_name": "Mark",
          "last_name": "Riddoch"
        },
        "chat": {
          "id": 1166366214,
          "first_name": "Mark",
          "last_name": "Riddoch",
          "type": "private"
        },
        "date": 1588328344,
        "text": "start",
        "entities": [
          {
            "offset": 0,
            "length": 6,
            "type": "bot_command"
          }
        ]
      }
    }
  ]
}
```

Once you have created your notification rule and move on to the delivery mechanism

- Select the Telegram plugin from the list of plugins
- Click *Next*

- Configure the delivery plugin
 - **Telegram BOT API token:** Paste the API token you received from botfather
 - **Telegram user chat_id:** Paste the id field form the chat
 - **Telegram BOT API url Prefix:** This is the fixed part of the URL used to send messages and should not be modified under normal circumstances.
- Enable the plugin and click *Next*
- Complete your notification setup

When the notification rule triggers you will receive messages Telegram application



8.5.19 Zendesk Ticket Creation

The *foglamp-notify-zendesk* delivery notification plugin allows notifications to be used to create tickets within Zendesk. The tickets are created within a specified project with a summary, description and other information supplied by FogLAMP.

To obtain an API token from Zendesk

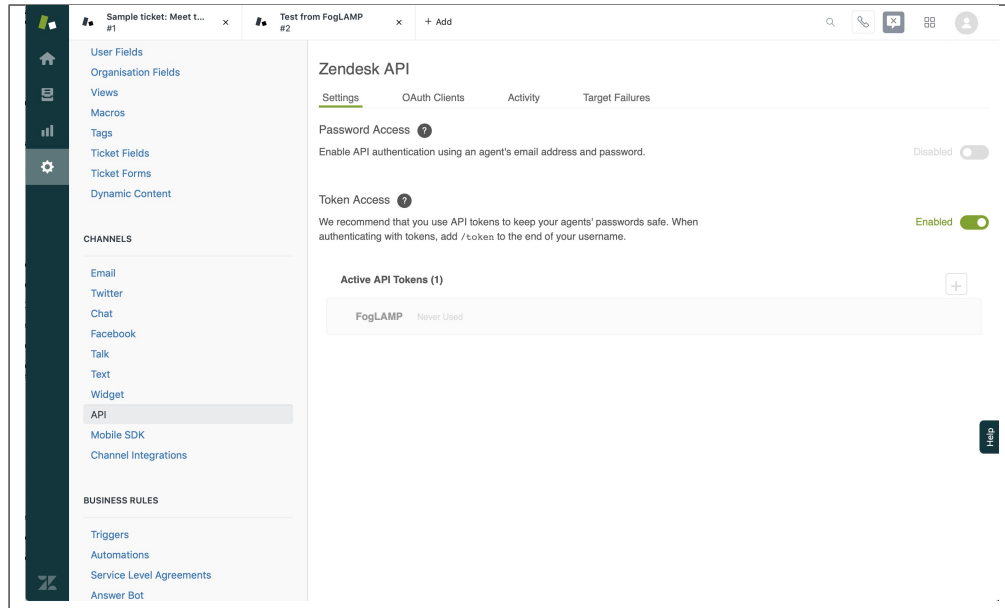
- Visit the page
- Select *Create API token*
- Enter a name for your application, this must be unique for each FogLAMP Zendesk application you create
- Click on Create

Once you have created your notification rule and move on to the delivery mechanism

- Select the *zendesk* plugin from the list of plugins
- Click *Next*

The screenshot displays the 'Delivery Channel' configuration step in the FogLAMP interface. A progress bar at the top indicates the sequence: 1. Notification Instance, 2. Rule, 3. Delivery Channel (active), and 4. Done. The main configuration area includes input fields for 'Subdomain', 'Subject', 'Email', 'API Token' (masked with a password icon), and 'Comment'. Below these is an 'Additional Fields' section with a table containing one row labeled '1' and a code editor icon. An 'Enabled' checkbox is located at the bottom left. Navigation buttons 'Previous' and 'Next' are at the bottom right.

- Configure the delivery plugin
 - **Subdomain:** The subdomain where your Zendesk instance is installed.
 - **Subject:** The subject for the new ticket that is created.
 - **Email:** Your Zendesk registered email address, this is the name of the account you used to create the API token
 - **API Token:** The API token for your email address. You must enable API token in your Zendesk account and create a token for FogLAMP to use.



- **Comment:** The text to add into the comment of the ticket, this may include text substitutions (see below).
- **Additional Fields:** This is a JSON document that contains a number of key/value pairs, each of these pairs is a field name and content to add to the ticket. Text substitutions may be applied here also.
- Enable the plugin and click *Next*
- Complete your notification setup

When the notification rule triggers a Zendesk ticket will be created.

Text Substitution

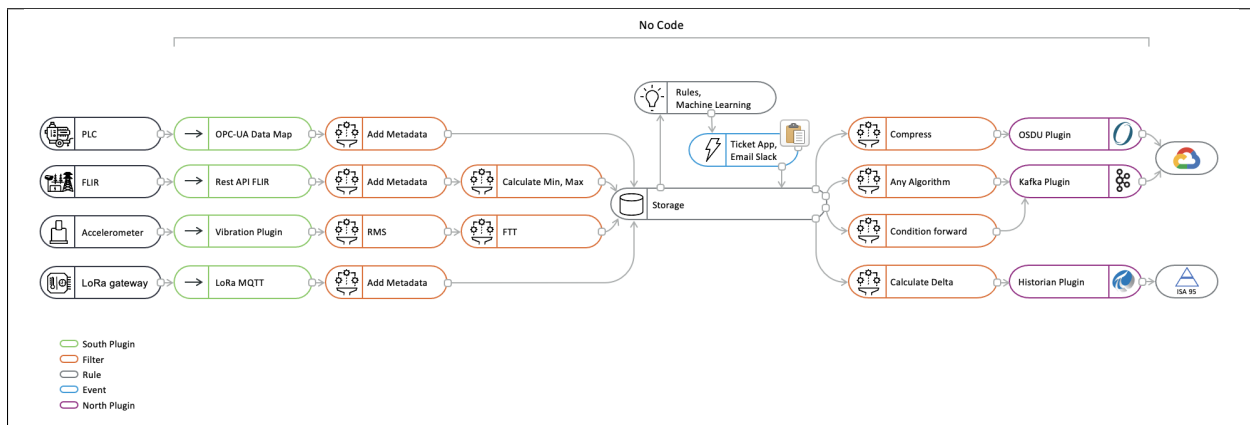
Text markers may be used to substitution text with the fields in the Zendesk ticket. The markers supported are

- **%MESSAGE%:** this is replaced with the message generated in the notification system
- **%REASON%:** this is replaced with the reason for the notification, it may be the string *triggered* or *cleared*.
- **%TIMESTAMP%:** this is replaced with the timestamp of the reading data that caused the notification to trigger.

DEVELOPING DATA PIPELINES

FogLAMP provides a system of data pipelines that allows data to flow from its point of ingest into the FogLAMP instance, the south plugin, to the storage layer in which it is buffered. The stages along this pipeline are FogLAMP processing filters, output of one filter becomes the input of the next. FogLAMP also supports pipelines on the egress as data flows from the storage layer to the north plugins and onward to the systems integrated upstream of the FogLAMP instance.

Operations in the south service are performed on the data from a single source, whilst operations in the north are performed on data going to a single destination. The filter pipeline in the north will have the data from sources flowing through the pipeline, this data will form a mixed stream that will contain all the data in date/time order.



9.1 Best Practices

It is possible with FogLAMP to support multiple data pipelines within a single FogLAMP instance, however if you have a well established FogLAMP instance with critical pipelines running on that instance it is perhaps not always the best practice to then develop a new, experimental pipeline on that same FogLAMP instance.

Looking first at south plugins; one reason for this is that data that enters the FogLAMP instance via your new pipeline will be sent to the storage system and then onward to the north destinations mixed with the data from other pipelines on your system. If your new pipeline is incorrect or generating poor data you are then left in a situation whereby that data has been sent to your existing upstream systems.

If it is unavoidable to use the same instance there are techniques that can be used to reduce the risk; namely to use an to block data from your new south pipeline entering your existing north pipelines and then being sent to your upstream systems. To do this you merely insert a filter at the start of each of your existing north pipelines and set it to exclude the named assets that will be ingested by your new, experimental pipeline. This will allow the data from the existing south service to still flow to the upstream systems, but prevent your new data from streaming out to these systems.

There are still risks associated with this approach, namely that the new service may produce assets of a different name to those you expect or may produce more assets than you expect. Data is still also sent to the notification service from your new pipeline, which may impact that service, although it is less likely than sending incorrect or unwanted data north. There is also the limitation that your new data will be discarded from the buffer and can not then be sent to the existing north pipelines if you subsequently decide the data is good. Data with your new asset names, from your new pipeline, will only be sent once you remove the from those pipelines in the north that send data to your upstream systems.

Developing new north pipelines is less risky, as the data that comes from the storage service and is destined for your new pipeline to upstream systems is effectively duplicated as it leaves the storage system. The main risk is that this new service will count as if the data has been sent up stream as far as the storage system is concerned and may make your data eligible for operation by the purge system sooner than would otherwise be the case. If you wish to prevent this you can update the purge configuration to insist the data is sent on all north channels before being considered sent for the purposes of the purge system. In most circumstances this is a precaution that can be ignored, however if you have configured your FogLAMP system for aggressive purging of data you may wish to consider this.

9.1.1 Incremental Development

The FogLAMP pipeline mechanism is designed for and lends itself to a modular development of the data processing requirement of your application. The pipeline is built from a collection of small, targeted filters that each perform a small, incremental process on the data. When building your pipelines, especially when using the filters that allow the application of scripts to the data, you should consider this approach and not build existing functionality that can be imported by applying an existing filter to the pipeline. Rather use that existing filter and add more steps to your pipeline, the FogLAMP environment is designed to provide minimal overhead when combining filters into a pipeline. Also the pipeline builder can make use of well used and tested filters, thus reducing the overheads to develop and test new functionality that is not needed.

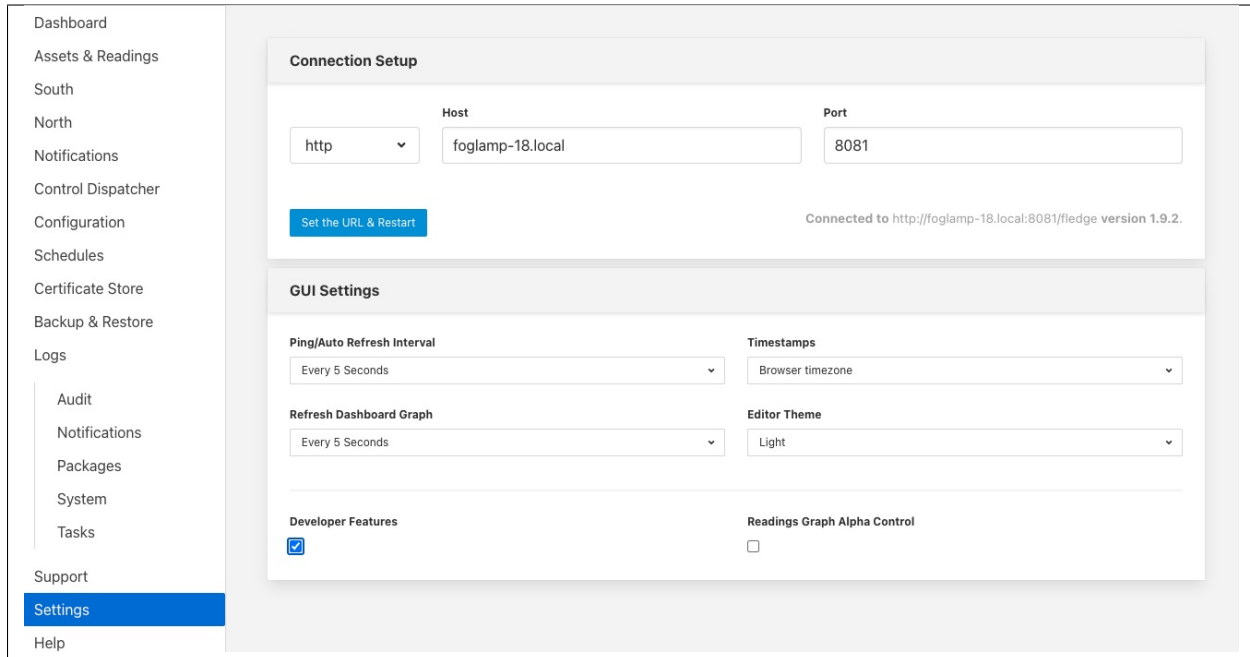
This piecemeal approach can also be adopted in the process of building the pipeline, especially if you use the to block data from progressing further through the FogLAMP system once it has been buffered in the storage layer. Simply add your south service, bring the service up and observe the data that is buffered from the service. You can now add another filter to the pipeline and observe how this alters the data that is being buffered. Since you have a block on the data flowing further within your system, this data will disappear as part of the normal purging process and will not end up in upstream systems to the north of FogLAMP.

If you are developing on a standalone FogLAMP instance, with no existing north services, and you still set your experimental data to disappear, this can be achieved by use of the purge process. Simply configure the purge process to frequently purge data and set the process to purge unsent data. This will mean that the data will remain in the buffer for you to examine for a short time before it is purged from that buffer. Simply adjust the purge interval to allow you enough time to view the data in the buffer. Provided all the experimental data has been purged before you make your system go live, you will not be troubled with your experimental data being sent upstream.

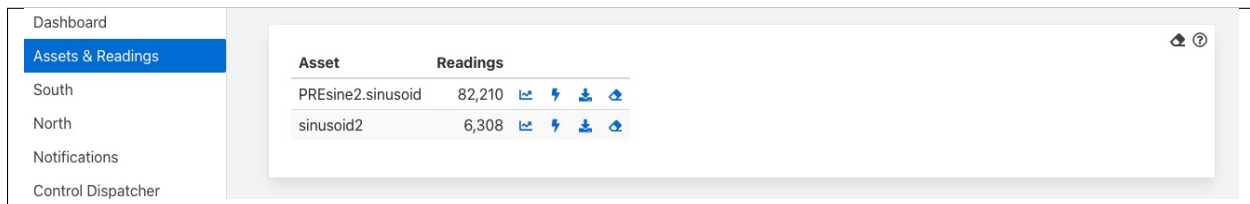
Remember of course to reconfigure the purge process to be more inline with the duration you wish to keep the data for and to turn off the purging of unsent data unless you are willing to loose data that can not be sent for a period of time greater than the purge interval.

Configuring a more aggressive purge system, with the purging of unsent data, is probably not something you would wish to do on an existing system with live data pipelines and should not be used as a technique for developing new pipelines on such a system.

An alternative approach for removing data from the system is to enable the *Developer Features* in the FogLAMP User Interface. This can be done by selecting the *Settings* page in the left hand menu and clicking the option on the bottom of that screen.



Amongst the extra features introduced by selecting *Developer Features* will be the ability to manually purge data from the FogLAMP data store. This on-demand purging can be either applied to a single asset or to all assets within the data store. The manual purge operations are accessed via the *Assets & Readings* item in the FogLAMP menu. A number of new icons will appear when the *Developer Features* are turned on, one per asset and one that impacts all assets.



Asset	Readings				
PREsine2.sinusoid	82,210				
sinusoid2	6,308				




These icons resemble erasers and are located in each row of the assets and also in the top right corner next to the help icon. Clicking on the eraser icon in each of the rows will purge the data for just that asset, leaving other assets untouched. Clicking on the icon in the top right corner will purge all the assets currently in the data store.

In both cases a confirmation dialog will be displayed to ensure against accidental use. If you choose to proceed the selected data within the FogLAMP buffer, either all or a specific asset, will be erased. There is no way to undo this operation or to retrieve the data once it has been purged.

Another consequence that may occur when developing new pipelines is that assets are created during the development process which are not required in the finished pipeline. The asset however remains associated with the service and the asset name and count of number of ingested readings will be displayed in the *South Services* page on the user interface.

sine2	enabled	sinusoid	1.9.2	sine2	734
				sine250	177

It is possible to deprecate the relationship between the service and the asset name using the developer features of the user interface. To do this you must first enable *Developer Features* in the user interface settings page. Now when you view the *South Services* page you will see an eraser icon next to each asset listed for a service.

<u>sine2</u>	enabled	sinusoid	1.9.2	 sine2	734
				 sine250	196

If you click on this icon you will be prompted to deprecate the relationship between the asset and the service. If you select *Yes* the relationship will be severed and the asset will no longer appear next to the service.

Deprecating the relationship will not remove the statistics for the asset, it will merely remove the relationship with the service and hence it will not be displayed against the service.

If an asset relationship is deprecated for an asset that is still in use, it will automatically be reinstated the next time a reading is ingested for that asset. Since the statistics were not deleted when the relationship was deprecated the previous readings will still be included in the statistics when the relationship is restored.

These *Developer Features* are designed to be of use when developing pipelines within FogLAMP, the functionality is not something that should be used in normal operation and the developer features should be turned off when pipelines are not being developed.

Sacrificial North System

Developing north pipelines in a piecemeal fashion can be more of an issue as you are unlikely to want to put poorly formatted data into your upstream systems. One approach to this is to have a sacrificial north system of some type that you can use to develop the pipeline and determine if you are performing the process you need to on that pipeline. This way it is unimportant if that system becomes polluted with data that is not in the form you require it. Ideally you would use a system of the same type to do your development and then switch to the production system when you are satisfied your pipeline is correct.

If this is not possible for some reason a second choice solution would be to use another FogLAMP instance as your test north system. Rather than configure the north plugin you ultimately wish to use you would install the north HTTP plugin and connect this to a second FogLAMP instance running an HTTP plugin. Your data would then be sent to your new FogLAMP instance where you can then examine the data to see what was sent by the first FogLAMP instance. You then build up your north pipeline on that first FogLAMP instance in the same way you did with your south pipeline. Once satisfied you will need to carefully recreate your north pipeline against the correct north plugin and then you may remove your experimental north pipeline and destroy your sacrificial FogLAMP instance that you used to buffer and view the data.

9.1.2 OMF Specific Considerations

Certain north plugins present specific problems to the incremental development approach as changing the format of data that is sent to them can cause them internal issues. The plugin that is used to send data to the Aveva PI Server is one such plugin.


The problem with the PI Server is that it is designed to store data in fixed formats, therefore having data that is not of a consistent type, i.e. made up of the set of attributes, can cause issues. In a PI server each new data type becomes a new tag, this is not a problem if you are happy to use tag naming that is flexible. However if you require that you used fixed name tags within the PI Server, using the filter, this can be an issue for incremental development of your pipeline. Changing the properties of the tag will result in a new name being required for the tag.

The simplest approach is to do all the initial development without the fixed name and then do the name mapping as the final step in developing the pipeline. Although not ideal it gives a relatively simple approach to resolving the problem.

Should you subsequently need to reuse the tag names with different types it becomes necessary to clear the type definitions from the PI Server by removing the element templates, the elements themselves and the cache. The PI Web API will then need to be restarted and the FogLAMP north plugin removed and recreated.

9.1.3 Examining Data

The easiest way to examine your data you have ingested via your new south pipeline is by use of the FogLAMP GUI to examine the data that currently resides within the buffer. You can view the data either via the graph feature of the Assets & Readings page, which will show the time series data.



If you have data that is not timeseries by nature, such as string, you may use the tabular displayed to show you non timeseries data, images if there are any or the download of the data to a spreadsheet view. This later view will not contain any image data in the readings.

sensehat_gyroscope-readings


timestamp	z	x	y
2020-05-04 14:30:49.145006	0.000792725	0.0010765493	0.0022465843
2020-05-04 14:30:48.145022	0.0010982286	-0.0004502609	0.000719551
2020-05-04 14:30:47.145006	0.0007928684	0.0032151192	-0.0011130939
2020-05-04 14:30:46.145008	-0.0013448559	0.0047423765	0.0001088944
2020-05-04 14:30:45.145000	-0.0004286431	0.0007723272	-0.0020291833
2020-05-04 14:30:44.144999	-0.0001233947	0.0013834909	0.0007194807
2020-05-04 14:30:43.145001	-0.000734292	-0.0001437888	0.0004143068

9.1.4 Examining Logs

It is important to view the logs for your service when building a pipeline, this is due to the FogLAMP goal that FogLAMP instances should run as unattended services and hence any errors or warnings generated are written to logs rather than to an interactive user session. The FogLAMP user interface does however provide a number of mechanisms for viewing the log data and filtering it to particular sources. You may view the log from the “System” item in the Log menu and then filter the source to your particular south or north service.

The screenshot shows the FogLAMP user interface. On the left is a sidebar with a navigation menu. Under the 'Logs' section, 'System' is selected. The main area is titled 'System Logs' and includes a search bar and filters for 'Service' (set to 'Simple') and 'Severity' (set to 'Info and above'). A list of log entries is displayed, each starting with a timestamp and followed by a message. Most entries are marked as 'ERROR' and state: 'The supplied Python script does not define a valid "convert" function'. One entry is marked as 'ERROR' and states: 'HTTP error during service registration: 400: A Service with the same name already exists'.

Alternatively if you display the north or south configuration page for your service you will find an icon in the bottom left of the screen that looks like a page of text with the corner folded over. Simply click on this icon and the log screen will be displayed and automatically filtered to view just the logs from the service whose configuration you were previously editing.



Log are displayed with the most recent entry first, with older entries shown as you move down the page. You may move to the next page to view older log entries. It is also possible to view different log severity; fatal, error, warning, info and debug. By default a service will not write info and debug messages to the log, it is possible to turn these levels on via the advanced configuration options of the service. This will then cause the log entries to be written, but before you can view them you must set the appropriate level of severity filtering and the user interface will filter out information and debug message by default.

It is important to turn the logging level back down to warning and above messages once you have finished your debugging session and failure to do this will cause excessive log entries to be written to the system log file.

Also note that the logs are written to the logging subsystem of the underlying Linux version, either syslog or the messages mechanism depending upon your Linux distribution. This means that these log files will be automatically rotated by the operating system mechanisms. This means the system will not, under normal circumstances, fill the storage subsystem. Older log files will be kept for a short time, but will be removed automatically after a few days. This should be borne in mind if you have information in the log that you wish to keep. Also the user interface will only allow you to view data in the most recent log file.

It is also possible to configure the syslog mechanism to write log files to non-standard files or remote machines. The FogLAMP mechanisms for viewing the system logs does require that the standard names for log files are used.

9.1.5 Enabling and Disabling Filters

It should be noted that each filter has an individual enable control, this has the advantage that it is easy to temporarily remove a filter from a pipeline during the development stage. However this does have the downside that it is easy to forget to enable a filter in the pipeline or accidentally add a filter in a disabled state.

9.1.6 Scripting Plugins

Where there is not an existing plugin that does what is required, either in a filter or in south plugins where the data payload of a protocol is highly variable, such as generic REST or MQTT plugins, FogLAMP offers the option of using a scripting language in order to extend the off the shelf plugin set.

This scripting is done via the Python scripting language, both Python 3 and Python 2 are supported by FogLAMP, however it is recommended that the Python 3 variant, be used by preference. The Python support allows external libraries to be used to extend the basic functionality of Python, however it should be noted currently that the Python libraries have to be manually installed on the FogLAMP host machine.

Scripting Guidelines

The user has the full range of Python functionality available to them within the script code they provides to this filter, however caution should be exercised as it is possible to adversely impact the functionality and performance of the FogLAMP system by misusing Python features to the detriment of FogLAMP's own features.

The general principles behind all FogLAMP filters apply to the scripts included in these filters;

- Do not duplicate existing functionality provided by existing filters.
- Keep the operations small and focused. It is better to have multiple filters each with a specific purpose than to create large, complex Python scripts.
- Do not buffer large quantities of data, this will effect the footprint of the service and also slow the data pipeline.

Importing Python Packages

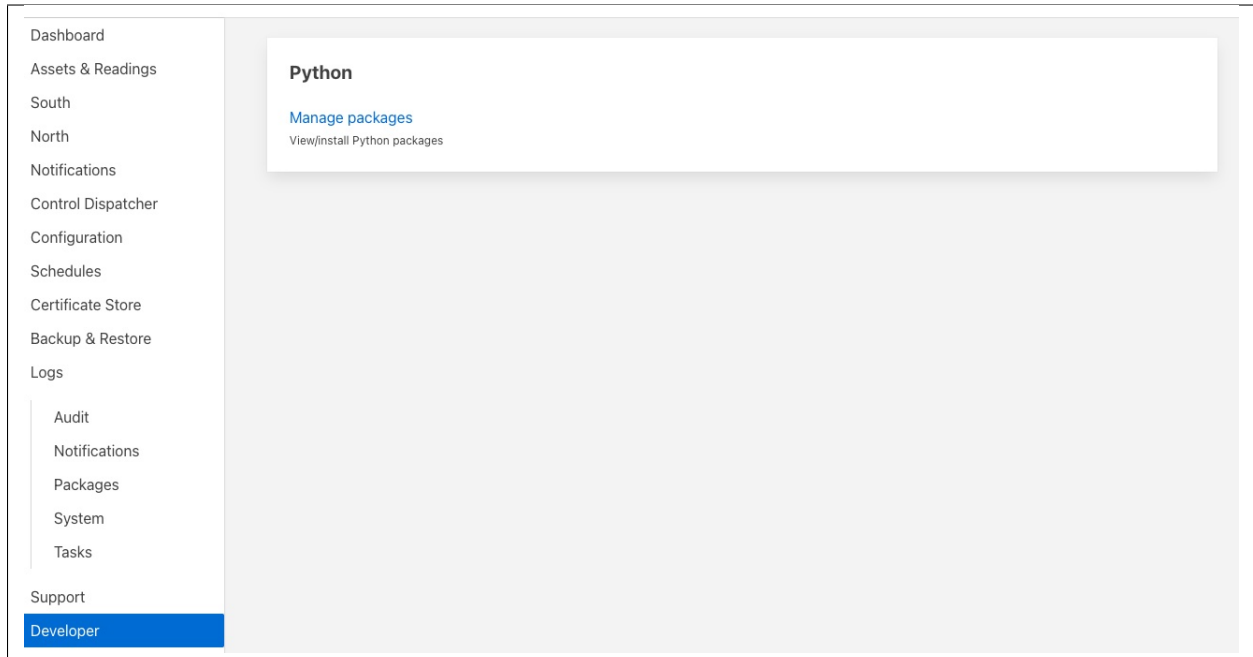
The user is free to import whatever packages they wish in a Python script, this includes the likes of the numpy packages and other that are limited to a single instance within a Python interpreter.

Do not import packages that you do not use or are not required. This adds an extra overhead to the filter and can impact the performance of FogLAMP. Only import packages you actually need.


Python does not provide a mechanism to remove a package that has previously been imported, therefore if you import a package in your script and then update your script to no longer import the package, the package will still be in memory from the previous import. This is because we reload updated scripts without closing down as restarting the Python interpreter. This is part of the sharing of the interpreter that is needed in order to allow packages such as numpy and scipy to be used. This can lead to misleading behavior as when the service gets restarted the package will not be loaded and the script may break because it makes use of the package still.

If you remove a package import from your script and you want to be completely satisfied that the script will still run without it, then you must restart the service in which you are using the plugin. This can be done by disabling and then re-enabling the service.

One of the *Developer Features* of the FogLAMP user interface allows the management of the installed Python Packages from within the user interface. This features is turned on via the *Developer features* toggle in the *Settings* page and will add a new menu item called *Developer*. Navigating to this page will give the the option of managing packages

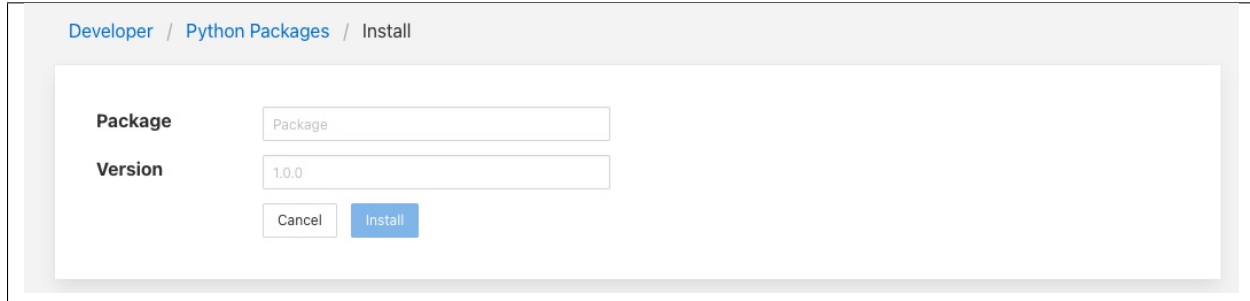


Clicking on *Manage packages* link will display the current set of Python packages that are installed on the machine.



To add a new package click on the *Add +* link in the top right corner. This will display a screen that allows you to

enter details of a Python package to install.



Developer / Python Packages / Install

Package

Version

Enter package name and an optional package version and then click on the *Install* button to install a new package via *pip3*.

Use of Global Variables

You may use global variables within your script and these globals will retain their value between invocations of the of processing function. You may use global variables as a method to keep information between executions and perform such operations as trend analysis based on data seen in previous calls to the filter function.

All Python code within a single service shares the same Python interpreter and hence they also share the same set of global variables. This means you must be careful as to how you name global variables and also if you need to have multiple instances of the same filter in a single pipeline you must be aware that the global variables will be shared between them. If your filter uses global variables it is normally not recommended to have multiple instances of them in the same pipeline.

It is tempting to use this sharing of global variables as a method to share information between filters, this is not recommended as should not be used. There are several reasons for this

- It provides data coupling between filters, each filter should be independent of each other filter.
- One of the filters sharing global variables may be disabled by the user with unexpected consequences.
- Filter order may be changed, resulting in data that is expected by a later filter in the chain not being available.
- Intervening filters may add or remove readings resulting in the data in the global variables not referring to the same reading, or set of readings that it was intended to reference.

If you wish one filter to pass data onto a later filter in the pipeline this is best done by adding data to the reading, as an extra data point. This data point can then be removed by the later filter. An example of this is the way FogLAMP adds to readings that are processed and removed by the north plugin.

For example let us assume we have calculated some value delta that we wish to pass to a later filter, we can add this as a data point to our reading which we will call *_hintDelta*.

```
def myPython(readings):  
    for elem in list(readings):  
        reading = elem['readings']  
        ...  
        reading['_hintDelta'] = delta  
        ...  
    return readings
```

This is far better than using a global as it is attached to the reading to which it refers and will remain attached to that reading until it is removed. It also means that it is independent of the number of readings that are processed per call, and resilient to readings being added or removed from the stream.

The name chosen for this data point in the example above has no significance, however it is good practice to choose a name that is unlikely to occur in the data normally and portrays the usage or meaning of the data.

File IO Operations

It is possible to make use of file operations within a Python filter function, however it is not recommended for production use for the following reasons;

- Pipelines may be moved to other hosts where files may not be accessible.
- Permissions may change dependent upon how FogLAMP systems are deployed in the various different scenarios.
- Edge devices may also not have large, high performance storage available, resulting in performance issues for FogLAMP or failure due to lack of space.
- FogLAMP is designed to be managed solely via the FogLAMP API and applications that use the API. There is no facility within that API to manage arbitrary files within the filesystem.

It is common to make use of files during development of a script to write information to in order to aid development and debugging, however this should be removed, along with associated imports of packages required to perform the file IO, when a filter is put into production.

Threads within Python

It is tempting to use threads within Python to perform background activity or to allow processing of data sets in parallel, however there is an issue with threading in Python, the Python Global Interpreter Lock or GIL. The GIL prevents two Python statements from being executed within the same interpreter by two threads simultaneously. Because we use a single interpreter for all Python code running in each service within FogLAMP, if a Python thread is created that performs CPU intensive work within it, we block all other Python code from running within that FogLAMP service.

We therefore avoid using Python threads within FogLAMP as a means to run CPU intensive tasks, only using Python threads to perform IO intensive tasks, using the asyncio mechanism of Python 3.5.3 or later. In older versions of FogLAMP we used multiple interpreters, one per filter, in order to workaround this issue, however that had the side effect that a number of popular Python packages, such as numpy, pandas and scipy, could not be used as they can not support multiple interpreters within the same address space. It was decided that the need to use these packages was greater than the need to support multiple interpreters and hence we have a single interpreter per service in order to allow the use of these packages.

Interaction with External Systems

Interaction with external systems, using network connections or any form of blocking communication should be avoided in a filter. Any blocking operation will cause data to be blocked in the pipeline and risks either large queues of data accumulating in the case of asynchronous south plugins or data being missed in the case of polled plugins.

Scripting Errors

If an error occurs in the plugin or Python script, including script coding errors and Python exception, details will be logged to the error log and data will not flow through the pipeline to the next filter or into the storage service.

Warnings raised will also be logged to the error log but will not cause data to cease flowing through the pipeline.

See *Examining Logs*: for details have how to access the system logs.

SECURING FOGLEAMP

The default installation of a FogLAMP service comes with security features turned off, there are several things that can be done to add security to FogLAMP. The REST API by default support unencrypted HTTP requests, it can be switched to require HTTPS to be used. The REST API and the GUI can be protected by requiring authentication to prevent users being able to change the configuration of the FogLAMP system. Authentication can be via username and password or by means of an authentication certificate.

10.1 Enabling HTTPS Encryption

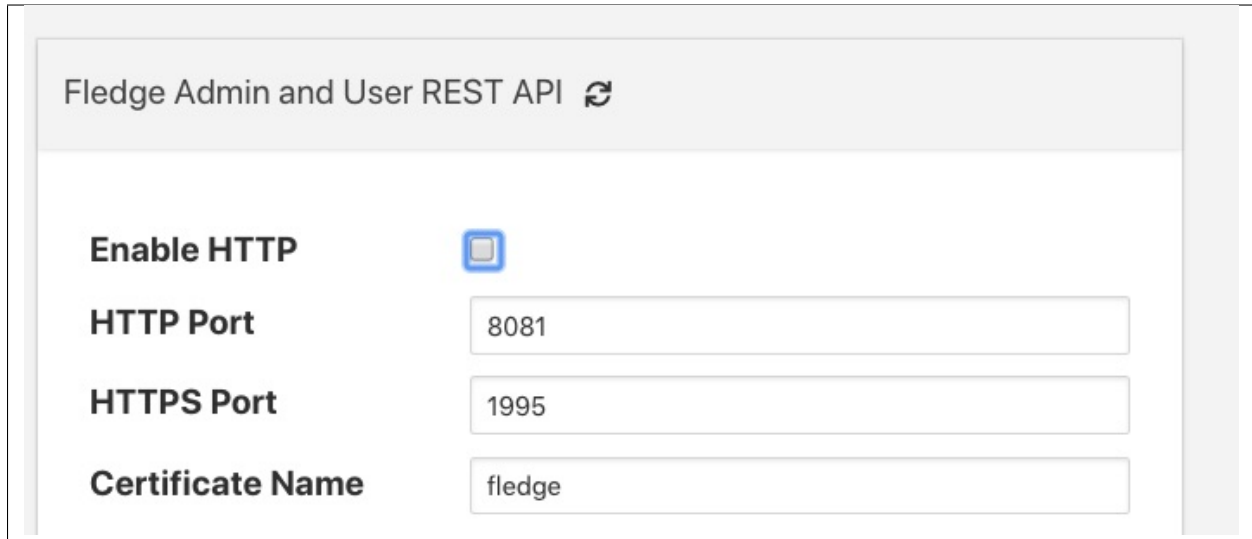
FogLAMP can support both HTTP and HTTPS as the transport for the REST API used for management, to switch between there two transport protocols select the *Configuration* option from the left-hand menu and the select *Admin API* from the configuration tree that appears,

The screenshot shows the FogLAMP web interface. The top bar displays 'FogLAMP', a green status indicator, 'aj/FogLAMP', and system statistics: 'Received: 6,246 Sent: 0 Uptime: 00:22:14'. On the left, a navigation menu includes 'Dashboard', 'Assets & Readings', 'South', 'North', 'Notifications', 'Configuration' (highlighted), 'Schedules', 'Certificate Store', 'Backup & Restore', 'Logs', 'Audit', 'Notifications', 'Packages', 'System', 'Tasks', 'Support', and 'Settings'. The main content area is titled 'FogLAMP Admin and User REST API' and contains the following settings:

- Enable HTTP:** ☒
- HTTP Port:** 8081
- HTTPS Port:** 1995
- Certificate Name:** foglamp
- Authentication:** optional
- Authentication method:** any
- Auth Certificate:** ca
- Allow Ping:** ☒
- Password Expiry Days:** 0
- Auth Providers:** A code editor showing a JSON configuration:

```
1 {  
2   "providers": [  
3     "username",  
4     "ldap"  
5   ]  
6 }
```

The first option you will see is a tick box labeled *Enable HTTP*, to select HTTPS as the protocol to use this tick box should be deselected.



Fledge Admin and User REST API ↻

Enable HTTP ☒

HTTP Port

HTTPS Port

Certificate Name

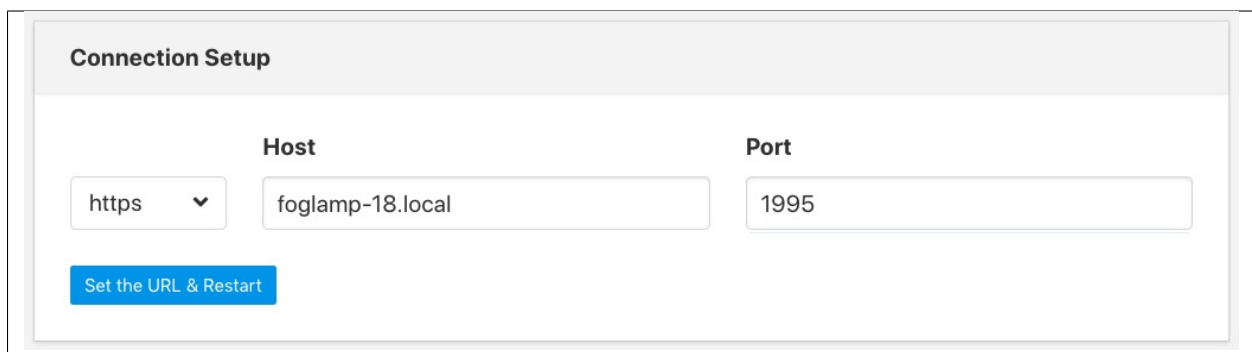
When this is unticked two options become active on the page, *HTTPS Port* and *Certificate Name*. The *HTTPS Port* is the port that FogLAMP will listen on for HTTPS requests, the default for this is port 1995.

The *Certificate Name* is the name of the certificate that will be used for encryption. The default is to use a self signed certificate called *foglamp* that is created as part of the installation process. This certificate is unique per foglamp installation but is not signed by a certificate authority. If you require the extra security of using a signed certificate you may use the FogLAMP [Certificate Store](#) functionality to upload a certificate that has been created and signed by a certificate authority.

After enabling HTTPS and selecting save you must restart FogLAMP in order for the change to take effect. You must also update the connection setting in the GUI to use the HTTPS transport and the correct port.

Note: if using the default self-signed certificate you might need to authorise the browser to connect to IP:PORT. Just open a new browser tab and type the URL `https://YOUR_FOGLAMP_IP:1995`

Then follow browser instruction in order to allow the connection and close the tab. In the FogLAMP GUI you should see the green icon (FogLAMP is running).



Connection Setup

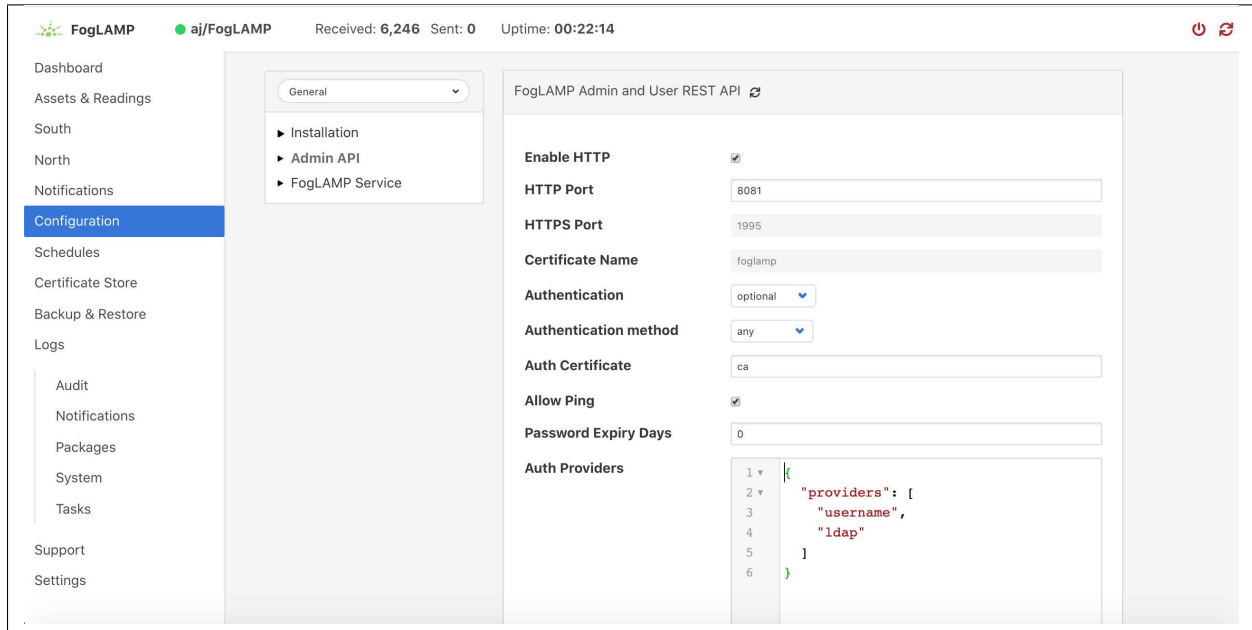
Host **Port**

https


[Set the URL & Restart](#)

10.2 Requiring User Login

In order to set the REST API and GUI to force users to login before accessing FogLAMP select the *Configuration* option from the left-hand menu and then select *Admin API* from the configuration tree that appears.



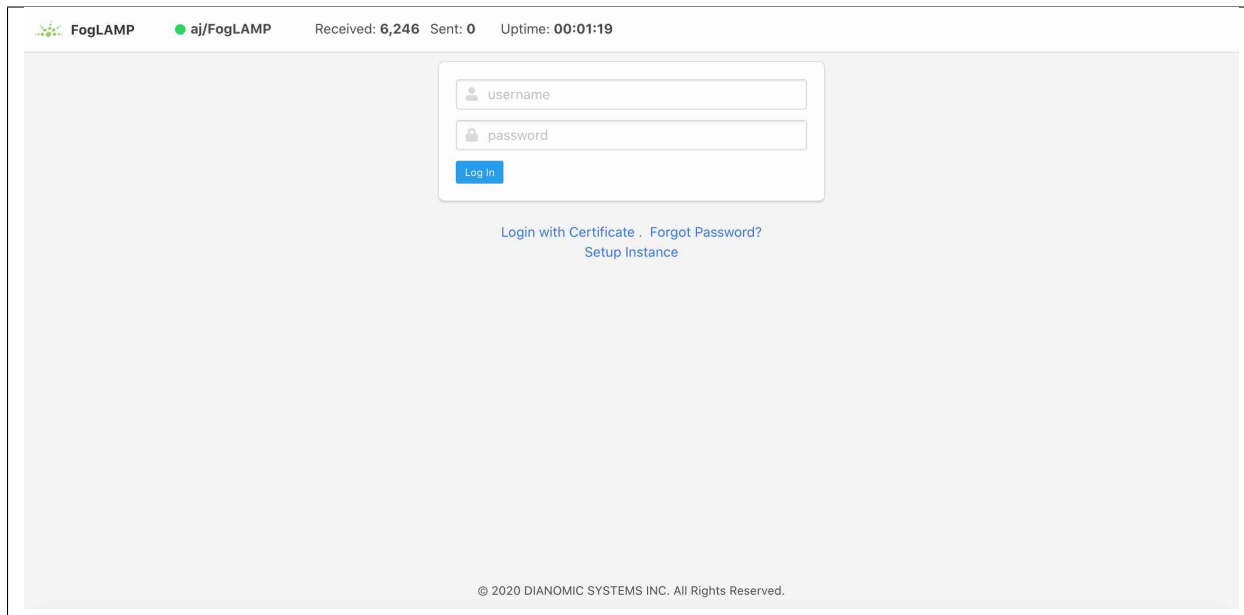
Two particular items are of interest in this configuration category that is then displayed; *Authentication* and *Authentication method*



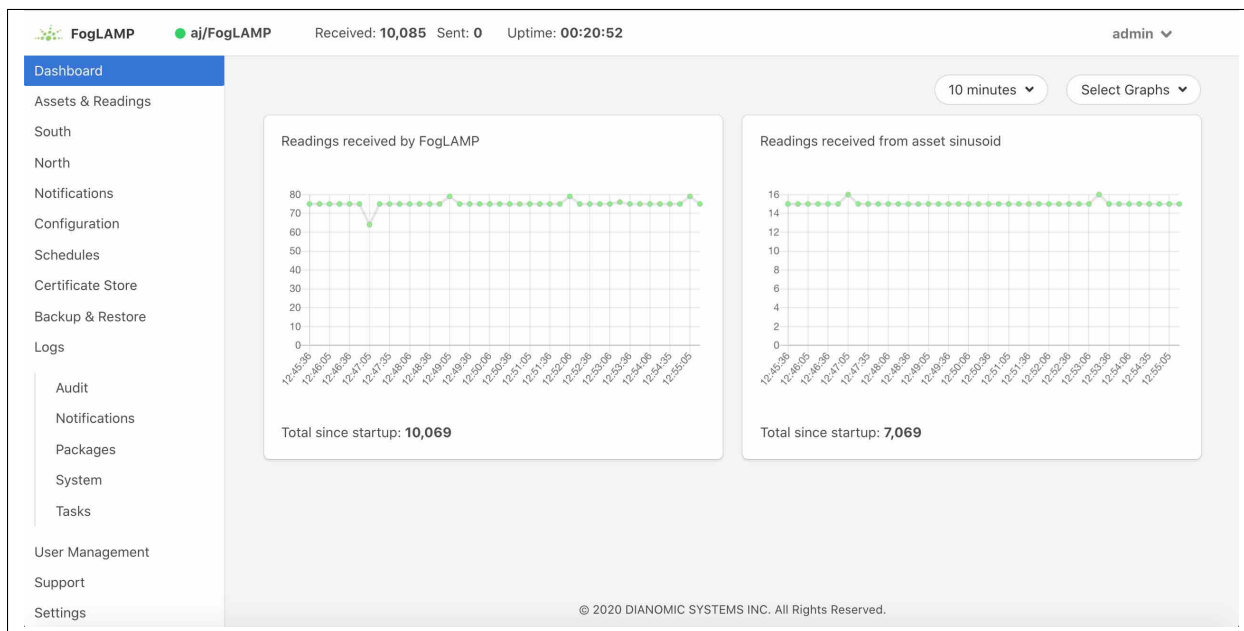
Select the *Authentication* field to be mandatory and the *Authentication method* to be password. Click on *Save* at the bottom of the dialog.

In order for the changes to take effect FogLAMP must be restarted, this can be done in the GUI by selecting the restart item in the top status bar of FogLAMP. Confirm the restart of FogLAMP and wait for it to be restarted.

Once restarted refresh your browser page. You should be presented with a login request.



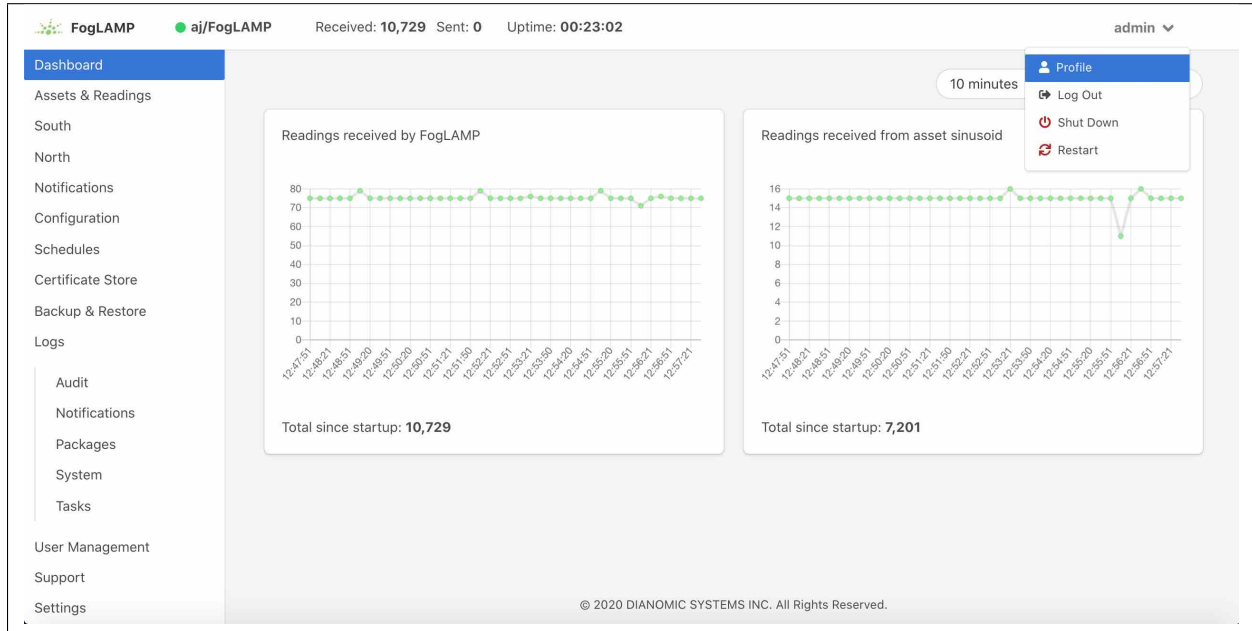
The default username is “admin” with a password of “foglamp”. Use these to login to FogLAMP, you should be presented with a slightly changed dashboard view.



The status bar now contains the name of the user that is currently logged in and a new option has appeared in the left-hand menu, *User Management*.

10.2.1 Changing Your Password

The top status bar of the FogLAMP GUI now contains the user name on the right-hand side and a pull down arrow, selecting this arrow gives a number of options including one labeled *Profile*.



Note: This pulldown menu is also where the *Shutdown* and *Restart* options have moved.

Selecting the *Profile* option will display the profile for the user.



Towards the bottom of this profile display the *change password* option appears. Click on this text and a new password dialog will appear.

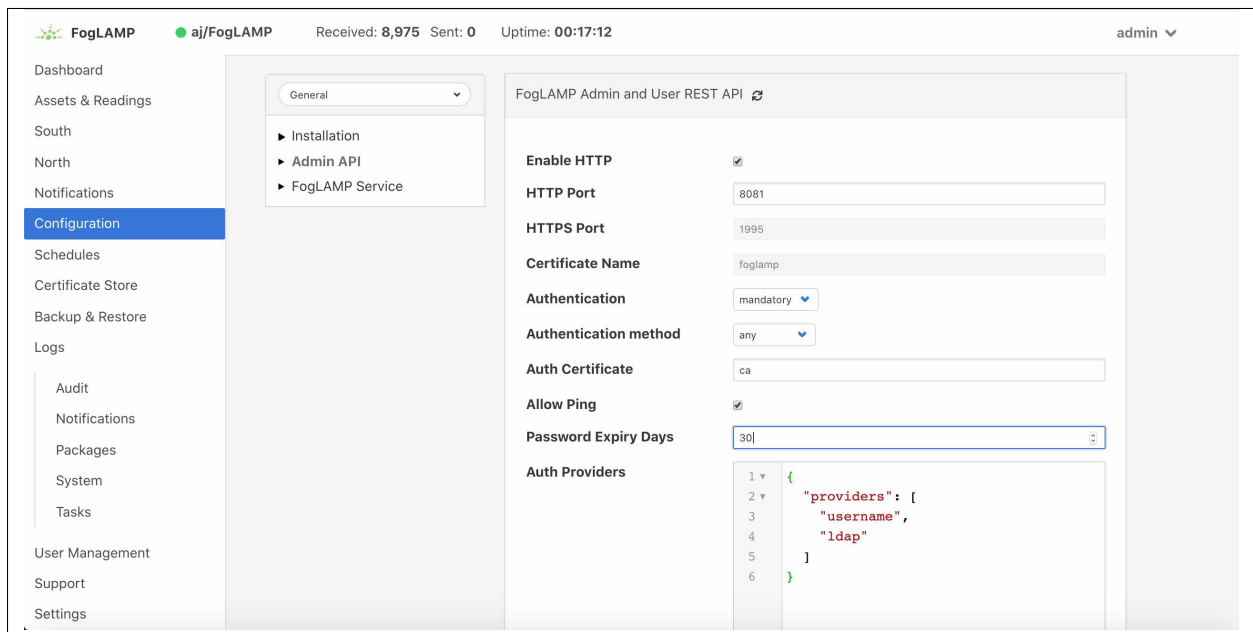


A modal dialog box titled "Reset Password" with a close button (X) in the top right corner. It contains three text input fields stacked vertically, labeled "current password", "new password", and "confirm password". At the bottom left of the dialog is a blue button labeled "Save".

This popup can be used to change your password. On successfully changing your password you will be logged out of the user interface and will be required to log back in using this new password.

10.2.2 Password Rotation Mechanism

FogLAMP provides a mechanism to limit the age of passwords in use within the system. A value for the maximum allowed age of a password is defined in the configuration page of the user interface.



The screenshot shows the FogLAMP web interface. The top navigation bar includes the FogLAMP logo, status indicators (green dot for 'aj/FogLAMP'), system metrics (Received: 8,975, Sent: 0, Uptime: 00:17:12), and a user dropdown (admin). The left sidebar contains a menu with items like Dashboard, Assets & Readings, South, North, Notifications, Configuration (highlighted), Schedules, Certificate Store, Backup & Restore, Logs, Audit, Notifications, Packages, System, Tasks, User Management, Support, and Settings. The main content area is titled "FogLAMP Admin and User REST API" and contains configuration options for the Admin API. These options include:

- Enable HTTP:** checked
- HTTP Port:** 8081
- HTTPS Port:** 1995
- Certificate Name:** foglamp
- Authentication:** mandatory (dropdown)
- Authentication method:** any (dropdown)
- Auth Certificate:** ca
- Allow Ping:** checked
- Password Expiry Days:** 30 (input field)
- Auth Providers:** A code editor showing a JSON configuration:



```

1 {
2   "providers": [
3     "username",
4     "ldap"
5   ]
6 }
```

Whenever a user logs into FogLAMP the age of their password is checked against the maximum allowed password age. If their password has reached that age then the user is not logged in, but is instead forced to enter a new password. They must then login with that new password. In addition the system maintains a history of the last three passwords the user has used and prevents them being reused.

10.3 User Management

Once mandatory authentication has been enabled and the currently logged in user has the role *admin*, a new option appears in the GUI, *User Management*.



The user management pages allows

- Adding new users.
- Deleting users.
- Resetting user passwords.
- Changing the role of a user.

FogLAMP currently supports two roles for users,

- **admin:** a user with admin role is able to fully configure FogLAMP and also manage FogLAMP users
- **user:** a user with this role is able to configure FogLAMP but can not manage users

10.3.1 Adding Users

To add a new user from the *User Management* page select the *Add User* icon in the top right of the *User Management* pane. a new dialog will appear that will allow you to enter details of that user.

A screenshot of a 'Create User' dialog box. The dialog has a title bar with 'Create User' and a close button (X). Inside, there is a section titled 'Role' with a dropdown menu showing 'user'. Below this is a text input field containing 'mark'. There are two password input fields, both masked with dots. The bottom of the dialog features a blue 'Save' button.

Create User

Role

user

mark

.....

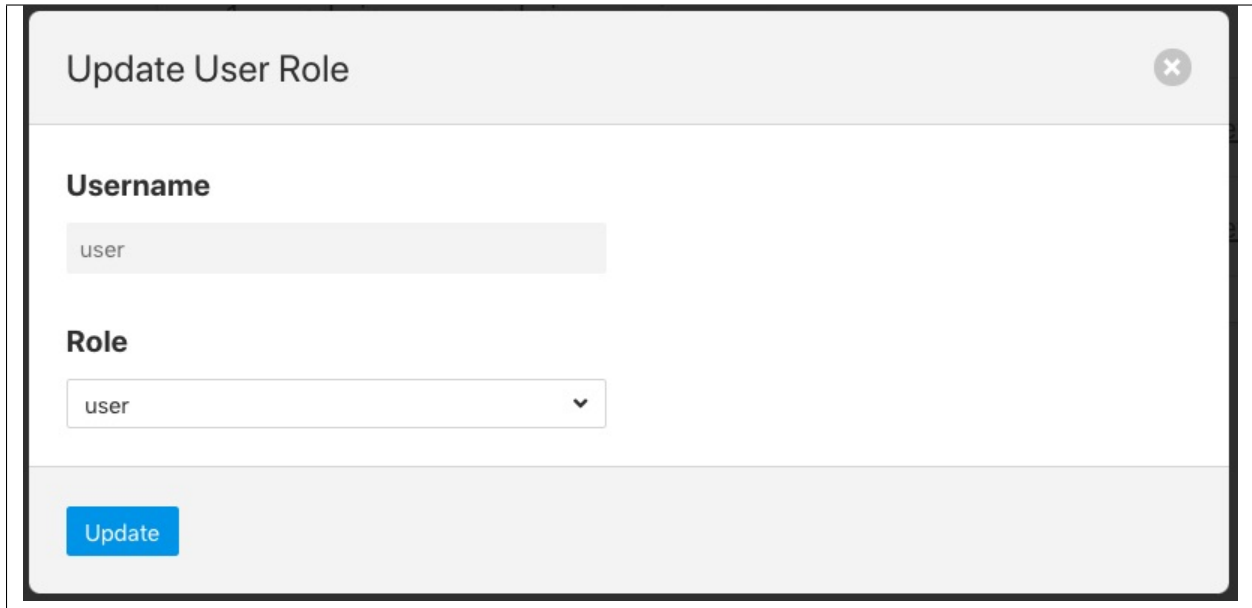
.....

Save

You can select a role for the new user, a user name and an initial password for the user. Only users with the role *admin* can add new users.

10.3.2 Changing User Roles

The role that a particular user has when the login can be changed from the *User Management* page. Simply select on the *change role* link next to the user you wish to change the role of.

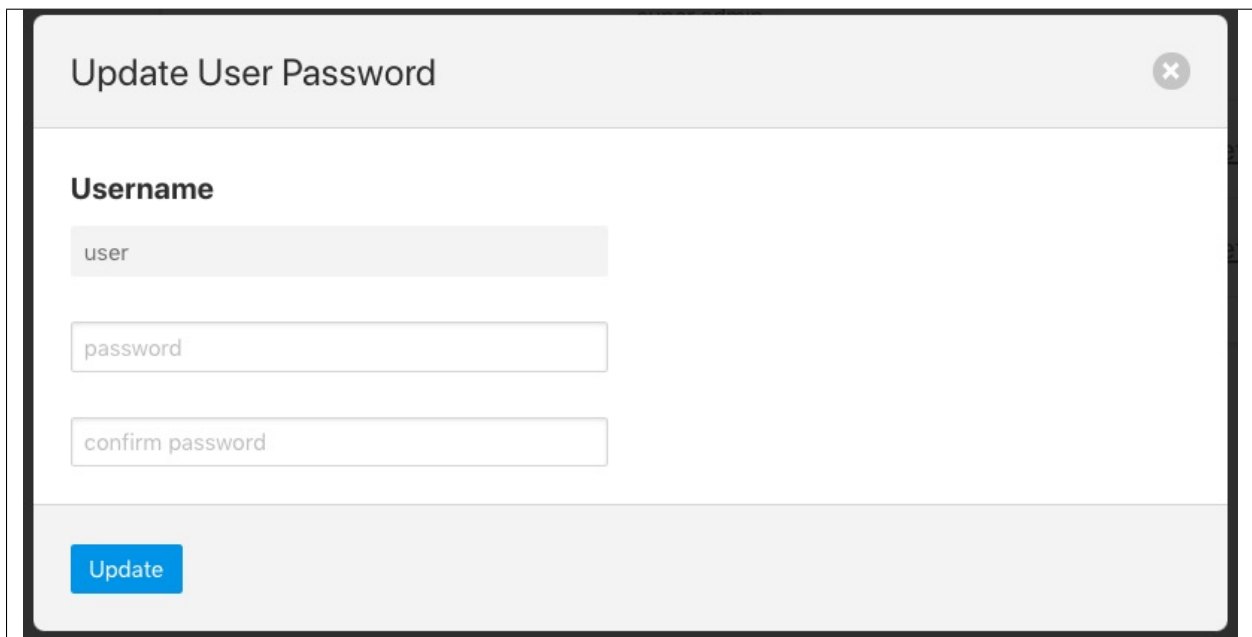


The image shows a modal dialog box titled "Update User Role" with a close button (X) in the top right corner. Inside the dialog, there is a section for "Username" with a text input field containing the value "user". Below this is a section for "Role" with a dropdown menu currently showing "user" and a downward arrow. At the bottom left of the dialog is a blue button labeled "Update".

Select the new role for the user from the drop down list and click on update. The new role will take effect the next time the user logs in.

10.3.3 Reset User Password

Users with the *admin* role may reset the password of other users. In the *User Management* page select the *reset password* link to the right of the user name of the user you wish to reset the password of. A new dialog will appear prompting for a new password to be created for the user.

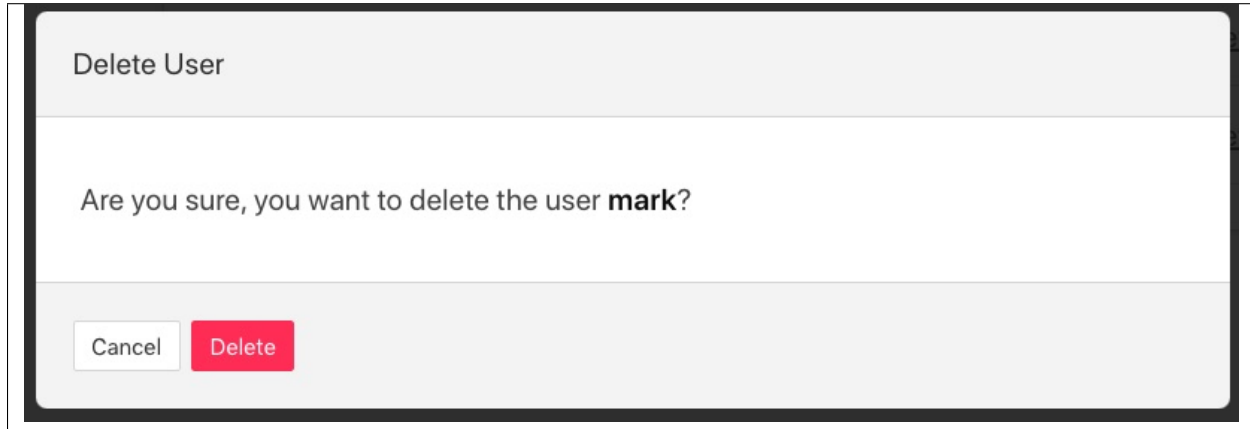


The image shows a modal dialog box titled "Update User Password" with a close button (X) in the top right corner. Inside the dialog, there is a section for "Username" with a text input field containing the value "user". Below this are two text input fields: "password" and "confirm password". At the bottom left of the dialog is a blue button labeled "Update".

Enter the new password and confirm that password by entering it a second time and click on *Update*.

10.3.4 Delete A User

Users may be deleted from the *User Management* page. Select the *delete* link to the right of the user you wish to delete. A confirmation dialog will appear. Select *Delete* and the user will be deleted.



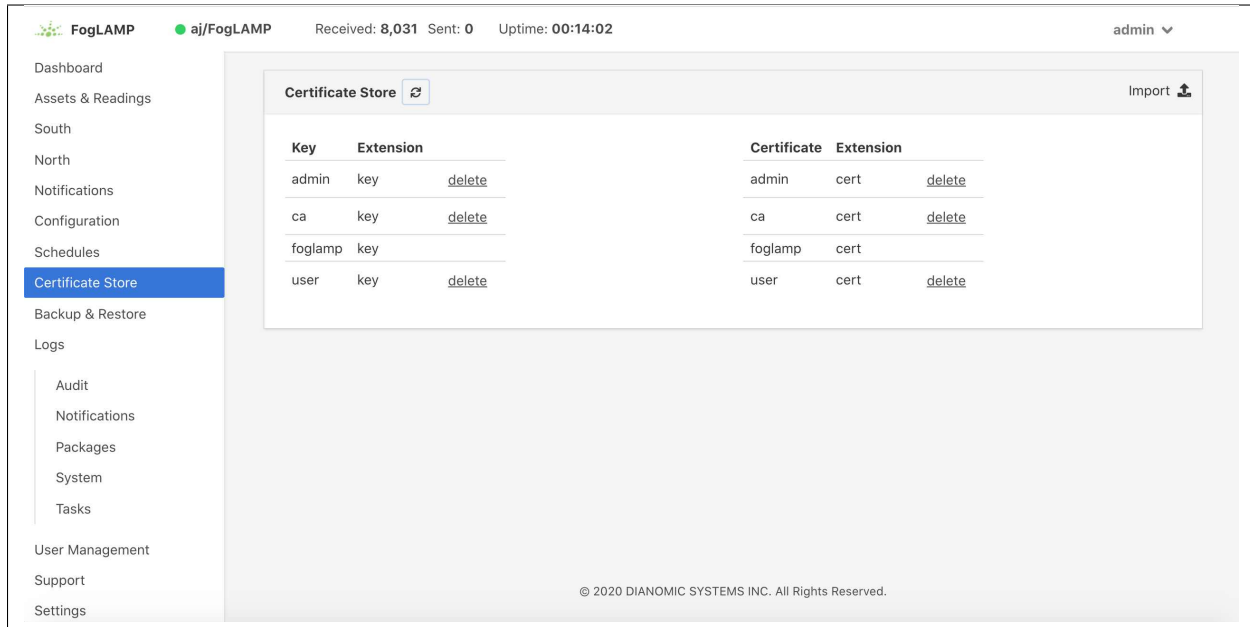
You can not delete the last user with role *admin* as this will prevent you from being able to manage FogLAMP.

10.4 Certificate Store

The FogLAMP *Certificate Store* allows certificates to be stored that may be referenced by various components within the system, in particular these certificates are used for the encryption of the REST API traffic and authentication. They may also be used by particular plugins that require a certificate of one type or another. A number of different certificate types are supported by the certificate store;

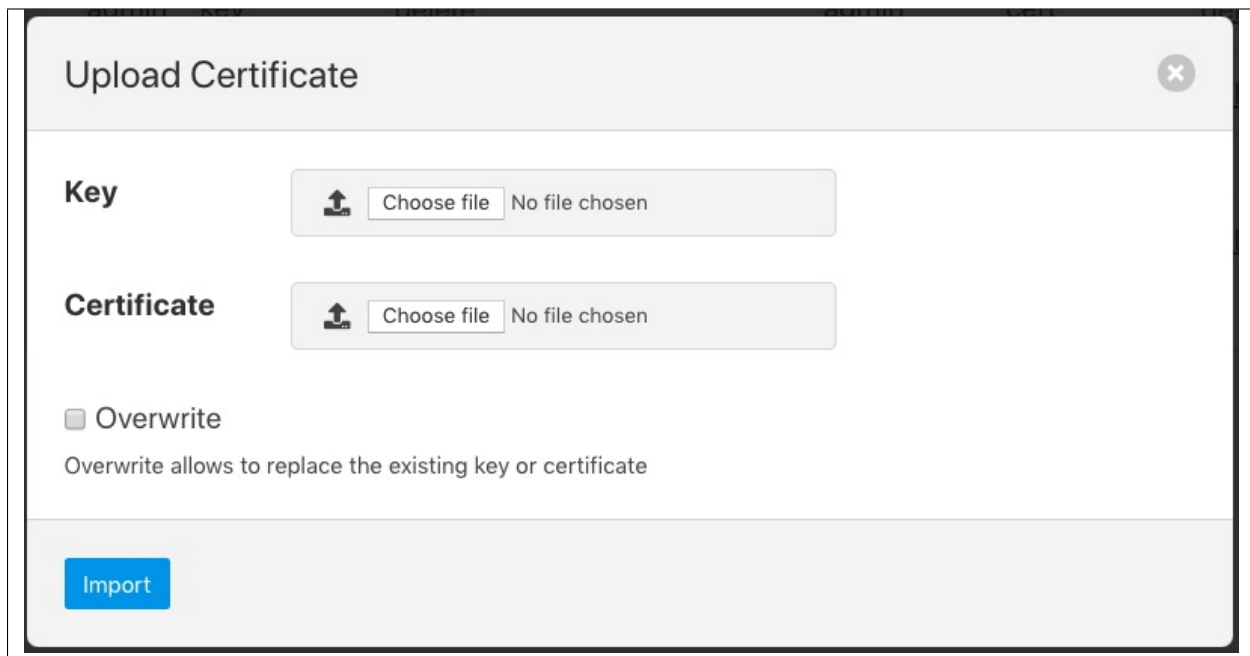
- PEM files as created by most certificate authorities
- CRT files as used by GlobalSign, VeriSign and Thawte
- Binary CER X.509 certificates
- JSON certificates as used by Google Cloud Platform

The *Certificate Store* functionality is available in the left-hand menu by selecting *Certificate Store*. When selected it will show the current content of the store.



Certificates may be removed by selecting the delete option next to the certificate name, note that the keys and certificates can be deleted independently. The self signed certificate that is created at installation time can not be deleted.

To add a new certificate select the *Import* icon in the top right of the certificate store display.



A dialog will appear that allows a key file and/or a certificate file to be selected and uploaded to the *Certificate Store*. An option allows to allow overwrite of an existing certificate. By default certificates may not be overwritten.

TUNING FOGLAMP

Many factors will impact the performance of a FogLAMP system

- The CPU, memory and storage performance of the underlying hardware
- The communication channel performance to the sensors
- The communications to the north systems
- The choice of storage system
- The external demands via the public REST API

Many of these are outside of the control of FogLAMP itself, however it is possible to tune the way FogLAMP will use certain resources to achieve better performance within the constraints of a deployment environment.

11.1 South Service Advanced Configuration

The south services within FogLAMP each have a set of advanced configuration options defined for them. These are accessed by editing the configuration of the south service itself. A screen with a set of tabbed panes will appear, select the tab labeled *Advanced Configuration* to view and edit the advanced configuration options.

lathe1004 South Service

Configuration **Advanced Configuration** Security Configuration

Maximum Reading Latency (mS) 5000

Maximum buffered Readings 100

Reading Rate 1

Throttle ☐

Reading Rate Per second

Minimum Log Level warning

Enabled ☒

Applications

Cancel Save

Service Info http://localhost:45437

Export Readings Delete Service

- *Maximum Reading Latency (mS)* - This is the maximum period of time for which a south service will buffer a reading before sending it onward to the storage layer. The value is expressed in milliseconds and it effectively defines the maximum time you can expect to wait before being able to view the data ingested by this south service.
- *Maximum buffered Readings* - This is the maximum number of readings the south service will buffer before attempting to send those readings onward to the storage service. This and the setting above work together to define the buffering strategy of the south service.
- *Reading Rate* - The rate at which polling occurs for this south service. This parameter only has effect if your south plugin is polled, asynchronous south services do not use this parameter. The units are defined by the setting of the *Reading Rate Per* item.
- *Throttle* - If enabled this allows the reading rate to be throttled by the south service. The service will attempt to poll at the rate defined by *Reading Rate*, however if this is not possible, because the readings are being forwarded out of the south service at a lower rate, the reading rate will be reduced to prevent the buffering in the south service from becoming overrun.
- *Reading Rate Per* - This defines the units to be used in the *Reading Rate* value. It allows the selection of per *second*, *minute* or *hour*.
- *Minimum Log Level* - This configuration option can be used to set the logs that will be seen for this service. It defines the level of logging that is sent to the syslog and may be set to *error*, *warning*, *info* or *debug*. Logs of the level selected and higher will be sent to the syslog. You may access the contents of these logs by selecting

the log icon in the bottom left of this screen.

11.1.1 Tuning Buffer Usage

The tuning of the south service allows the way the buffering is used within the south service to be controlled. Setting the latency value low results in frequent calls to send data to the storage service and therefore means data is more quickly available. However sending small quantities of data in each call the the storage system does not result in the most optimal use of the communications or of the storage engine itself. Setting a higher latency value results in more data being sent per transaction with the storage system and a more efficient system. The cost of this is the requirement for more in-memory storage within the south service.

Setting the *Maximum buffers Readings* value allows the user to place a cap on the amount of memory used to buffer within the south service, since when this value is reach, regardless of the age of the data and the setting of the latency parameter, the data will be sent to the storage service. Setting this to a smaller value allows tighter control on the memory footprint at the cost of less efficient use of the communication and storage service.

Tuning between performance, latency and memory usage is always a balancing act, there are situations where the performance requirements mean that a high latency will need to be incurred in order to make the most efficient use of the communications between the micro services and the transnational performance of the storage engine. Likewise the memory resources available for buffering may restrict the performance obtainable.

11.2 North Advanced Configuration

In a similar way to the south services, north services and tasks also have advanced configuration that can be used to tune the operation of the north side of FogLAMP. The north advanced configuration is accessed in much the same way as the south, select the North page and open the particular north service or task. A tabbed screen will be shown which contains an *Advanced Configuration* tab.

OMF North Service

Configuration **Advanced Configuration** Security Configuration

Minimum Log Level warning

Data block size 100

Enabled ☐

Applications

Cancel Save

Delete Service

- *Minimum Log Level* - This configuration option can be used to set the logs that will be seen for this service or task. It defines the level of logging that is sent to the syslog and may be set to *error*, *warning*, *info* or *debug*. Logs of the level selected and higher will be sent to the syslog. You may access the contents of these logs by selecting the log icon in the bottom left of this screen.
- *Data block size* - This defines the number of readings that will be sent to the north plugin for each call to the *plugin_send* entry point. This allows the performance of the north data pipeline to be adjusted, with larger blocks sizes increasing the performance, by reducing overhead, but at the cost of requiring more memory in the north service or task to buffer the data as it flows through the pipeline. Setting this value too high may cause issues for certain of the north plugins that have limitations on the number of messages they can handle within a single block.

11.3 Health Monitoring

The FogLAMP core monitors the health of other services within FogLAMP, this is done with the *Service Monitor* within FogLAMP and can be configured via the *Configuration* menu item in the FogLAMP user interface. In the configuration page select the *Advanced* options and then the *Service Monitor* section.

The screenshot shows the FogLAMP Configuration page. On the left is a sidebar menu with items: Dashboard, Assets & Readings, South, North, Notifications, Control Dispatcher, Configuration (highlighted), Schedules, Certificate Store, Backup & Restore, and Logs. The main content area has a top bar with 'Advanced' and a dropdown arrow. Below this is a 'Service Monitor' section with a link icon. The settings are as follows:

Setting	Value
Health Check Interval (In seconds)	5
Ping Timeout	1
Max Attempts To Check Heartbeat	15
Restart Failed	auto

A 'Save' button is located at the bottom right of the Service Monitor settings panel.

- *Health Check Interval* - This setting determines how often FogLAMP will send a health check request to each of the microservices within the FogLAMP instance. The value is expressed in seconds. Making this value small will decrease the amount of time it will take to detect a failure, but will increase the load on the system for performing health checks. Making this too frequent is likely to increase the occurrence of false failure detection.
- *Ping Timeout* - Amount of time to wait, in seconds, before declaring that a health check request has failed. Failure for a health check response to be seen within this time will make a service as unresponsive. Small values can result in busy services becoming suspect erroneously.
- *Max Attempts To Check Heartbeat* - This is the number of heartbeat requests that must fail before the core determines that the service has failed and attempts any restorative action. Reducing this value will cause the service to be declared as failed sooner and hence recovery can be performed sooner. If this value is too small then it can result in multiple instances of a service running or frequent restarts occurring. Making this too long results in loss of data.
- *Restart Failed* - Determine what action should be taken when a service is detected as failed. Two options are available, *Manual*, in which case not automatic action will be taken, or *Auto*, in which case the service will be automatically restarted.

11.4 Scheduler

The FogLAMP core contains a scheduler that is used for running periodic tasks, this scheduler has a couple of tuning parameters. To access these parameters from the FogLAMP User Interface, in the configuration page select the *Advanced* options and then the *Scheduler* section.

- *Max Running Tasks* - Specifies the maximum number of tasks that can be running at any one time. This parameter is designed to stop runaway tasks adversely impacting the performance of the system. When this number is reached no new tasks will be created until one or more of the currently running tasks terminated. Set this too low and you will not be able to run all the task you require in parallel. Set it too high and the system is more at risk from runaway tasks.
- *Max Age of Task* - Specifies, in days, how long a task can run for. Tasks that run longer than this will be killed by the system.

Note: Individual tasks have a setting that they may use to stop multiple instances of the same task running in parallel. This also helps protect the system from runaway tasks.

11.5 Storage

The storage layer is perhaps one of the areas that most impacts the overall performance of the FogLAMP instance as it is the end point for the data pipelines; the location at which all ingest pipelines in the south terminate and the point of origin for all north pipelines to external systems.

The storage system in FogLAMP serves two purposes

- The storage of configuration and persistent state of FogLAMP itself
- The buffering of reading data as it traverses the FogLAMP instance

The physical storage is managed by plugins that are loaded dynamically into the storage service in the same way as with other services in FogLAMP. In the case of the storage service it may have either one or two plugins loaded. If a single plugin is loaded this will be used for the storage of both configuration and readings; if two plugins are loaded then one will be used for storing the configuration and the other for storing the readings. Not all plugins support both classes of data.

11.5.1 Choosing A Storage Plugin

FogLAMP comes with a number of storage plugins that may be used, each one has its benefits and limitations, below is an overview of each of the plugins that are currently included with FogLAMP.

sqlite The default storage plugin that is used. It is implemented using the *SQLite* database and is capable of storing both configuration and reading data. It is optimized to allow parallelism when multiple assets are being ingested into the FogLAMP instance. It does however have limitations on the number of different assets that can be ingested within an instance. The precise limit is dependent upon a number of other factors, but is of the order of 900 unique asset names per instance. This is a good general purpose storage plugin and can manage reasonably high rates of data reading.

sqlitelb This is another *SQLite* based plugin able to store both readings and configuration data. It is designed for lower bandwidth data, hence the name suffix *lb*. It does not have the same parallelism optimization as the default *sqlite* plugin, and is therefore less good when high rate data spread across multiple assets is being ingested. However it does perform well when ingesting high rates of a single asset or low rates of a very large number of assets. It does not have any limitations on the number of different assets that can be stored within the FogLAMP instance.

sqlitememory This is a *SQLite* based plugin that uses in memory tables and can only be used to store reading data, it must be used in conjunction with another plugin that will be used to store the configuration. Reading data is stored in tables in memory and thus very high bandwidth data can be supported. If FogLAMP is shutdown however the data stored in these tables will be lost.

postgres This plugin is implemented using the *PostgreSQL* database and supports the storage of both configuration and reading data. It uses the standard Postgres storage engine and benefits from the additional features of Postgres for security and replication. It is capable of high levels of concurrency however has slightly less overall performance than the *sqlite* plugins. Postgres also does not work well with certain types of storage media, such as SD cards as it has a higher wear rate on the media.

In most cases the default *sqlite* storage plugin is perfectly acceptable, however if very high data rates, or huge volumes of data (i.e. large images at a reasonably high rate) are ingested this plugin can start to exhibit issues. This usually exhibits itself by large queues building in the south service or in extreme cases by transaction failure messages in the log for the storage service. If this happens then the recommended course of action is to either switch to a plugin that stores data in memory rather than on external storage, *sqlitememory*, or investigate the media where the data is stored. Low performance storage will adversely impact the *sqlite* plugin.

The *sqlite* plugin may also prove less than optimal if you are ingesting many hundreds of different assets in the same FogLAMP instance. The *sqlite* plugin has been optimized to allow concurrent south services to write to the storage in parallel. This is done by the use of multiple databases to improve the concurrency, however there is a limit, imposed by the number of open databases that can be supported. If this limit is exceeded it is recommended to switch to the *sqlitelb* plugin. There are configuration options regarding how these databases are used that can change the point at which it becomes necessary to switch to the other plugin.

Configuring Storage Plugins

The storage plugins to use can be selected in the *Advanced* section of the *Configuration* page. Select the *Storage* category from the category tree display and the following will be displayed.

The screenshot shows the FogLAMP Configuration page. On the left is a sidebar menu with options: Dashboard, Assets & Readings, South, North, Notifications, Control Dispatcher, Configuration (highlighted), Schedules, Certificate Store, Backup & Restore, and Logs. The main content area is titled 'Storage configuration' and contains a dropdown menu set to 'Advanced'. Below this, a list of categories is shown: Storage (selected), Service Monitor, and Scheduler. The configuration fields include: Storage Plugin (sqlite), Readings Plugin (empty), Database threads (1), Manage Storage (checkbox), Service Port (0), and Management Port (0). A 'Save' button is at the bottom right.

- **Storage Plugin:** The name of the storage plugin to use. This will be used to store the configuration data and must be one of the supported storage plugins.

Note: This can not be the *sqlitememory* plugin as that plugin does not support the storage of configuration.

- **Reading Plugin:** The name of the storage plugin that will be used to store the readings data. If left blank then the *Storage Plugin* above will be used to store both configuration and readings.
- **Database threads:** Increase the number of threads used within the storage service to manage the database activity. This is not the number of threads that can be used to read or write the database and increasing this will not improve the throughput of the data.
- **Manage Storage:** This is used when an external storage application, such as the Postgres database is used that requires separate initialization. If this external process is not run by default setting this to true will cause FogLAMP to start the storage process. Normally this is not required as Postgres should be run as a system service and SQLite does not require it.
- **Service Port:** Normally the storage service will dynamically create a service port that will be used by the storage service. Setting this to a value other than 0 will cause a fixed port to be used. This can be useful when developing a new storage plugin or to allow access to a non-foglamp application to the storage layer. This should only be changed with extreme caution.
- **Management Port:** Normally the storage service will dynamically create a management port that will be used by the storage service. Setting this to a value other than 0 will cause a fixed port to be used. This can be useful when developing a new storage plugin.

Changing will be saved once the *save* button is pressed. FogLAMP uses a mechanism whereby this data is not only saved in the configuration database, but also cached to a file called *storage.json* in the *etc* directory of the data directory. This is required such that FogLAMP can find the configuration database during the boot process. If the configuration becomes corrupt for some reason simply removing this file and restarting FogLAMP will cause the default configuration to be restored. The location of the FogLAMP data directory will depend upon how you installed FogLAMP and the environment variables used to run FogLAMP.

- Installation from a package will usually put the data directory in */usr/local/foglamp/data*. However this can be overridden by setting the *\$FOGLAMP_DATA* environment variable to point at a different location.
- When running a copy of FogLAMP built from source the data directory can be found in

`${FOGLAMP_ROOT}/data`. Again this may be overridden by setting the `$FOGLAMP_DATA` environment variable.

Note: When changing the storage service a reboot of the FogLAMP instance is required before the new storage plugins will be used. Also, data is not migrated from one plugin to another and hence if there is unsent data within the database this will be lost when changing the storage plugin. The `sqlite` and `sqlitelb` plugin however share the same configuration data tables and hence configuration will be preserved when changing between these databases but reading data will not.

sqlite Plugin Configuration

The storage plugin configuration can be found in the *Advanced* section of the *Configuration* page. Select the *Storage* category from the category tree display and the plugin name from beneath that category. In the case of the *sqlite* storage plugin the following will be displayed.

The screenshot shows the FogLAMP configuration interface. On the left is a sidebar menu with options: Dashboard, Assets & Readings, South, North, Notifications, Control Dispatcher, Configuration (highlighted), Schedules, Certificate Store, Backup & Restore, Logs, Audit, and Notifications. The main content area is titled 'Storage Plugin' and shows a configuration form for the 'sqlite' plugin. The form includes the following fields: 'Pool Size' (value 5), 'No. Readings per database' (value 15), 'No. databases to allocate in advance' (value 3), 'Database allocation threshold' (value 1), 'Database allocation size' (value 2), and 'Purge Exclusions' (empty). A 'Save' button is located at the bottom right of the form.

- **Pool Size:** The storage service uses a connection pool to communicate with the underlying database, it is this pool size that determines how many parallel operations can be invoked on the database.

This pool size is only the initial size, the storage service will grow the pool if required, however setting a realistic initial pool size will improve the ramp up performance of FogLAMP.


Note: Although the pool size denotes the number of parallel operations that can take place, database locking considerations may reduce the number of actual operations in progress at any point in time.

- **No. Readings per database:** The *sqlite* plugin support multiple readings databases, with the name of the asset used to determine which database to store the readings in. This improves the level of parallelism by reducing the lock contention when data is being written. Setting this value to 1 will cause only a single asset name to be stored within a single readings database, resulting in no contention between assets. However there is a limit on the number of databases, therefore setting this to 1 will limit the number of different assets that can be ingested into the instance.
- **No. databases to allocate in advance:** This controls how many reading databases FogLAMP should initially created. Creating databases is a slow process and thus is best achieved before data starts to flow through FogLAMP. Setting this too high will cause FogLAMP to allocate a large number of databases than required and waste open database connections. Ideally set this to the number of different assets you expect to ingest divided by the number of readings per database configuration above. This should give you sufficient databases to store the data you require.

- **Database allocation threshold:** The allocation of a new database is a slow process, therefore rather than wait until there are no available databases before allocating new ones, it is possible to pre-allocate database as the number of free databases becomes low. This value allows you to set the point at which to allocate more databases. As soon as the number of free databases declines to this value the plugin will allocate more databases.
- **Database allocation size:** The number of new databases to create whenever an allocation occurs. This effectively denotes the size of the free pool of databases that should be created.
- **Purge Exclusion:** This is not a performance settings, but allows a number of assets to be exempted from the purge process. This value is a comma separated list of asset names that will be excluded from the purge operation.

sqlitelb Configuration

The storage plugin configuration can be found in the *Advanced* section of the *Configuration* page. Select the *Storage* category from the category tree display and the plugin name from beneath that category. In the case of the *sqlitelb* storage plugin the following will be displayed.



Note: The *sqlite* configuration is still present and selectable since this instance has run that storage plugin in the past and the configuration is preserved when switching between *sqlite* and *sqlitelb* plugins.


- **Pool Size:** The storage service uses a connection pool to communicate with the underlying database, it is this pool size that determines how many parallel operations can be invoked on the database.

This pool size is only the initial size, the storage service will grow the pool if required, however setting a realistic initial pool size will improve the ramp up performance of FogLAMP.

Note: Although the pool size denotes the number of parallel operations that can take place, database locking considerations may reduce the number of actual operations in progress at any point in time.

postgres Configuration

The storage plugin configuration can be found in the *Advanced* section of the *Configuration* page. Select the *Storage* category from the category tree display and the plugin name from beneath that category. In the case of the *postgres* storage plugin the following will be displayed.




- **Pool Size:** The storage service uses a connection pool to communicate with the underlying database, it is this pool size that determines how many parallel operations can be invoked on the database.

This pool size is only the initial size, the storage service will grow the pool if required, however setting a realistic initial pool size will improve the ramp up performance of FogLAMP.

Note: Although the pool size denotes the number of parallel operations that can take place, database locking considerations may reduce the number of actual operations in progress at any point in time.

sqlitememory Configuration

The storage plugin configuration can be found in the *Advanced* section of the *Configuration* page. Select the *Storage* category from the category tree display and the plugin name from beneath that category. Since this plugin only supports the storage of readings there will always be at least one other reading plugin displayed. Selecting the *sqlitememory* storage plugin the following will be displayed.



- **Pool Size:** The storage service uses a connection pool to communicate with the underlying database, it is this pool size that determines how many parallel operations can be invoked on the database.

This pool size is only the initial size, the storage service will grow the pool if required, however setting a realistic initial pool size will improve the ramp up performance of FogLAMP.

Note: Although the pool size denotes the number of parallel operations that can take place, database locking considerations may reduce the number of actual operations in progress at any point in time.

TROUBLESHOOTING THE PI SERVER INTEGRATION


This section describes how to troubleshoot issues with the PI Server integration using FogLAMP version $\geq 1.9.1$ and PI Web API 2019 SP1 1.13.0.6518

- *Log files*
- *How to check the PI Web API is installed and running*
- *Commands to check the PI Web API*
- *Error messages and causes*
- *Possible solutions to common problems*

12.1 Log files

FogLAMP logs messages at error and warning levels by default, it is possible to increase the verbosity of messages logged to include information and debug messages also. This is done by altering the minimum log level setting for the north service or task. To change the minimal log level within the graphical user interface select the north service or task, click on the advanced settings link and then select a new minimal log level from the option list presented. The name of the north instance should be used to extract just the logs about the PI Server integration, as in this example:

screenshot from the FogLAMP GUI



```
$ sudo cat /var/log/syslog | grep North_Readings_to_PI
```

Sample message:


```
user.info, 6,1,Mar 15 08:29:57,localhost,FogLAMP, North_Readings_to_PI[15506]: INFO: SendingPro-  
cess is starting
```

Another sample message:

North_Readings_to_PI[20884]: WARNING: Error in retrieving the PIWebAPI version, The PI Web API server is not reachable, verify the network reachability

12.2 How to check the PI Web API is installed and running

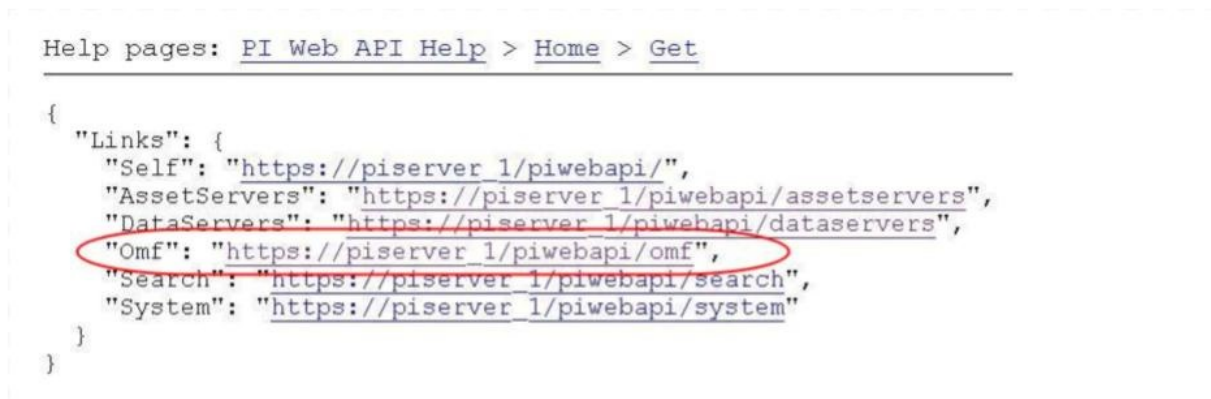
Open the URL `https://piserver_1/piwebapi` in the browser, substituting `piserver_1` with the name/address of your PI Server, to verify the reachability and proper installation of PI Web API. If PI Web API is configured for Basic authentication a prompt, similar to the one shown below, requesting entry of the user name and password will be displayed



NOTE:

- Enter the user name and password which you set in your FogLAMP configuration.

The *PI Web API OMF* plugin must be installed to allow the integration with FogLAMP, in this screenshot the 4th row shows the proper installation of the plugin:



Select the item *System* to verify the installed version:

Help pages: [PI Web API Help](#) > [System](#) > [Landing](#)

```
{
  "ProductTitle": "PI Web API 2019 SP1",
  "ProductVersion": "1.13.0.6518",
  "Links": {
    "Self": "https://piserver\_1/piwebapi/system",
    "CacheInstances": "https://piserver\_1/piwebapi/system/cacheinstances",
    "Configuration": "https://piserver\_1/piwebapi/system/configuration",
    "UserInfo": "https://piserver\_1/piwebapi/system/userinfo",
    "Versions": "https://piserver\_1/piwebapi/system/versions",
    "Status": "https://piserver\_1/piwebapi/system/status",
    "InstanceConfiguration": "https://piserver\_1/piwebapi/system/instanceconfiguration"
  }
}
```

12.3 Commands to check the PI WEB API

Open the PI Web API URL and drill down into the Data Archive and the Asset Framework hierarchies to verify the proper configuration on the PI Server side. Also confirm that the correct permissions have been granted to access these hierarchies.

Data Archive drill down

Following the path *DataServers* -> *Points*:

Help pages: [PI Web API Help](#) > [Home](#) > [Get](#)

```
{
  "Links": {
    "Self": "https://piserver_1/piwebapi/",
    "AssetServers": "https://piserver_1/piwebapi/assetservers",
    "DataServers": "https://piserver_1/piwebapi/dataservers",
    "Omf": "https://piserver_1/piwebapi/omf",
    "Search": "https://piserver_1/piwebapi/search",
    "System": "https://piserver_1/piwebapi/system"
  }
}
```

Help pages: [PI Web API Help](#) > [DataServer](#) > [List](#)

```

"Links": {},
"Items": [
  {
    "WebId": "F1DSqEe26YHa2kewUjKUX4XLfwV01OLTMjMjhNVTFTMDVQ",
    "Id": "e9b647a8-da81-47da-b052-32945f85cb7f",
    "Name": "WIN-3228MU1S05P",
    "Path": "\\\\" + PIServers["WIN-3228MU1S05P"],
    "IsConnected": false,
    "ServerVersion": "",
    "ServerTime": null,
    "Links": {
      "Self": "https://piserver-1/piwabapi/dataservers/F1DSqEe26YHa2kewUjKUX4XLfwV01OLTMjMjhNVTFTMDVQ",
      "Points": "https://piserver-1/piwabapi/dataservers/F1DSqEe26YHa2kewUjKUX4XLfwV01OLTMjMjhNVTFTMDVQ/points",
      "EnumerationSets": "https://piserver-1/piwabapi/dataservers/F1DSqEe26YHa2kewUjKUX4XLfwV01OLTMjMjhNVTFTMDVQ/enumerationsets"
    }
  }
]
}

```

You should be able to browse the *PI Points* page and see your *PI Points* if some data was already sent:

Help pages: [PI Web API Help](#) > [DataServer](#) > [GetPoints](#)

```
{
  "Links": {},
  "Items": [
    {
      "WebId": "F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE",
      "Id": 2935,
      "Name": "427300550797094880_measurement_asset_1",
      "Path": "\\WIN-3229M01S05P\\427300550797094880_measurement_asset_1",
      "PointClass": "classic",
      "PointType": "Float64",
      "DigitalSetName": "",
      "EngineeringUnits": "",
      "Span": 100.0,
      "Zero": 0.0,
      "Step": false,
      "Future": false,
      "DisplayDigits": -5,
      "Links": {
        "Self": "https://piserver 1/piwebapi/points/F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE",
        "DataServer": "https://piserver 1/piwebapi/dataservers/F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE",
        "Attributes": "https://piserver 1/piwebapi/points/F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE/attributes",
        "InterpolatedData": "https://piserver 1/piwebapi/streams/F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE/interpolated",
        "RecordedData": "https://piserver 1/piwebapi/streams/F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE/recorded",
        "PlotData": "https://piserver 1/piwebapi/streams/F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE/plot",
        "SummaryData": "https://piserver 1/piwebapi/streams/F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE/summary",
        "Value": "https://piserver 1/piwebapi/streams/F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE/value",
        "EndValue": "https://piserver 1/piwebapi/streams/F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE/end"
      }
    },
    {
      "WebId": "F1Dq8e26Yha2kew0jKX4KLfwbwAAV010LThy6t1MVTFTMDVQXQyNzNmMDUIMdc5NzcwOTQ4ODBM01PQVNVKVRUS0X0FTU0VUXzE",
      "Id": 2936,
      "Name": "427300550797094880_measurement_asset_2.pl",
      "Path": "\\WIN-3229M01S05P\\427300550797094880_measurement_asset_2.pl"
    }
  ]
}
```

Asset Framework drill down

Following the path *AssetServers* -> Select the *Instance* -> Select the proper *Databases* -> drill down into the AF hierarchy up to the required level -> *Elements*:

Help pages: [PI Web API Help](#) > [Home](#) > [Get](#)

```
{
  "Links": {
    "Self": "https://piserver 1/piwebapi/",
    "AssetServers": "https://piserver 1/piwebapi/assetservers",
    "DataServers": "https://piserver 1/piwebapi/dataservers",
    "Omf": "https://piserver 1/piwebapi/omf",
    "Search": "https://piserver 1/piwebapi/search",
    "System": "https://piserver 1/piwebapi/system"
  }
}
```

selecting the instance

Help pages: [PI Web API Help](#) > [AssetServer](#) > [List](#)

```
{
  "Links": {},
  "Items": [
    {
      "WebId": "FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ",
      "Id": "4afc0662-56f7-496e-aa2-1804b5b04a16",
      "Name": "WIN-3228MU1805P",
      "Description": "",
      "Path": "\\WIN-3228MU1805P",
      "IsConnected": true,
      "ServerVersion": "2.10.8.440",
      "ServerTime": "2021-03-15T13:27:31.0903158Z",
      "ExtendedProperties": {},
      "Links": {
        "Self": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ",
        "Databases": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/assetdatabases",
        "NotificationContactTemplates": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/notificationcontacttemplates",
        "NotificationPlugins": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/notificationplugins",
        "SecurityIdentities": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/securityidentities",
        "SecurityMappings": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/securitymappings",
        "UnitClasses": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/unitclasses",
        "AnalysisRulePlugins": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/analysisruleplugins",
        "TimeRulePlugins": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/timeruleplugins",
        "Security": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/security",
        "SecurityEntries": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ/securityentries"
      }
    },
    {
      "WebId": "FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ",
      "Id": "6745c0ba-5224-40bc-a166-d3c876fde90",
      "Name": "pi-server-1",
      "Description": "",
      "Path": "\\pi-server-1",
      "IsConnected": false,
      "ServerVersion": "",
      "ServerTime": null,
      "ExtendedProperties": {},
      "Links": {
        "Self": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ",
        "Databases": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/assetdatabases",
        "NotificationContactTemplates": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/notificationcontacttemplates",
        "NotificationPlugins": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/notificationplugins",
        "SecurityIdentities": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/securityidentities",
        "SecurityMappings": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/securitymappings",
        "UnitClasses": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/unitclasses",
        "AnalysisRulePlugins": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/analysisruleplugins",
        "TimeRulePlugins": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/timeruleplugins",
        "Security": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/security",
        "SecurityEntries": "https://pi-server-1/piwebapi/assetserver/FIR5uxsFzYsVbChZt08h2_ekAUElTRVJWRVJfMQ/securityentries"
      }
    }
  ]
}
```

selecting the database

Help pages: [PI Web API Help](#) > [AssetServer](#) > [GetDatabases](#)

```
{
  "Links": {},
  "Items": [
    {
      "WebId": "FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04",
      "Id": "de1a62d-9eac-4945-82e1-4a1e9ba988e",
      "Name": "Configuration",
      "Description": "A store for configuration data.",
      "Path": "\\WIN-3228MU1805P\\Configuration",
      "ExtendedProperties": {},
      "Links": {
        "Self": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04",
        "Elements": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/elements",
        "EventFrames": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/eventframes",
        "AssetServer": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ",
        "ElementCategories": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/elementcategories",
        "AttributeCategories": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/attributecategories",
        "TableCategories": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/tablecategories",
        "AnalysisCategories": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/analysiscategories",
        "AnalysisTemplates": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/analysistemplates",
        "EnumerationData": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/enumerations",
        "Tables": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/tables",
        "Security": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/security",
        "SecurityEntries": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLabh3qyeRhmC40ue6qdgjgV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/securityentries"
      }
    },
    {
      "WebId": "FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04",
      "Id": "a395ed2d-58c8-49e1-9cd1-6fce75e11e34",
      "Name": "FogLAMP",
      "Description": "",
      "Path": "\\WIN-3228MU1805P\\FogLAMP",
      "ExtendedProperties": {},
      "Links": {
        "Self": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04",
        "Elements": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/elements",
        "EventFrames": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/eventframes",
        "AssetServer": "https://pi-server-1/piwebapi/assetserver/FIR5Ygb8SvdWbkmg4hgEtbBKFgV010LTHyMjMNVFTMDVQ",
        "ElementCategories": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/elementcategories",
        "AttributeCategories": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/attributecategories",
        "TableCategories": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/tablecategories",
        "AnalysisCategories": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/analysiscategories",
        "AnalysisTemplates": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/analysistemplates",
        "EnumerationData": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/enumerations",
        "Tables": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/tables",
        "Security": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/security",
        "SecurityEntries": "https://pi-server-1/piwebapi/assetdatabases/FIR0Ygb8SvdWbkmg4hgEtbBKFgLeZv08hY4UmcW_0dWEeNAV010LTHyMjMNVFTMDVQKBNFTkZJR1VQVbJT04/securityentries"
      }
    }
  ]
}
```

Proceed with the drill down operation up to the desired level/asset.

12.4 Error messages and causes

Some error messages and causes:

Message	Cause
North_Readings_to_PI[20884]: WARNING: Error in retrieving the PIWebAPI version, The PI Web API server is not reachable , verify the network reachability	FogLAMP is not able to reach the machine in which PI Server is running due to a network problem or a firewall restriction.
North_Readings_to_PI[5838]: WARNING: Error in retrieving the PIWebAPI version, 503 Service Unavailable	FogLAMP is able to reach the machine in which PI Server is executing but the PI Web API is not running.
North_Readings_to_PI[24485]: ERROR: Sending JSON data error : Container not found. 4273005507977094880_1measurement_sin_4816_asset_1 - WIN-4M7ODKB0RH2:443 /piwebapi/omf	FogLAMP is able to interact with PI Web API but there is an attempt to store data in a PI Point that does not exist.

12.5 OMF Plugin Data

The OMF north plugin must create type information within the OMF subsystem of the PI Server before any data can be sent. This type information is persisted within the PI Server between sessions and must also be persisted within FogLAMP for each connection to a PI Server. This is done using the plugin data persistence features of the FogLAMP north plugin.

This results in an important connection between a north service or task and a PI Server, which does add extra constraints as to what may be done at each end. It is very important this data is kept synchronized between the two ends. In normal circumstances this is not a problem, but there are some actions that can cause problems and require action on both ends.

Delete a north service or task using the OMF plugin If a north service or task using the OMF plugin is deleted then the persisted data of the plugin is also lost. This is FogLAMP's record of what types have been created in the PI Server and is no longer synchronized following the deletion of the north service. Any new service or task that is created and connected to the same PI Server will receive duplicate type errors from the PI Server. There are two possible solutions to this problem;

- Remove the type data from the PI Server such that neither end has the type information.
- Before deleting the north service or task export the plugin persisted data and import that data into the new service or task.

Cleanup a PI Server and reuse and existing OMF North service or task This is the opposite problem to that stated above, the plugin will try to send data thinking that the types have already been created in the PI Server and receive an error. FogLAMP will automatically correct for this and create new types. These new types however will be created with new names, which may not be the desired behavior. Type names are created using a fixed algorithm. To re-use the previous names, stopping the north service and deleting the plugin persisted data will reset the algorithm and recreate the types using the names that had been previously used.

Taking an existing FogLAMP north task or service and moving it to a new PI Server This new PI Server will not have the type information from the old and we will once again get errors when sending data due to these missing types. FogLAMP will automatically correct for this and create new types. These new types however will be created with new names, which may not be the desired behavior. Type names are created using a fixed algorithm. To re-use the previous names, stopping, the north service and deleting the plugin persisted data will reset the algorithm and recreate the types using the names that had been previously used.

12.5.1 Managing Plugin Persisted Data

This is not a feature that users would ordinarily need to be concerned with, however it is possible to enable *Developer Features* in the FogLAMP User Interface that will provide a mechanism to manage this data.

Enable Develop Features

Navigate to the *Settings* page of the GUI and toggle on the *Developer Features* check box on the bottom left of the page.

Viewing Persisted Data


In order to view the persisted data for the plugins of a service open either the *North* or *South* page on the user interface and select your service or task. An page will open that allows you to update the configuration of the plugin. This contains a set of tabs that may be selected, when *Developer Features* are enabled one of these tabs will be labeled *Developer*.

The screenshot shows the 'OMF North Service' configuration page in the FogLAMP GUI. The 'Developer' tab is active, showing a list of configuration parameters. The 'Configuration' tab is also visible, and the 'Advanced Configuration' and 'Security Configuration' tabs are disabled. The 'Developer' tab contains the following settings:


Configuration Parameter	Value
Endpoint	PI Web API
Send full structure	<input checked="" type="checkbox"/>
Naming Scheme	Concise
Server hostname	54.160.93.60
Server port, 0=use the default	0
Producer Token	omf_north_0001
Data Source	readings
Static Data	
Sleep Time Retry	1
Maximum Retry	3
HTTP Timeout	10
Integer Format	int64
Number Format	float64
Compression	<input type="checkbox"/>
Default Asset Framework Location	/fledge/data_piwebapi/mark
Asset Framework hierarchy rules	1 {}

The *Developer* tab will allow the viewing of the persisted data for all of the plugins in that service, filters and either north or south plugins, for which data is persisted.


Persisted data is only written when a plugin is shutdown, therefore in order to get the most up to date view of the data it is recommended that service is disabled before viewing the persisted data. It is possible to view the persisted data of a running service, however this will be a snapshot taken from the last time the service was shutdown.



It is possible for more than one plugin within a pipeline to persist data, in order to select between the plugins that have persisted data a menu is provided in the top left which will list all those plugins for which data can be viewed.



As well as viewing the persisted data it is also possible to perform other actions, such as *Delete*, *Export* and *Import*. These actions are available via a menu that appears in the top right of the screen.



Note: The service must be disabled before use of the Delete or Import features and to get the latest values when performing an Export.

12.5.2 Understanding The OMF Persisted Data

The persisted data takes the form of a JSON document, the following is an example for a FogLAMP instance configured with just the Sinusoid plugin.

```
{
  "sentDataTypes": [
    {
      "sinusoid": {
        "type-id": 1,
        "dataTypesShort": "0x101",
        "hintChecksum": "0x0",
        "namingScheme": 0,
        "afhHash": "15489826335467873671",
        "afHierarchy": "foglamp/data_piwebapi/mark",
        "afHierarchyOrig": "foglamp/data_piwebapi/mark",
        "dataTypes": {
          "sinusoid": {
            "type": "number",
            "format": "float64"
          }
        }
      }
    }
  ]
}
```

The *SentDataTypes* is a JSON array of object, with each object representing one data type that has been sent to the PI

Server. The key/value pairs within the object are as follow

Key	Description
type-id	An index of the different types sent for this asset. Each time a new type is sent to the PI Server for this asset this index will be incremented.
dataType-sShort	A summary of the types in the datatypes of the asset. The value is an encoded number that contains the count of each of base types, integer, float and string, in the datapoints of this asset.
hintChecksum	A checksum of the OMFHints used to create this type. 0 if no OMF Hint was used.
nam-ingScheme	The current OMF naming scheme when the type was sent.
afhHash	A Hash of the AF settings for the type.
afHierarchy	The AF Hierarchy location.
afHierarchyOrig	The original setting of AF Hierarchy. This may differ from the above if specific AF rules are in place.
dataTypes	The data type sent to the PI Server. This is an actually OMF type definition and is the exact type definition sent to the PI Web API endpoint.

12.6 Possible solutions to common problems

Recreate a single or a sets of PI Server objects and resend all the data for them to the PI Server on the Asset Framework hierarchy

procedure:

- disable the 1st north instance
- delete the objects in the PI Server, AF + Data archive, that are to be recreated or were partially sent.
- create a new **DISABLED** north instance using a new, unique name and having the same AF hierarchy as the 1st north instance
- install *foglamp-filter-asset* on the new north instance
- configure *foglamp-filter-asset* with a rule like the following one

```
{
  "rules": [
    {
      "asset_name": "asset_4",
      "action": "include"
    }
  ],
  "defaultAction": "exclude"
}
```

- enable the 2nd north instance
- let the 2nd north instance send the desired amount of data and then disable it
- enable the 1st north instance

note:

- the 2nd north instance will be used only to recreate the objects and resend the data
- the 2nd north instance will resend all the data available for the specified *included* assets

- there will be some data duplicated for the recreated assets because part of the information will be managed by both the north instances

Recreate all the PI Server objects and resend all the data to the PI Server on a different Asset Framework hierarchy level**procedure:**

- disable the 1st north instance
- create a new north instance using a new, unique name and having a new AF hierarchy (North option 'Asset Framework hierarchies tree')

note:

- this solution will create a set of new objects unrelated to the previous ones
- all the data stored in FogLAMP will be sent

Recreate all the PI Server objects and resend all the data to the PI Server on the same Asset Framework hierarchy level of the 1**procedure:**

- disable the 1st north instance
- delete properly the objects on the PI Server, AF + Data archive, that were eventually partially deleted
- stop / start PI Web API
- create a new north instance 2nd using the same AF hierarchy (North option 'Asset Framework hierarchies tree')

note:

- all the types will be recreated on the PI Server. If the structure of each asset, number and types of the properties, does not change the data will be accepted and laced into the PI Server without any error. PI Web API 2019 SP1 1.13.0.6518 will accept the data.
- Using PI Web API 2019 SP1 1.13.0.6518 the PI Server creates objects with the compression feature disabled. This will cause any data that was previously loaded and is still present in the Data Archive, to be duplicated.

Recreate all the PI Server objects and resend all the data to the PI Server on the same Asset Framework hierarchy level of the 1**procedure:**

- disable the 1st north instance
- delete all the objects on the PI Server side, both in the AF and in the Data Archive, sent by the 1st north instance
- stop / start PI Web API
- create a new north instance using the same AF hierarchy (North option 'Asset Framework hierarchies tree')

note:

- all the data stored in FogLAMP will be sent

PLUGIN DEVELOPER GUIDE

FogLAMP makes extensive use of plugin components to extend the base functionality of the platform. In particular, plugins are used to;

- Extend the set of sensors and actuators that FogLAMP supports.
- Extend the set of services to which FogLAMP will push accumulated data gathered from those sensors.
- The mechanism by which FogLAMP buffers data internally.
- Filter plugins may be used to augment, edit or remove data as it flows through FogLAMP.
- Rule plugins extend the rules that may trigger the delivery of notifications at the edge.
- Notification delivery plugins allow for new delivery mechanisms to be integrated into FogLAMP.

This chapter presents the plugins that are bundled with FogLAMP, how to write and use new plugins to support different sensors, protocols, historians and storage devices. It will guide you through the process and entry points that are required for the various different types of plugin.

There are also numerous plugins that are available as separate packages or in separate repositories that may be used with FogLAMP.

13.1 Plugins

In this version of FogLAMP you have six types of plugins:

- **South Plugins** - They are responsible for communication between FogLAMP and the sensors and actuators they support. Each instance of a FogLAMP South microservice will use a plugin for the actual communication to the sensors or actuators that that instance of the South microservice supports.
- **North Plugins** - They are responsible for taking reading data passed to them from the South bound service and doing any necessary conversion to the data and providing the protocol to send that converted data to a north-side task.
- **Storage Plugins** - They sit between the Storage microservice and the physical data storage mechanism that stores the FogLAMP configuration and readings data. Storage plugins differ from other plugins in that they are written exclusively in C/C++, however they share the same common attributes and entry points that the other filter must support.
- **Filter Plugins** - Filter plugins are used to modify data as it flows through FogLAMP. Filter plugins may be combined into a set of ordered filters that are applied as a pipeline to either the south ingress service or the north egress task that sends data to external systems.
- **Notification Rule Plugins** - These are used by the optional notification service in order to evaluate data that flows into the notification service to determine if a notification should be sent.

- **Notification Delivery Plugins** - These plugins are used by the optional notification service to deliver a notification to a system when a notification rule has triggered. These plugins allow the mechanisms to deliver notifications to be extended.

13.1.1 Plugins in this version of FogLAMP

This version of FogLAMP provides the following plugins in the main repository:

Type	Name	Initial Status	Description	Availability	Notes
Storage	SQLite	Enabled	SQLite storage for data and metadata	Ubuntu: x86_64 Ubuntu Core: x86, ARM Raspbian	
Storage	Postgres	Disabled	PostgreSQL storage for data and metadata	Ubuntu: x86_64 Ubuntu Core: x86, ARM Raspbian	
North	OMF	Disabled	OSIssoft Message Format sender to PI Connector Relay OMF	Ubuntu: x86_64 Ubuntu Core: x86, ARM Raspbian	It works with PI Connector Relay OMF 1.2.X and 2.2. The plugin also works against EDS and OCS.

In addition to the plugins in the main repository, there are many other plugins available in separate repositories, a list of the is maintained within this document.

13.1.2 Installing New Plugins

As a general rule and unless the documentation states otherwise, plugins should be installed in two ways:

- When the plugin is available as **package**, it should be installed when **FogLAMP is running**. This is the required method because the package executed pre and post-installation tasks that require FogLAMP to run.
- When the plugin is available as **source code**, it should be installed when **FogLAMP is either running or not**. You will want to manually move the plugin code into the right location where FogLAMP is installed, add prerequisites and execute the REST commands necessary to start the plugin **after** you have started FogLAMP if it is not running when you start this process.

For example, this is the command to use to install the *OpenWeather* South plugin:

```
$ sudo systemctl status foglamp.service
foglamp.service - LSB: FogLAMP
   Loaded: loaded (/etc/init.d/foglamp; bad; vendor preset: enabled)
   Active: active (running) since Wed 2018-05-16 01:32:25 BST; 4min 1s ago
     Docs: man:systemd-sysv-generator(8)
    CGroup: /system.slice/foglamp.service
            └─13741 python3 -m foglamp.services.core
               └─13746 /usr/local/foglamp/services/storage --address=0.0.0.0 --port=40138

May 16 01:36:09 ubuntu python3[13741]: FogLAMP[13741] INFO: scheduler: foglamp.
↪services.core.scheduler.scheduler: Process started: Schedule 'stats collection'
↪process 'stats coll
                                     ['tasks/statistics', '--port=40138', '--
↪address=127.0.0.1', '--name=stats collector']
...
```

(continues on next page)

(continued from previous page)

```

FogLAMP v1.3.1 running.
FogLAMP Uptime: 266 seconds.
FogLAMP records: 0 read, 0 sent, 0 purged.
FogLAMP does not require authentication.
=== FogLAMP services:
foglamp.services.core
=== FogLAMP tasks:
$
$ sudo cp foglamp-south-openweathermap-1.2-x86_64.deb /var/cache/apt/archives/.
$ sudo apt install /var/cache/apt/archives/foglamp-south-openweathermap-1.2-x86_64.deb
Reading package lists... Done
Building dependency tree
Reading state information... Done
Note, selecting 'foglamp-south-openweathermap' instead of '/var/cache/apt/archives/
↳foglamp-south-openweathermap-1.2-x86_64.deb'
The following packages were automatically installed and are no longer required:
  linux-headers-4.4.0-109 linux-headers-4.4.0-109-generic linux-headers-4.4.0-119_
↳linux-headers-4.4.0-119-generic linux-headers-4.4.0-121 linux-headers-4.4.0-121-
↳generic
  linux-image-4.4.0-109-generic linux-image-4.4.0-119-generic linux-image-4.4.0-121-
↳generic linux-image-extra-4.4.0-109-generic linux-image-extra-4.4.0-119-generic
  linux-image-extra-4.4.0-121-generic
Use 'sudo apt autoremove' to remove them.
The following NEW packages will be installed
  foglamp-south-openweathermap
0 to upgrade, 1 to newly install, 0 to remove and 0 not to upgrade.
Need to get 0 B/3,404 B of archives.
After this operation, 0 B of additional disk space will be used.
Selecting previously unselected package foglamp-south-openweathermap.
(Reading database ... 211747 files and directories currently installed.)
Preparing to unpack .../foglamp-south-openweathermap-1.2-x86_64.deb ...
Unpacking foglamp-south-openweathermap (1.2) ...
Setting up foglamp-south-openweathermap (1.2) ...
openweathermap plugin installed.
$
$ foglamp status
FogLAMP v1.3.1 running.
FogLAMP Uptime: 271 seconds.
FogLAMP records: 36 read, 0 sent, 0 purged.
FogLAMP does not require authentication.
=== FogLAMP services:
foglamp.services.core
foglamp.services.south --port=42066 --address=127.0.0.1 --name=openweathermap
=== FogLAMP tasks:
$

```

You may also install new plugins directly from within the FogLAMP GUI, however you will need to have setup your Linux machine to include the FogLAMP package repository in the list of repositories the Linux package manager searches for new packages.

13.2 Representing Data

The key purpose of FogLAMP and the plugins is the manipulation of data, that data is passed around the system and represented in a number of ways. This section will introduce the data representation formats used at various locations within the FogLAMP system. Conceptually the unit of data that we use is a reading. The reading represents the state of a monitored device at a point in time and has a number of elements.

Name	Description
asset	The name of the asset or device to which the data refers
timestamp	The point in time at which these values were observed.
data points	A set of named values for the data held for the asset

There are actually two timestamps within a reading and these may be different. There is a *user_ts*, which is the time the plugin assigned to the reading data and may come from the device itself and the *ts*. The *ts* timestamp is set by the system when the data is read into FogLAMP. Unless the plugin is able to determine a timestamp from the device the *user_ts* is usually the same as the *ts*.

The data points themselves are a set of name and value pairs, with the values supporting a number of different data types. These will be described below.

Reading data is nominally stored and passed between the APIs using JSON, however for convenience it is access in different ways within the different languages that can be used to implement FogLAMP components and plugins. In JSON a reading is represented as a JSON DICT whereas in C++ a Reading is a class, as is a data point. The way the different data point types are represented is outline below.

Type	JSON	C++	Python
Integer	An integer	An int	An integer
Floating Point	A floating point value	A double	A floating point
Boolean	A string either "true" or "false"	A bool	A boolean
String	A string	A std::string pointer	A string
List of numbers	An array of floating point values	A std::vector<double>	A list of floating point values
2 Dimensional list of numbers	A list of lists of floating point values	A std::vector of std::vector<double> pointers	A list of lists of floating point values
Data buffer	A base64 encoded string with a header	A Databuffer class	A 1 dimensional numpy array of values
Image	A base64 encoded string with a header	A DPLImage class	A 2 dimensional numpy array of pixels. In the case of RGB images each pixels is an array

13.3 Writing and Using Plugins

A plugin has a small set of external entry points that must exist in order for FogLAMP to load and execute that plugin. Currently plugins may be written in either Python or C/C++, the set of entry points is the same for both languages. The entry points detailed here will be presented for both languages, a more in depth discussion of writing plugins in C/C++ will then follow.

13.3.1 Common FogLAMP Plugin API

Every plugin provides at least one common API entry point, the *plugin_info* entry point. It is used to obtain information about a plugin before it is initialized and used. It allows FogLAMP to determine what type of plugin it is, e.g. a South bound plugin or a North bound plugin, obtain default configuration information for the plugin and determine version information.

Plugin Information

The information entry point is implemented as a call, *plugin_info*, that takes no arguments. Data is returned from this API call as a JSON document with certain well known properties.

A typical Python implementation of this would simply return a fixed dictionary object that encodes the required properties.

```
def plugin_info():
    """ Returns information about the plugin.

    Args:
    Returns:
        dict: plugin information
    Raises:
    """

    return {
        'name': 'DHT11 GPIO',
        'version': '1.0',
        'mode': 'poll',
        'type': 'south',
        'interface': '1.0',
        'config': _DEFAULT_CONFIG
    }
```

These are the properties returned by the JSON document:

- **name** - A textual name that will be used for reporting purposes for this plugin.
- **version** - This property allows the version of the plugin to be communicated to the plugin loader. This is used for reporting purposes only and has no effect on the way FogLAMP interacts with the plugin.
- **mode** - A set of options that defines how the plugin operates. Multiple values can be given, the different options are separated from each other using the | symbol.
- **type** - The type of the plugin, used by the plugin loader to determine if the plugin is being used correctly. The type is a simple string and may be *south*, *north*, *filter*, *rule* or *delivery*.

Note: If you browse the FogLAMP code you may find old plugins with type *device*: this was the type used to indicate a South plugin and it is now deprecated.

- **interface** - This property reports the version of the plugin API to which this plugin was written. It allows FogLAMP to support upgrades of the API whilst being able to recognise the version that a particular plugin is compliant with. Currently all interfaces are version 1.0.
- **configuration** - This allows the plugin to return a JSON document which contains the default configuration of the plugin. This is in line with the extensible plugin mechanism of FogLAMP, each plugin will return a set of configuration items that it wishes to use, this will then be used to extend the set of FogLAMP configuration items. This structure, a JSON document, includes default values but no actual values for each configuration

option. The first time FogLAMP's configuration manager sees a category it will register the category and create values for each item using the default value in the configuration document. On subsequent calls the value already in the configuration manager will be used. This mechanism allows the plugin to extend the set of configuration variables whilst giving the user the opportunity to modify the value of these configuration items. It also allow new versions of plugins to add new configuration items whilst retaining the values of previous items. And new items will automatically be assigned the default value for that item. As an example, a plugin that wishes to maintain two configuration variables, say a GPIO pin to use and a polling interval, would return a configuration document that looks as follows:

```
{
  'pollInterval': {
    'description': 'The interval between poll calls to the device poll routine,
    ↪expressed in milliseconds.',
    'type': 'integer',
    'default': '1000'
  },
  'gpiopin': {
    'description': 'The GPIO pin into which the DHT11 data pin is connected',
    'type': 'integer',
    'default': '4'
  }
}
```

The various values that may appear in the *mode* item are shown in the table below

Mode	Description
poll	The plugin is a polled plugin and <i>plugin_poll</i> will be called periodically to obtain new values.
async	The plugin is an asynchronous plugin, <i>plugin_poll</i> will not be called and the plugin will be supplied with a callback function that it calls each time it has a new value to pass to the system. The <i>plugin_register_ingest</i> entry point will be called to register the callback with the plugin. The <i>plugin_start</i> call will be called once to initiate the asynchronous delivery of data.
none	This is equivalent to poll.
control	The plugin support a control flow to the device the plugin is connected to. The must supply the control entry points <i>plugin_write</i> and <i>plugin_operation</i> .

A C/C++ plugin returns the same information as a structure, this structure includes the JSON configuration document as a simple C string.

```
#include <plugin_api.h>

extern "C" {

/**
 * The plugin information structure
 */
static PLUGIN_INFORMATION info = {
    "MyPlugin",           // Name
    "1.0.1",              // Version
    0,                    // Flags
    PLUGIN_TYPE_SOUTH,    // Type
    "1.0.0",              // Interface version
    default_config         // Default configuration
};
```

(continues on next page)

(continued from previous page)

```

/**
 * Return the information about this plugin
 */
PLUGIN_INFORMATION *plugin_info()
{
    return &info;
}

```

In the above example the constant *default_config* is a string that contains the JSON configuration document. In order to make the JSON easier to manage a special macro is defined in the *plugin_api.h* header file. This macro is called *QUOTE* and is designed to ease the quoting requirements to create this JSON document.

```

const char *default_config = QUOTE({
    "plugin" : {
        "description" : "My example plugin in C++",
        "type" : "string",
        "default" : "MyPlugin",
        "readonly" : "true"
    },
    "asset" : {
        "description" : "The name of the asset the plugin will produce",
        "type" : "string",
        "default" : "MyAsset"
    }
});

```

The *flags* items contains a bitmask of flag values used to pass information regarding the behavior and requirements of the plugin. The flag values currently supported are shown below

Flag Name	Description
SP_COMMON	Used exclusively by storage plugins. The plugin supports the common table access needed to store configuration
SP_READINGS	Used exclusively by storage plugins. The plugin supports the storage of reading data
SP_ASYNC	The plugin is an asynchronous plugin, <i>plugin_poll</i> will not be called and the plugin will be supplied with a callback function that it calls each time it has a new value to pass to the system. The <i>plugin_register_ingest</i> entry point will be called to register the callback with the plugin. The <i>plugin_start</i> call will be called once to initiate the asynchronous delivery of data. This applies only to south plugins.
SP_PERSISTENT_DATA	The plugin wishes to persist data between executions
SP_INGEST	Non-south plugin wishes to ingest new data into the system. Used by notification plugins
SP_GET_MANAGEMENT	The plugin requires access to the management API interface for the service
SP_GET_STORAGE	The plugin requires access to the storage service
SP_DEPRECATED	The plugin should be considered to be deprecated. New service can not use this plugin, but existing services may continue to use it
SP_BUILT_IN	The plugin is not implemented as an external package but is built into the system
SP_CONTROL	The plugin implement control features

These flag values may be combined by use of the or operator where more than one of the above options is supported.

Plugin Initialization

The plugin initialization is called after the service that has loaded the plugin has collected the plugin information and resolved the configuration of the plugin but before any other calls will be made to the plugin. The initialization routine is called with the resolved configuration of the plugin, this includes values as opposed to the defaults that were returned in the *plugin_info* call.

This call is used by the plugin to do any initialization or state creation it needs to do. The call returns a handle which will be passed into each subsequent call of the plugin. The handle allows the plugin to have state information that is maintained and passed to it whilst allowing for multiple instances of the same plugin to be loaded by a service if desired. It is equivalent to a this or self pointer for the plugin, although the plugin is not defined as a class.

In Python a simple example of a sensor that reads a GPIO pin for data, we might choose to use that configured GPIO pin as the handle we pass to other calls.

```
def plugin_init(config):
    """ Initialise the plugin.

    Args:
        config: JSON configuration document for the device configuration category
    Returns:
        handle: JSON object to be used in future calls to the plugin
    Raises:
        """

    handle = config['gpiopin']['value']
    return handle
```

A C/C++ plugin should return a value in a *void* pointer that can then be dereferenced in subsequent calls. A typical C++ implementation might create an instance of a class and use that instance as the handle for the plugin.

```
/**
 * Initialise the plugin, called to get the plugin handle
 */
PLUGIN_HANDLE plugin_init(ConfigCategory *config)
{
    MyPluginClass *plugin = new MyPluginClass();

    plugin->configure(config);

    return (PLUGIN_HANDLE)plugin;
}
```

It should also be observed in the above C/C++ example the *plugin_init* call is passed a pointer to a *ConfigCategory* class that encapsulates the JSON configuration category for the plugin. Details of the *ConfigCategory* class are available in the section .

Plugin Shutdown

The plugin shutdown method is called as part of the shutdown sequence of the service that loaded the plugin. It gives the plugin the opportunity to do any cleanup operations before terminating. As with all calls it is passed the handle of our plugin instance. Plugins can not prevent the shutdown and do not have to implement any actions. In our simple sensor example there is nothing to do in order to shutdown the plugin.

A C/C++ plugin might use this *plugin_shutdown* call to delete the plugin class instance it created in the corresponding *plugin_init* call.

```
/**
 * Shutdown the plugin
 */
void plugin_shutdown(PLUGIN_HANDLE *handle)
{
    MyPluginClass *plugin = (MyPluginClass *)handle;

    delete plugin;
}
```

Plugin Reconfigure

The plugin reconfigure method is called whenever the configuration of the plugin is changed. It allows for the dynamic reconfiguration of the plugin whilst it is running. The method is called with the handle of the plugin and the updated configuration document. The plugin should take whatever action it needs to and return a new or updated copy of the handle that will be passed to future calls.

The plugin reconfigure method is shared between most but not all plugin types. In particular it does not exist for the shorted lived plugins that are created to perform a single operation and then terminated. These are the north plugins and the notification delivery plugins.

Using a simple Python example of our sensor reading a GPIO pin, we extract the new pin number from the new configuration data and return that as the new handle for the plugin instance.

```
def plugin_reconfigure(handle, new_config):
    """ Reconfigures the plugin, it should be called when the configuration of the
    ↪ plugin is changed during the
        operation of the device service.
        The new configuration category should be passed.

    Args:
        handle: handle returned by the plugin initialisation call
        new_config: JSON object representing the new configuration category for the
    ↪ category
    Returns:
        new_handle: new handle to be used in the future calls
    Raises:
        """

    new_handle = new_config['gpiopin']['value']
    return new_handle
```

In C/C++ the *plugin_reconfigure* method is very similar, note however that the *plugin_reconfigure* call is passed the JSON configuration category as a string and not a *ConfigCategory*, it is easy to parse and create the C++ class however, a name for the category must be given however.

```
/**
 * Reconfigure the plugin
 */
void plugin_reconfigure(PLUGIN_HANDLE *handle, string& newConfig)
{
    ConfigCategory      config("newConfiguration", newConfig);
    MyPluginClass       *plugin = (MyPluginClass *)*handle;

    plugin->configure(&config);
}
```

It should be noted that the *plugin_reconfigure* call may be delivered in a separate thread for a C/C++ plugin and that the plugin should implement any mutual exclusion mechanisms that are required based on the actions of the *plugin_reconfigure* method.

13.3.2 Configuration Lifecycle

FogLAMP has a very particular way of handling configuration, there are a number of design aims that have resulted in the configuration system within FogLAMP.

- A desire to allow the plugins to define their own configuration elements.
- Dynamic configuration that allows for maximum uptime during configuration changes.
- A descriptive way to define the configuration such that user interfaces can be built without prior knowledge of the elements to be configured.
- A common approach that will work across many different languages.

FogLAMP divides its configuration in categories. A category being a collection of configuration items. A category is also the smallest item of configuration that can be subscribed to by the code. This subscription mechanism is the way that FogLAMP facilitates dynamic reconfiguration. It allows a service to subscribe to one or more configuration categories, whenever an item within a category changes the central configuration manager will call a handler to pass the newly updated configuration category. This handler may be within a service or between services using the micro service management API that every service must support. The mechanism however is transparent to the code involved.

The configuration items within a category are JSON object, the object key is the name of the configuration item, the object itself contains data about that item. As an example, if we wanted to have a configuration item called *MaxRetries* that is an integer with a default value of 5, then we would configure it using the JSON object

```
"MaxRetries" : {
    "type" : "integer",
    "default" : "5"
}
```

We have used the properties *type* and *default* to define properties of the configuration item *MaxRetries*. These are not the only properties that a configuration item can have, the full set of properties are

Property	Description
default	The default value for the configuration item. This is always expressed as a string regardless of the type of the configuration item.
deprecated	A boolean flag to indicate that this item is no longer used and will be removed in a future release.
description	A description of the configuration item used in the user interface to give more details of the item. Commonly used as a mouse over help prompt.
display-name	The string to use in the user interface when presenting the configuration item. Generally a more user friendly form of the item name. Item names are referenced within the code.
length	The maximum length of the string value of the item.
mandatory	A boolean flag to indicate that this item can not be left blank.
maximum	The maximum value for a numeric configuration item.
minimum	The minimum value for a numeric configuration item.
options	Only used for enumeration type elements. This is a JSON array of string that contains the options in the enumeration.
order	Used in the user interface to give an indication of how high up in the dialogue to place this item.
read-only	A boolean property that can be used to include items that can not be altered by the API.
rule	A validation rule that will be run against the value. This must evaluate to true for the new value to be accepted by the API
type	The type of the configuration item. The list of types supported are; integer, float, string, password, enumeration, boolean, JSON, URL, IPV4, IPV6, script, code, X509 certificate and northTask.
validity	An expression used to determine if the configuration item is valid. Used in the UI to gray out one value based on the value of others.
value	The current value of the configuration item. This is not included when defining a set of default configuration in, for example, a plugin.

Of the above properties of a configuration item *type*, *default* and *description* are mandatory, all other may be omitted.

Configuration data is stored by the storage service and is maintained by the configuration in the core FogLAMP service. When code requires configuration it would create a configuration category with a set of items as a JSON document. It would then register that configuration category with the configuration manager. The configuration manager is responsible for storing the data in the storage layer, as it does this it first checks to see if there is already a configuration category from a previous execution of the code. If one does exist then the two are merged, this merging process allows updates to the software to extend the configuration category whilst maintaining any changes in values made by the user.

Dynamic reconfiguration within FogLAMP code is supported by allowing code to subscribe for changes in a configuration category. The services that load plugin will automatically register for the plugin configuration category and when changes are seen will call the *plugin_reconfigure* entry point of the plugin with the new configuration. This allows the plugins to receive the updated configuration and take what actions it must in order to honour the changes to configuration. This allows for configuration to be changed without the need to stop and restart the services, however some plugins may need to close connections and reopen them, which may cause a slight interruption in the process of gathering data. That choice is up to the developers of the individual plugins.

Discovery

It is possible using this system to do a limited amount of discovery and tailoring of plugin configuration. A typical case when discovery might be used is to discover devices on a network that can be monitored. This can be achieved by putting the discovery code in the *plugin_info* entry point and having that discovery code alter the default configuration that is returned as part of the plugin information structure.

Any example of this might be to have an enumeration in the configuration that enumerates the devices to be monitored. The discovery code would then populate the enumerations options item with the various devices it discovered when the *plugin_info* call was made.

An example of the *plugin_info* entry point that does this might be as follows

```
/**
 * Return the information about this plugin
 */
PLUGIN_INFORMATION *plugin_info()
{
    DeviceDiscovery discover;

    char *config = discover.discover(default_config, "discovered");
    info.config = config;
    return &info;
}
```

The configuration in *default_config* is assumed to have an enumeration item called *discovered*

```
"discovered" : {
    "description" : "The discovered devices, select 'Manual' to manually enter an
↪IP address",
    "type" : "enumeration",
    "options" : [ "Manual" ],
    "default" : "Manual",
    "displayName": "Devices",
    "mandatory": "true",
    "order" : "2"
},
"IP" : {
    "description" : "The IP address of your device, used to add a device that
↪could not be discovered",
    "type" : "string",
    "default" : "127.0.0.1",
    "displayName": "IP Address",
    "mandatory": "true",
    "order" : "3",
    "validity" : "discovered == \"Manual\""
},
```

Note the use of the *Manual* option to allow entry of devices that could not be discovered.

The *discover* method does the actual discovery and manipulates the JSON configuration to add the *options* element of the configuration item.

The code that connects to the device should then look at the *discovered* configuration item, if it finds it set to *Manual* then it will get an IP address from the *IP* configuration item. Otherwise it uses the information in the *discovered* item to connect, note that this need not just be an IP address, you can format the data in a way that is more user friendly and have the connection code extract what it needs or create a table in the *discover* method to allow for user meaningful strings to be mapped to network addresses.

The example here was written in C++, there is nothing that is specific to C++ however and the same approach can be taken in Python.

One thing to note however, the *plugin_info* call is used in the display of available plugins, discovery code that is very slow will impact the performance of plugin selection.

13.4 South Plugins

South plugins are used to communicate with sensors and actuators, there are two modes of plugin operation; *asyncio* and *polled*.

13.4.1 Polled Mode

Polled mode is the simplest form of South plugin that can be written, a poll routine is called at an interval defined in the plugin configuration. The South service determines the type of the plugin by examining at the mode property in the information the plugin returns from the *plugin_info* call.

Plugin Poll

The plugin *poll* method is called periodically to collect the readings from a poll mode sensor. As with all other calls the argument passed to the method is the handle returned by the initialization call, the return of the method should be the JSON payload of the readings to return.

The JSON payload returned, as a Python dictionary, should contain the properties; asset, timestamp, key and readings.

Property	Description
asset	The asset key of the sensor device that is being read
timestamp	A timestamp for the reading data
key	A UUID which is the unique key of this reading
readings	The reading data itself as a JSON object

It is important that the *poll* method does not block as this will prevent the proper operation of the South microservice. Using the example of our simple DHT11 device attached to a GPIO pin, the *poll* routine could be:

```
def plugin_poll(handle):
    """ Extracts data from the sensor and returns it in a JSON document as a Python_
    ↪dict.

    Available for poll mode only.

    Args:
        handle: handle returned by the plugin initialisation call
    Returns:
        returns a sensor reading in a JSON document, as a Python dict, if it is_
    ↪available
        None - If no reading is available
    Raises:
        DataRetrievalError
    """

    try:
        humidity, temperature = Adafruit_DHT.read_retry(Adafruit_DHT.DHT11, handle)
```

(continues on next page)

(continued from previous page)

```

    if humidity is not None and temperature is not None:
        time_stamp = str(datetime.now(tz=timezone.utc))
        readings = { 'temperature': temperature , 'humidity' : humidity }
        wrapper = {
            'asset': 'dht11',
            'timestamp': time_stamp,
            'key': str(uuid.uuid4()),
            'readings': readings
        }
        return wrapper
    else:
        return None

except Exception as ex:
    raise exceptions.DataRetrievalError(ex)

return None

```

13.4.2 Async IO Mode

In asyncio mode the plugin inserts itself into the event processing loop of the South Service itself. This is a more complex mechanism and is intended for plugins that need to block or listen for incoming data via a network.

Plugin Start

The *plugin_start* method, as with other plugin calls, is called with the plugin handle data that was returned from the *plugin_init* call. The *plugin_start* call will only be called once for a plugin, it is the responsibility of *plugin_start* to install the plugin code into the python event handling system for asyncio. Assuming an example whereby the interface to a sensor is via HTTP and the sensor will make HTTP POST calls to our plugin in order to send data into FogLAMP, a *plugin_start* for this scenario would create a web application endpoint for reception of the POST command.

```

loop = asyncio.get_event_loop()
app = web.Application(middlewares=[middleware.error_middleware])
app.router.add_route('POST', '/', SensorPhoneIngest.render_post)
handler = app.make_handler()
coro = loop.create_server(handler, host, port)
server = asyncio.ensure_future(coro)

```

This code first gets the event loop for this Python execution, it then creates the web application and adds a route for the POST request. In this case it is calling the *render_post* method of the object *SensorPhone*. It then goes on to create the handler and install the web server instance into the event system.

Async Data Callback

The async data callback is used for incoming sensor data and passing that reading data into the FogLAMP ingest process. Unlike the poll mechanism, this is done from within the callback rather than by passing the data back to the South service itself. A plugin entry point, *plugin_register_ingest* is called by the south service before the plugin is started to register the callback with the plugin. The plugin would usually save the callback function and the reference data for later use.

```
def plugin_register_ingest(handle, callback, ingest_ref):
    """Required plugin interface component to communicate to South C server

    Args:
        handle: handle returned by the plugin initialisation call
        callback: C opaque object required to passed back to C->ingest method
        ingest_ref: C opaque object required to passed back to C->ingest method
    """
    global c_callback, c_ingest_ref
    c_callback = callback
    c_ingest_ref = ingest_ref
```

The plugin then uses these saved references when it has data to be ingested. A new reading is constructed and passed to the callback function using *async_ingest* object that should be imported by the plugin.

```
import async_ingest
```

Then for each reading to be ingested the data is sent to the ingest thread of the south plugin using the following construct.

```
data = {
    'asset': self.asset_name,
    'timestamp': utils.local_timestamp(),
    'readings': reads
}
async_ingest.ingest_callback(c_callback, c_ingest_ref, data)
```

```
message['status'] = code
return web.json_response(message)
```

13.4.3 Set Point Control

South plugins can also be used to exert control on the underlying device to which they are connected. This is not intended for use as a substitute for real time control systems, but rather as a mechanism to make non-time critical changes to a device or to trigger an operation on the device.

To make a south plugin support control features there are two steps that need to be taken

- Tag the plugin as supporting control
- Add the entry points for control

Enable Control

A plugin enables control features by means of the mode field in the plugin information dict which is returned by the *plugin_info* entry point of the plugin. The flag value *control* should be added to the mode field of the plugin. Multiple flag values are separated by the pipe symbol '|'.

```
# plugin information dict
{
    'name': 'Sinusoid Poll plugin',
    'version': '1.9.2',
    'mode': 'poll|control',
    'type': 'south',
    'interface': '1.0',
```

(continues on next page)

(continued from previous page)

```
'config': _DEFAULT_CONFIG
}
```

Adding this flag will cause the south service to do a number of things when it loads the plugin;

- The south service will attempt to resolve the two control entry points.
- A toggle will be added to the advanced configuration category of the service that will permit the disabling of control services.
- A security category will be added to the south service that contains the access control lists and permissions associated with the service.

Control Entry Points

Two entry points are supported for control operations in the south plugin

- **plugin_write**: which is used to set the value of a parameter within the plugin or device
- **plugin_operation**: which is used to perform an operation on the plugin or device

The south plugin can support one or both of these entry points as appropriate for the plugin.

Write Entry Point

The write entry point is used to set data in the plugin or write data into the device.

The plugin write entry point is defined as follows

```
def plugin_write(handle, name, value)
```

Where the parameters are;

- **handle** the handle of the plugin instance
- **name** the name of the item to be changed
- **value** a string presentation of the new value to assign to the item

The return value defines if the write was successful or not. True is returned for a successful write.

```
def plugin_write(handle, name, value):
    """ Setpoint write operation

    Args:
        handle: handle returned by the plugin initialisation call
        name: Name of parameter to write
        value: Value to be written to that parameter
    Returns:
        bool: Result of the write operation
    """
    _LOGGER.info("plugin_write(): name={}, value={}".format(name, value))
    return True
```

In this case we are merely printing the parameter name and the value to be set for this parameter. Normally control would be used for making a change with the connected device itself, such as changing a PLC register value. This is simply an example to demonstrate the API.

Operation Entry Point

The plugin will support an operation entry point. This will execute the given operation synchronously, it is expected that this operation entry point will be called using a separate thread, therefore the plugin should implement operations in a thread safe environment.

The plugin write operation entry point is defined as follows

```
def plugin_operation(handle, operation, params)
```

Where the parameters are;

- **handle** the handle of the plugin instance
- **operation** the name of the operation to be executed
- **params** a list of name/value tuples that are passed to the operation

The *operation* parameter should be used by the plugin to determine which operation is to be performed. The actual parameters are passed in a list of key/value tuples as strings.

The return from the call is a boolean result of the operation, a failure of the operation or a call to an unrecognized operation should be indicated by returning a false value. If the operation succeeds a value of true should be returned.

The following example shows the implementation of the plugin operation entry point.

```
def plugin_operation(handle, operation, params):
    """ Setpoint control operation

    Args:
        handle: handle returned by the plugin initialisation call
        operation: Name of operation
        params: Parameter list
    Returns:
        bool: Result of the operation
    """
    _LOGGER.info("plugin_operation(): operation={}, params={}".format(operation,
↪params))
    return True
```

In the case of a real machine the operation would most likely cause an action on a machine, for example a request to the machine to re-calibrate itself. Above example is just a demonstration of the API.

13.4.4 A South Plugin Example In Python: the DHT11 Sensor

Let's try to put all the information together and write a plugin. We can continue to use the example of an inexpensive sensor, the DHT11, used to measure temperature and humidity, directly wired to a Raspberry PI. This plugin is available on github, .

First, here is a set of links where you can find more information regarding this sensor:

-
-
-

The Hardware

The DHT sensor is directly connected to a Raspberry PI 2 or 3. You may decide to buy a sensor and a resistor and solder them yourself, or you can buy a ready-made circuit that provides the correct output to wire to the Raspberry PI. shows a DHT11 with resistor that you can buy online.

The sensor can be directly connected to the Raspberry PI GPIO (General Purpose Input/Output). An introduction to the GPIO and the pinset is available . In our case, you must connect the sensor on these pins:

- **VCC** is connected to PIN #2 (5v Power)
- **GND** is connected to PIN #6 (Ground)
- **DATA** is connected to PIN #7 (BCM 4 - GPCLK0)

shows the sensor wired to the Raspberry PI and is a zoom into the wires used.

The Software

For this plugin we use the ADAFruit Python Library (links to the GitHub repository are above). First, you must install the library (in future versions the library will be provided in a ready-made package):

```
$ git clone https://github.com/adafruit/Adafruit_Python_DHT.git
Cloning into 'Adafruit_Python_DHT'...
remote: Counting objects: 249, done.
remote: Total 249 (delta 0), reused 0 (delta 0), pack-reused 249
Receiving objects: 100% (249/249), 77.00 KiB | 0 bytes/s, done.
Resolving deltas: 100% (142/142), done.
$ cd Adafruit_Python_DHT
$ sudo apt-get install build-essential python-dev
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following NEW packages will be installed:
build-essential python-dev
...
$ sudo python3 setup.py install
running install
running bdist_egg
running egg_info
creating Adafruit_DHT.egg-info
...
$
```

The Plugin

This is the code for the plugin:

```
# -*- coding: utf-8 -*-

# FOGLAMP_BEGIN
# See: http://foglamp.readthedocs.io/
# FOGLAMP_END

""" Plugin for a DHT11 temperature and humidity sensor attached directly
    to the GPIO pins of a Raspberry Pi
```

(continues on next page)

(continued from previous page)

This plugin uses the Adafruit DHT library, to install this perform the following steps:

```
git clone https://github.com/adafruit/Adafruit_Python_DHT.git
cd Adafruit_Python_DHT
sudo apt-get install build-essential python-dev
sudo python setup.py install
```

To access the GPIO pins foglamp must be able to access /dev/gpiomem, the default access for this is owner and group read/write. Either FogLAMP must be added to the group or the permissions altered to allow FogLAMP access to the device.

"""

```
from datetime import datetime, timezone
import uuid
```

```
from foglamp.common import logger
from foglamp.services.south import exceptions
```

```
__author__ = "Mark Riddoch"
__copyright__ = "Copyright (c) 2017 OSIsoft, LLC"
__license__ = "Apache 2.0"
__version__ = "${VERSION}"
```

```
_DEFAULT_CONFIG = {
    'plugin': {
        'description': 'Python module name of the plugin to load',
        'type': 'string',
        'default': 'dht11'
    },
    'pollInterval': {
        'description': 'The interval between poll calls to the device poll routine_
↳expressed in milliseconds.',
        'type': 'integer',
        'default': '1000'
    },
    'gpiopin': {
        'description': 'The GPIO pin into which the DHT11 data pin is connected',
        'type': 'integer',
        'default': '4'
    }
}
```

```
_LOGGER = logger.setup(__name__)
""" Setup the access to the logging system of FogLAMP """
```

```
def plugin_info():
    """ Returns information about the plugin.

    Args:
    Returns:
        dict: plugin information
    Raises:
```

(continues on next page)

(continued from previous page)

```

"""

return {
    'name': 'DHT11 GPIO',
    'version': '1.0',
    'mode': 'poll',
    'type': 'south',
    'interface': '1.0',
    'config': _DEFAULT_CONFIG
}

def plugin_init(config):
    """ Initialise the plugin.

    Args:
        config: JSON configuration document for the device configuration category
    Returns:
        handle: JSON object to be used in future calls to the plugin
    Raises:
        """

    handle = config['gpiopin']['value']
    return handle

def plugin_poll(handle):
    """ Extracts data from the sensor and returns it in a JSON document as a Python_
    ↪dict.

    Available for poll mode only.

    Args:
        handle: handle returned by the plugin initialisation call
    Returns:
        returns a sensor reading in a JSON document, as a Python dict, if it is_
    ↪available
        None - If no reading is available
    Raises:
        DataRetrievalError
        """

    try:
        humidity, temperature = Adafruit_DHT.read_retry(Adafruit_DHT.DHT11, handle)
        if humidity is not None and temperature is not None:
            time_stamp = str(datetime.now(tz=timezone.utc))
            readings = {'temperature': temperature, 'humidity': humidity}
            wrapper = {
                'asset': 'dht11',
                'timestamp': time_stamp,
                'key': str(uuid.uuid4()),
                'readings': readings
            }
            return wrapper
        else:
            return None

```

(continues on next page)

(continued from previous page)

```

except Exception as ex:
    raise exceptions.DataRetrievalError(ex)

return None

def plugin_reconfigure(handle, new_config):
    """ Reconfigures the plugin, it should be called when the configuration of the
    ↪ plugin is changed during the
        operation of the device service.
        The new configuration category should be passed.

    Args:
        handle: handle returned by the plugin initialisation call
        new_config: JSON object representing the new configuration category for the
    ↪ category
    Returns:
        new_handle: new handle to be used in the future calls
    Raises:
        """

    new_handle = new_config['gpiopin']['value']
    return new_handle

def plugin_shutdown(handle):
    """ Shutdowns the plugin doing required cleanup, to be called prior to the device
    ↪ service being shut down.

    Args:
        handle: handle returned by the plugin initialisation call
    Returns:
    Raises:
        """
    pass

```

Building FogLAMP and Adding the Plugin

If you have not built FogLAMP yet, follow the steps described . After the build, you can optionally install FogLAMP following steps.

- If you have started FogLAMP from the build directory, copy the structure of the *foglamp-south-dht11/python/* directory into the *python* directory:

```

$ cd ~/FogLAMP
$ cp -R ~/foglamp-south-dht11/python/foglamp/plugins/south/dht11 python/foglamp/
↪ plugins/south/
$

```

- If you have installed FogLAMP by executing `sudo make install`, copy the structure of the *foglamp-south-dht11/python/* directory into the installed *python* directory:

```

$ sudo cp -R ~/foglamp-south-dht11/python/foglamp/plugins/south/dht11 /usr/local/
↪ foglamp/python/foglamp/plugins/south/
$

```

Note: If you have installed FogLAMP using an alternative *DESTDIR*, remember to add the path to the destination directory to the `cp` command.

- Add service

```
$ curl -sX POST http://localhost:8081/foglamp/service -d '{"name": "dht11", "type":  
→"south", "plugin": "dht11", "enabled": true}'
```

Note: Each plugin repo has its own debian packaging script and documentation, And that is the recommended way to go! As above method(s) may need explicit action for linux and/or python dependencies installation.

Using the Plugin

Once south plugin is added as an enabled service, You are ready to use the DHT11 plugin.

```
$ curl -X GET http://localhost:8081/foglamp/service | jq
```

Let's see what we have collected so far:

```
$ curl -s http://localhost:8081/foglamp/asset | jq  
[  
  {  
    "count": 158,  
    "asset_code": "dht11"  
  }  
]  
$
```

Finally, let's extract some values:

```
$ curl -s http://localhost:8081/foglamp/asset/dht11?limit=5 | jq  
[  
  {  
    "timestamp": "2017-12-30 14:41:39.672",  
    "reading": {  
      "temperature": 19,  
      "humidity": 62  
    }  
  },  
  {  
    "timestamp": "2017-12-30 14:41:35.615",  
    "reading": {  
      "temperature": 19,  
      "humidity": 63  
    }  
  },  
  {  
    "timestamp": "2017-12-30 14:41:34.087",  
    "reading": {  
      "temperature": 19,  
      "humidity": 62  
    }  
  },  
]
```

(continues on next page)

(continued from previous page)

```
{
  "timestamp": "2017-12-30 14:41:32.557",
  "reading": {
    "temperature": 19,
    "humidity": 63
  }
},
{
  "timestamp": "2017-12-30 14:41:31.028",
  "reading": {
    "temperature": 19,
    "humidity": 63
  }
}
]
$
```

Clearly we will not see many changes in temperature or humidity, unless we place our thumb on the sensor or we blow warm breathe on it :-)

```
$ curl -s http://localhost:8081/foglamp/asset/dht11?limit=5 | jq
[
  {
    "timestamp": "2017-12-30 14:43:16.787",
    "reading": {
      "temperature": 25,
      "humidity": 95
    }
  },
  {
    "timestamp": "2017-12-30 14:43:15.258",
    "reading": {
      "temperature": 25,
      "humidity": 95
    }
  },
  {
    "timestamp": "2017-12-30 14:43:13.729",
    "reading": {
      "temperature": 24,
      "humidity": 95
    }
  },
  {
    "timestamp": "2017-12-30 14:43:12.201",
    "reading": {
      "temperature": 24,
      "humidity": 95
    }
  },
  {
    "timestamp": "2017-12-30 14:43:05.616",
    "reading": {
      "temperature": 22,
      "humidity": 95
    }
  }
]
```

(continues on next page)

(continued from previous page)

```

    }
}
$

```

Needless to say, the North plugin will send the buffered data to the PI system using the OMF plugin or any other north system using the appropriate north plugin.

13.5 South Plugins in C

South plugins written in C/C++ are no different in use to those written in Python, it is merely a case that they are implemented in a different language. The same options of polled or asynchronous methods still exist and the enduser of FogLAMP is not aware in which language the plugin has been written.

13.5.1 Polled Mode

Polled mode is the simplest form of South plugin that can be written, a poll routine is called at an interval defined in the plugin advanced configuration. The South service determines the type of the plugin by examining the mode property in the information the plugin returns from the *plugin_info* call.

Plugin Poll

The plugin *poll* method is called periodically to collect the readings from a poll mode sensor. As with all other calls the argument passed to the method is the handle returned by the *plugin_init* call, the return of the method should be a *Reading* instance that contains the data read.

The *Reading* class consists of

Property	Description
assetName	The asset key of the sensor device that is being read
userTimestamp	A timestamp for the reading data
datapoints	The reading data itself as a set of datapoint instances

More detail regarding the *Reading* class can be found in the section .

It is important that the *poll* method does not block as this will prevent the proper operation of the South microservice. Using the example of our simple DHT11 device attached to a GPIO pin, the *poll* routine could be:

```

/**
 * Poll for a plugin reading
 */
Reading plugin_poll(PLUGIN_HANDLE *handle)
{
    DHT11 *dht11 = (DHT11*)handle;
    return dht11->takeReading();
}

```

Where our *DHT11* class has a method *takeReading* as follows

```

/**
 * Take reading from sensor
 *
 * @param firstReading This flag indicates whether this is the first reading to be
↳taken from sensor,
 * if so get it reliably even if takes multiple retries.
↳Subsequently (firstReading=false),
 * if reading from sensor fails, last good reading is returned.
 */
Reading DHT11::takeReading(bool firstReading)
{
    static uint8_t sensorData[4] = {0,0,0,0};

    bool valid = false;
    unsigned int count=0;
    do {
        valid = readSensorData(sensorData);
        count++;
    } while(!valid && firstReading && count < MAX_SENSOR_READ_RETRIES);

    if (firstReading && count >= MAX_SENSOR_READ_RETRIES)
        Logger::getLogger()->error("Unable to get initial valid reading from
↳DHT11 sensor connected to pin %d even after %d tries", m_pin, MAX_SENSOR_READ_
↳RETRIES);

    vector<Datapoint *> vec;

    ostringstream tmp;
    tmp << ((unsigned int)sensorData[0]) << "." << ((unsigned int)sensorData[1]);
    DatapointValue dpv1(stod(tmp.str()));
    vec.push_back(new Datapoint("Humidity", dpv1));

    ostringstream tmp2;
    tmp2 << ((unsigned int)sensorData[2]) << "." << ((unsigned_
↳int)sensorData[3]);
    DatapointValue dpv2(stod(tmp2.str()));
    vec.push_back(new Datapoint ("Temperature", dpv2));

    return Reading(m_assetName, vec);
}

```

We are creating two *DatapointValues* for the Humidity and Temperature values returned by reading the DHT11 sensor.

Plugin Poll Returning Multiple Values

It is possible in a C/C++ plugin to have a plugin that returns multiple readings in a single call to a poll routine. This is done by setting the interface version of 2.0.0 rather than 1.0.0. In this interface version the *plugin_poll* call returns a vector of *Reading* rather than a single *Reading*.

```

/**
 * Poll for a plugin reading
 */
std::vector<Reading *> *plugin_poll(PLUGIN_HANDLE *handle)
{
    Modbus *modbus = (Modbus *)handle;

```

(continues on next page)

(continued from previous page)

```

    if (!handle)
        throw runtime_error("Bad plugin handle");
    return modbus->takeReading();
}

```

13.5.2 Async IO Mode

In asyncio mode the plugin runs either a separate thread or uses some incoming event from a device or callback mechanism to trigger sending data to FogLAMP. The asynchronous mode uses two additional entry points to the plugin, one to register a callback on which the plugin sends data, *plugin_register_ingest* and another to start the asynchronous behavior *plugin_start*.

Plugin Register Ingest

The *plugin_register_ingest* call is used to allow the south service to pass a callback function to the plugin that the plugin uses to send data to the service every time the plugin has some new data.

```

/**
 * Register ingest callback
 */
void plugin_register_ingest(PLUGIN_HANDLE *handle, INGEST_CB cb, void *data)
{
    MyPluginClass *plugin = (MyPluginClass *)handle;

    if (!handle)
        throw new exception();
    plugin->registerIngest(data, cb);
}

```

The plugin should store the callback function pointer and the data associated with the callback such that it can use that information to pass a reading to the south service. The following code snippets show how a plugin class might store the callback and data and then use it to send readings into FogLAMP at a later stage.

```

/**
 * Record the ingest callback function and data in member variables
 *
 * @param data The Ingest function data
 * @param cb The callback function to call
 */
void MyPluginClass::registerIngest(void *data, INGEST_CB cb)
{
    m_ingest = cb;
    m_data = data;
}

/**
 * Called when a data is available to send to the south service
 *
 * @param points The points in the reading we must create
 */
void MyPluginClass::ingest(Reading& reading)
{

```

(continues on next page)

(continued from previous page)

```

        (*m_ingest)(m_data, reading);
    }

```

Plugin Start

The *plugin_start* method, as with other plugin calls, is called with the plugin handle data that was returned from the *plugin_init* call. The *plugin_start* call will only be called once for a plugin, it is the responsibility of *plugin_start* to take whatever action is required in the plugin in order to start the asynchronous actions of the plugin. This might be to start a thread, register an endpoint for a remote connection or call an entry point in a third party library to start asynchronous processing.

```

/**
 * Start the Async handling for the plugin
 */
void plugin_start (PLUGIN_HANDLE *handle)
{
    MyPluginClass *plugin = (MyPluginClass *)handle;

    if (!handle)
        return;
    plugin->start();
}

/**
 * Start the asynchronous processing thread
 */
void MyPluginClass::start()
{
    m_running = true;
    m_thread = new thread(threadWrapper, this);
}

```

13.5.3 Set Point Control

South plugins can also be used to exert control on the underlying device to which they are connected. This is not intended for use as a substitute for real time control systems, but rather as a mechanism to make non-time critical changes to a device or to trigger an operation on the device.

To make a south plugin support control features there are two steps that need to be taken

- Tag the plugin as supporting control
- Add the entry points for control

Enable Control

A plugin enables control features by means of the flags in the plugin information data structure which is returned by the *plugin_info* entry point of the plugin. The flag value *SP_CONTROL* should be added to the flags of the plugin.

```
/**
 * The plugin information structure
 */
static PLUGIN_INFORMATION info = {
    PLUGIN_NAME,          // Name
    VERSION,              // Version
    SP_CONTROL,           // Flags - add control
    PLUGIN_TYPE_SOUTH,    // Type
    "1.0.0",              // Interface version
    CONFIG                // Default configuration
};
```

Adding this flag will cause the south service to do a number of things when it loads the plugin;

- The south service will attempt to resolve the two control entry points.
- A toggle will be added to the advanced configuration category of the service that will permit the disabling of control services.
- A security category will be added to the south service that contains the access control lists and permissions associated with the service.

Control Entry Points

Two entry points are supported for control operations in the south plugin

- **plugin_write**: which is used to set the value of a parameter within the plugin or device
- **plugin_operation**: which is used to perform an operation on the plugin or device

The south plugin can support one or both of these entry points as appropriate for the plugin.

Write Entry Point

The write entry point is used to set data in the plugin or write data into the device.

The plugin write entry point is defined as follows

```
bool plugin_write(PLUGIN_HANDLE *handle, string name, string value)
```

Where the parameters are;

- **handle** the handle of the plugin instance
- **name** the name of the item to be changed
- **value** a string presentation of the new value to assign top the item

The return value defines if the write was successful or not. True is returned for a successful write.

```
bool plugin_write(PLUGIN_HANDLE *handle, string& name, string& value)
{
    Random *random = (Random *)handle;
```

(continues on next page)

(continued from previous page)

```

    return random->write(operation, name, value);
}

```

In this case the main logic of the write operation is implemented in a class that contains all the plugin logic. Note that the assumption here, and a design pattern often used by plugin writers, is that the *PLUGIN_HANDLE* is actually a pointer to a C++ class instance.

In this case the implementation in the plugin class is as follows:

```

bool Random::write(string& name, string& value)
{
    if (name.compare("mode") == 0)
    {
        if (value.compare("relative") == 0)
        {
            m_mode = RELATIVE_MODE;
        }
        else if (value.compare("absolute") == 0)
        {
            m_mode = ABSOLUTE_MODE;
        }
        Logger::getLogger()->error("Unknown mode requested '%s' ignored.",
↪value.c_str());
        return false;
    }
    else
    {
        Logger::getLogger()->error("Unknown control item '%s' ignored.", name.c_
↪str());
        return false;
    }
    return true;
}

```

In this case the code is relatively simple as we assume there is a single control parameter that can be written, the mode of operation. We look for the known name and if a different name is passed an error is logged and false is returned. If the correct name is passed in we then check the value and take the appropriate action. If the value is not a recognized value then an error is logged and we again return false.

In this case we are merely setting a value within the plugin, this could equally well be done via configuration and would in that case be persisted between restarts. Normally control would not be used for this, but rather for making a change with the connected device itself, such as changing a PLC register value. This is simply an example to demonstrate the mechanism.

Operation Entry Point

The plugin will support an operation entry point. This will execute the given operation synchronously, it is expected that this operation entry point will be called using a separate thread, therefore the plugin should implement operations in a thread safe environment.

The plugin write operation entry point is defined as follows

```

bool plugin_operation(PLUGIN_HANDLE *handle, string& operation, int count, PLUGIN_
↪PARAMETER **params)

```

Where the parameters are;

- **handle** the handle of the plugin instance
- **operation** the name of the operation to be executed
- **count** the number of parameters
- **params** a set of name/value pairs that are passed to the operation

The *operation* parameter should be used by the plugin to determine which operation is to be performed, that operation may also be passed a number of parameters. The count of these parameters are passed to the plugin in the *count* argument and the actual parameters are passed in an array of key/value pairs as strings.

The return from the call is a boolean result of the operation, a failure of the operation or a call to an unrecognized operation should be indicated by returning a false value. If the operation succeeds a value of true should be returned.

The following example shows the implementation of the plugin operation entry point.

```
bool plugin_operation(PLUGIN_HANDLE *handle, string& operation, int count, PLUGIN_
↳PARAMETER **params)
{
    Random *random = (Random *)handle;

    return random->operation(operation, count, params);
}
```

In this case the main logic of the operation is implemented in a class that contains all the plugin logic. Note that the assumption here, and a design pattern often used by plugin writers, is that the *PLUGIN_HANDLE* is actually a pointer to a C++ class instance.

In this case the implementation in the plugin class is as follows:

```
/**
 * SetPoint operation. We support reseeding the random number generator
 */
bool Random::operation(const std::string& operation, int count, PLUGIN_PARAMETER_
↳**params)
{
    if (operation.compare("seed") == 0)
    {
        if (count)
        {
            if (params[0]->name.compare("seed"))
            {
                long seed = strtol(params[0]->value.c_str(), NULL, _
↳10);

                srand(seed);
            }
            else
            {
                return false;
            }
        }
        else
        {
            srand(time(0));
        }
        Logger::getLogger()->info("Reseeded random number generator");
        return true;
    }
    Logger::getLogger()->error("Unrecognised operation %s", operation.c_str());
}
```

(continues on next page)

(continued from previous page)

```

    return false;
}

```

In this example, the operation method checks the name of the operation to perform, only a single operation is supported by this plugin. If this operation name differs the method will log an error and return false. If the operation is recognized it will check for any arguments passed in, retrieve and use it. In this case an optional *seed* argument may be passed.

There is no actual machine connected here, therefore the operation occurs within the plugin. In the case of a real machine the operation would most likely cause an action on a machine, for example a request to the machine to re-calibrate itself.

13.5.4 A South Plugin Example In C/C++: the DHT11 Sensor

Using the same example as before, the DHT11 temperature and humidity sensor, let's look at how to create the plugin in C/C++.

The Software

For this plugin we use the wiringpi C library to connect to the hardware of the Raspberry Pi

```

$ sudo apt-get install wiringpi
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following NEW packages will be installed:
wiringpi
...
$

```

The Plugin

This is the code for the plugin.cpp file that provides the plugin API:

```

/*
 * FogLAMP south plugin.
 *
 * Copyright (c) 2018 OSisoft, LLC
 *
 * Released under the Apache 2.0 Licence
 *
 * Author: Amandeep Singh Arora
 */
#include <dht11.h>
#include <plugin_api.h>
#include <stdio.h>
#include <stdlib.h>
#include <strings.h>
#include <string>
#include <logger.h>
#include <plugin_exception.h>
#include <config_category.h>
#include <rapidjson/document.h>

```

(continues on next page)

(continued from previous page)

```

#include <version.h>

using namespace std;
#define PLUGIN_NAME "dht11_V2"

/**
 * Default configuration
 */
const static char *default_config = QUOTE({
    "plugin" : {
        "description" : "DHT11 C south plugin",
        "type" : "string",
        "default" : PLUGIN_NAME,
        "readonly": "true"
    },
    "asset" : {
        "description" : "Asset name",
        "type" : "string",
        "default" : "dht11",
        "order": "1",
        "displayName": "Asset Name",
        "mandatory" : "true"
    },
    "pin" : {
        "description" : "Rpi pin to which DHT11 is attached",
        "type" : "integer",
        "default" : "7",
        "displayName": "Rpi Pin"
    }
});

/**
 * The DHT11 plugin interface
 */
extern "C" {

/**
 * The plugin information structure
 */
static PLUGIN_INFORMATION info = {
    PLUGIN_NAME,           // Name
    VERSION,               // Version
    0,                     // Flags
    PLUGIN_TYPE_SOUTH,     // Type
    "1.0.0",               // Interface version
    default_config          // Default configuration
};

/**
 * Return the information about this plugin
 */
PLUGIN_INFORMATION *plugin_info()
{
    return &info;
}

```

(continues on next page)

(continued from previous page)

```

/**
 * Initialise the plugin, called to get the plugin handle
 */
PLUGIN_HANDLE plugin_init(ConfigCategory *config)
{
    unsigned int pin;

    if (config->itemExists("pin"))
    {
        pin = stoul(config->getValue("pin"), nullptr, 0);
    }

    DHT11 *dht11= new DHT11(pin);

    if (config->itemExists("asset"))
        dht11->setAssetName(config->getValue("asset"));
    else
        dht11->setAssetName("dht11");

    Logger::getLogger()->info("m_assetName set to %s", dht11->getAssetName());

    return (PLUGIN_HANDLE) dht11;
}

/**
 * Poll for a plugin reading
 */
Reading plugin_poll(PLUGIN_HANDLE *handle)
{
    DHT11 *dht11 = (DHT11*) handle;
    return dht11->takeReading();
}

/**
 * Reconfigure the plugin
 */
void plugin_reconfigure(PLUGIN_HANDLE *handle, string& newConfig)
{
    ConfigCategory conf("dht", newConfig);
    DHT11 *dht11 = (DHT11*) *handle;

    if (conf.itemExists("asset"))
        dht11->setAssetName(conf.getValue("asset"));
    if (conf.itemExists("pin"))
    {
        unsigned int pin = stoul(conf.getValue("pin"), nullptr, 0);
        dht11->setPin(pin);
    }
}

/**
 * Shutdown the plugin
 */
void plugin_shutdown(PLUGIN_HANDLE *handle)
{
    DHT11 *dht11 = (DHT11*) handle;
    delete dht11;
}

```

(continues on next page)

(continued from previous page)

```
}  
};
```

The full source code, including the *DHT11* class can be found in [GitHub](#)

Building FogLAMP and Adding the Plugin

If you have not built FogLAMP yet, follow the steps described . After the build, you can optionally install FogLAMP following steps.

- Clone the *foglamp-south-dht* repository

```
$ git clone https://github.com/fledge-iot/fledge-south-dht.git  
...  
$
```

- Set the environment variable `FOGLAMP_ROOT` to the directory in which you built FogLAMP

```
$ export FOGLAMP_ROOT=~/.foglamp  
$
```

- Go to the location in which you cloned the *foglamp-south-dht* repository and create a build directory and run `cmake` in that directory

```
$ cd ~/.foglamp-south-dht  
$ mkdir build  
$ cd build  
$ cmake ..  
...  
$
```

- Now make the plugin

```
$ make  
$
```

- If you have started FogLAMP from the build directory, copy the plugin into the destination directory

```
$ mkdir -p $FOGLAMP_ROOT/plugins/south/dht  
$ cp libdht.so $FOGLAMP_ROOT/plugins/south/dht  
$
```

- If you have installed FogLAMP by executing `sudo make install`, copy the plugin into the destination directory

```
$ sudo mkdir -p /usr/local/foglamp/plugins/south/dht  
$ sudo cp libdht.so /usr/local/foglamp/plugins/south/dht  
$
```

Note: If you have installed FogLAMP using an alternative *DESTDIR*, remember to add the path to the destination directory to the `cp` command.

- Add service


```
$ curl -sX POST http://localhost:8081/foglamp/service -d '{"name": "dht", "type":  
↪ "south", "plugin": "dht", "enabled": true}'
```

You may now use the C/C++ plugin in exactly the same way as you used a Python plugin earlier.

13.6 C++ Support Classes

A number of support classes exist within the common library that forms part of every FogLAMP plugin.

13.6.1 Reading

The *Reading* class and the associated *Datapoint* and *DatapointValue* classes provide the mechanism within C++ classes to manipulated the reading asset data. The public part of the *Reading* class is currently defined as follows;

```
class Reading {  
    public:  
        Reading(const std::string& asset, Datapoint *value);  
        Reading(const std::string& asset, std::vector<Datapoint *> values);  
        Reading(const std::string& asset, std::vector<Datapoint *> values,   
↪const std::string& ts);  
        Reading(const Reading& orig);  
  
        ~Reading();  
        void addDatapoint(Datapoint *value);  
        Datapoint *removeDatapoint(const std::string&   
↪name);  
  
        std::string toJSON(bool minimal = false) const;  
        std::string getDatapointsJSON() const;  
        // Return AssetName  
        const std::string& getAssetName() const { return m_asset;  
↪ };  
  
        // Set AssetName  
        void setAssetName(std::string assetName) {   
↪m_asset = assetName; };  
        unsigned int getDatapointCount() { return m_values.  
↪size(); };  
  
        void removeAllDatapoints();  
        // Return Reading datapoints  
        const std::vector<Datapoint *> getReadingData() const { return m_  
↪values; };  
  
        // Return reference to Reading datapoints  
        std::vector<Datapoint *>& getReadingData() { return m_values; };  
        unsigned long getId() const { return m_id; };  
        unsigned long getTimestamp() const { return   
↪(unsigned long)m_timestamp.tv_sec; };  
        unsigned long getUserTimestamp() const { return   
↪(unsigned long)m_userTimestamp.tv_sec; };  
        void setId(unsigned long id) { m_id = id; }  
↪;  
        void setTimestamp(unsigned long ts) { m_  
↪timestamp.tv_sec = (time_t)ts; };  
        void setTimestamp(struct timeval tm) { m_  
↪timestamp = tm; };  
        void setTimestamp(const std::string&   
↪timestamp);
```

(continues on next page)

(continued from previous page)

```

        void                                getTimestamp(struct timeval *tm) {
↪ *tm = m_timestamp; };
        void                                setUserTimestamp(unsigned long uTs) {
↪ m_userTimestamp.tv_sec = (time_t)uTs; };
        void                                setUserTimestamp(struct timeval tm) {
↪ m_userTimestamp = tm; };
        void                                setUserTimestamp(const std::string&
↪ timestamp);
        void                                getUserTimestamp(struct timeval *tm)
↪ { *tm = m_userTimestamp; };

        typedef enum dateTimeFormat { FMT_DEFAULT, FMT_STANDARD, FMT_ISO8601 }
↪ readingTimeFormat;

        // Return Reading asset time - ts time
        const std::string getAssetDateTime(readingTimeFormat dateTimeFmt =
↪ FMT_DEFAULT, bool addMs = true) const;
        // Return Reading asset time - user_ts time
        const std::string getAssetDateUserTime(readingTimeFormat dateTimeFmt
↪ = FMT_DEFAULT, bool addMs = true) const;
    }

```

The *Reading* class contains a number of items that are mapped to the JSON representation of data that is sent to the FogLAMP storage service and are used by the various services and plugins within FogLAMP.

- **Asset Name:** The name of the asset. The asset name is set in the constructor of the reading and retrieved via the *getAssetName()* method.
- **Timestamp:** The timestamp when the reading was first seen within FogLAMP.
- **User Timestamp:** The timestamp for the actual data in the reading. This may differ from the value of Timestamp if the device itself is able to supply a timestamp value.
- **Datapoints:** The actual data of a reading stored in a Datapoint class.

The *Datapoint* class provides a name for each data point within a *Reading* and the tagged type data for the reading value. The public definition of the *Datapoint* class is as follows;

```

class Datapoint {
    public:
        /**
         * Construct with a data point value
         */
        Datapoint(const std::string& name, DatapointValue& value) : m_
↪ name(name), m_value(value);
        ~Datapoint();
        /**
         * Return asset reading data point as a JSON
         * property that can be included within a JSON
         * document.
         */
        std::string toJSONProperty();
        const std::string getName() const;
        void setName(std::string name);
        const DatapointValue getData() const;
        DatapointValue& getData();
    }

```

Closely associated with the *Datapoint* is the *DatapointValue* which uses a tagged union to store the values. The public definition of the *DatapointValue* is as follows;

```
class DatapointValue {
public:
    /**
     * Construct with a string
     */
    DatapointValue(const std::string& value)
    {
        m_value.str = new std::string(value);
        m_type = T_STRING;
    };
    /**
     * Construct with an integer value
     */
    DatapointValue(const long value)
    {
        m_value.i = value;
        m_type = T_INTEGER;
    };
    /**
     * Construct with a floating point value
     */
    DatapointValue(const double value)
    {
        m_value.f = value;
        m_type = T_FLOAT;
    };
    /**
     * Construct with an array of floating point values
     */
    DatapointValue(const std::vector<double>& values)
    {
        m_value.a = new std::vector<double>(values);
        m_type = T_FLOAT_ARRAY;
    };
    /**
     * Construct with an array of Datapoints
     */
    DatapointValue(std::vector<Datapoint*>& values, bool isDict)
    {
        m_value.dpa = values;
        m_type = isDict? T_DP_DICT : T_DP_LIST;
    }
    /**
     * Construct with an Image
     */
    DatapointValue(const DPIImage& value)
    {
        m_value.image = new DPIImage(value);
        m_type = T_IMAGE;
    }
    /**
     * Construct with a DataBuffer

```

(continues on next page)

(continued from previous page)

```

    */
    DatapointValue(const DataBuffer& value)
    {
        m_value.dataBuffer = new DataBuffer(value);
        m_type = T_DATABUFFER;
    }

    /**
     * Construct with an Image Pointer, the
     * image becomes owned by the datapointValue
     */
    DatapointValue(DPImage *value)
    {
        m_value.image = value;
        m_type = T_IMAGE;
    }

    /**
     * Construct with a DataBuffer
     */
    DatapointValue(DataBuffer *value)
    {
        m_value.dataBuffer = value;
        m_type = T_DATABUFFER;
    }

    /**
     * Construct with a 2 dimensional array of floating point values
     */
    DatapointValue(const std::vector< std::vector<double> >& values)
    {
        m_value.a2d = new std::vector< std::vector<double> >();
        for (auto row : values)
        {
            m_value.a2d->push_back(std::vector<double>(row));
        }
        m_type = T_2D_FLOAT_ARRAY;
    };

    /**
     * Copy constructor
     */
    DatapointValue(const DatapointValue& obj);

    /**
     * Assignment Operator
     */
    DatapointValue& operator=(const DatapointValue& rhs);

    /**
     * Destructor
     */
    ~DatapointValue();

    /**
     * Set the value of a datapoint, this may

```

(continues on next page)

(continued from previous page)

```

    * also cause the type to be changed.
    * @param value      An integer value to set
    */
    void setValue(long value)
    {
        m_value.i = value;
        m_type = T_INTEGER;
    }

    /**
     * Set the value of a datapoint, this may
     * also cause the type to be changed.
     * @param value      A floating point value to set
     */
    void setValue(double value)
    {
        m_value.f = value;
        m_type = T_FLOAT;
    }

    /** Set the value of a datapoint to be an image
     * @param value The image to set in the data point
     */
    void setValue(const DPImage& value)
    {
        m_value.image = new DPImage(value);
        m_type = T_IMAGE;
    }

    /**
     * Return the value as a string
     */
    std::string toString() const;

    /**
     * Return string value without trailing/leading quotes
     */
    std::string toStringValue() const { return *m_value.str; };

    /**
     * Return long value
     */
    long toInt() const { return m_value.i; };

    /**
     * Return double value
     */
    double toDouble() const { return m_value.f; };

    // Supported Data Tag Types
    typedef enum DatapointTag
    {
        T_STRING,
        T_INTEGER,
        T_FLOAT,
        T_FLOAT_ARRAY,
        T_DP_DICT,
        T_DP_LIST,
    }

```

(continues on next page)

(continued from previous page)

```

        T_IMAGE,
        T_DATABUFFER,
        T_2D_FLOAT_ARRAY
    } dataTagType;

    /**
     * Return the Tag type
     */
    dataTagType getType() const
    {
        return m_type;
    }

    std::string getTypeStr() const
    {
        switch(m_type)
        {
            case T_STRING: return std::string("STRING");
            case T_INTEGER: return std::string("INTEGER");
            case T_FLOAT: return std::string("FLOAT");
            case T_FLOAT_ARRAY: return std::string("FLOAT_ARRAY");
            case T_DP_DICT: return std::string("DP_DICT");
            case T_DP_LIST: return std::string("DP_LIST");
            case T_IMAGE: return std::string("IMAGE");
            case T_DATABUFFER: return std::string("DATABUFFER");
            case T_2D_FLOAT_ARRAY: return std::string("2D_FLOAT_
→ARRAY");

            default: return std::string("INVALID");
        }
    }

    /**
     * Return array of datapoints
     */
    std::vector<Datapoint*>& getDpVec()
    {
        return m_value.dpa;
    }

    /**
     * Return array of float
     */
    std::vector<double*>& getDpArr()
    {
        return m_value.a;
    }

    /**
     * Return 2D array of float
     */
    std::vector<std::vector<double> >& getDp2DArr()
    {
        return m_value.a2d;
    }

    /**
     * Return the Image

```

(continues on next page)

(continued from previous page)

```

    */
    DImage *getImage()
    {
        return m_value.image;
    }

    /**
     * Return the DataBuffer
     */
    DataBuffer *getDataBuffer()
    {
        return m_value.dataBuffer;
    }
};

```

The *DatapointValue* can store data in as a number of types

Type	C++ Representation
T_STRING	Pointer to std::string
T_INTEGER	long
T_FLOAT	double
T_FLOAT_ARRAY	Pointer to std::vector<double>
T_2D_FLOAT_ARRAY	Pointer to std::vector<std::vector<double>>
T_DP_DICT	Pointer to std::vector<Datapoint *>
T_DP_LIST	Pointer to std::vector<Datapoint *>
T_IMAGE	Pointer to DImage
T_DATABUFFER	Pointer to DataBuffer

13.6.2 Configuration Category

The *ConfigCategory* class is a support class for managing configuration information within a plugin and is passed to the plugin entry points. The public definition of the class is as follows;

```

class ConfigCategory {
public:
    enum ItemType {
        UnknownType,
        StringItem,
        EnumerationItem,
        JsonItem,
        BoolItem,
        NumberItem,
        DoubleItem,
        ScriptItem,
        CategoryType,
        CodeItem
    };

    ConfigCategory(const std::string& name, const std::string& json);
    ConfigCategory() {};
    ConfigCategory(const ConfigCategory& orig);
    ~ConfigCategory();
    void addItem(const std::string& name,
    ↪const std::string description,

```

(continues on next page)

(continued from previous page)

```

↪const std::string def,
        void
↪const std::string description,
        void
↪std::string& value,
        void
↪options);
        void
        void
        void
        bool
↪subCategories);
        void
↪description);
        std::string
        std::string
        unsigned int
        bool
↪const;
        bool
↪name, const std::string& displayName);
        std::string
↪const;
        std::string
↪const;
        std::string
↪name) const;
        std::string
↪const;
        bool
↪const std::string& value);
        std::string
↪name) const;
        std::vector<std::string>
↪const;
        std::string
↪const;
        std::string
↪const;
        std::string
↪const;
        bool
↪const;
        bool
↪name) const;
        bool
        bool
        bool
↪const;
        bool
↪const;
        bool
↪const;
        std::string
        std::string
↪const;

        const std::string& type,
        const std::string& value);
addItem(const std::string& name,
        const std::string def, const
        const std::vector<std::string>
removeItems();
removeItemsType(ItemType type);
keepItemsType(ItemType type);
extractSubcategory(ConfigCategory &
setDescription(const std::string&
getName() const;
getDescription() const;
getCount() const;
itemExists(const std::string& name)
setItemDisplayName(const std::string&
getValue(const std::string& name)
getType(const std::string& name)
getDescription(const std::string&
getDefault(const std::string& name)
setDefault(const std::string& name,
getDisplayname(const std::string&
getOptions(const std::string& name)
getLength(const std::string& name)
getMinimum(const std::string& name)
getMaximum(const std::string& name)
isString(const std::string& name)
isEnumeration(const std::string&
isJSON(const std::string& name) const;
isBool(const std::string& name) const;
isNumber(const std::string& name)
isDouble(const std::string& name)
isDeprecated(const std::string& name)
toJSON(const bool full=false) const;
itemsToJSON(const bool full=false)

```

(continues on next page)

(continued from previous page)

```

        ConfigCategory&
        ConfigCategory&
        void
        void
        std::string
↪itemName) const;
        enum ItemAttribute
↪MANDATORY_ATTR, FILE_ATTR};
        std::string
↪itemName,
                                ItemAttribute_
↪itemAttribute) const;
    }

```

```

        operator=(ConfigCategory const& rhs);
        operator+=(ConfigCategory const& rhs);
        setItemsValueFromDefault();
        checkDefaultValuesOnly() const;
        itemToJSON(const std::string&_
        { ORDER_ATTR, READONLY_ATTR,_
        getItemAttribute(const std::string&_

```

Although *ConfigCategory* is a complex class, only a few of the methods are commonly used within a plugin

- **itemExists:** - used to test if an expected configuration item exists within the configuration category.
- **getValue:** - return the value of a configuration item from within the configuration category
- **isBool:** - tests if a configuration item is of boolean type
- **isNumber:** - tests if a configuration item is a number
- **isDouble:** - tests if a configuration item is valid to be represented as a double
- **isString:** - tests if a configuration item is a string

13.6.3 Logger

The *Logger* class is used to write entries to the syslog system within FogLAMP. A singleton *Logger* exists which can be obtained using the following code snippet;

```

Logger *logger = Logger::getLogger();
logger->error("An error has occurred within the plugin processing");

```

It is then possible to log messages at one of five different log levels; *debug*, *info*, *warn*, *error* or *fatal*. Messages may be logged using standard printf formatting strings. The public definition of the *Logger* class is as follows;

```

class Logger {
public:
    Logger(const std::string& application);
    ~Logger();
    static Logger *getLogger();
    void debug(const std::string& msg, ...);
    void printLongString(const std::string&);
    void info(const std::string& msg, ...);
    void warn(const std::string& msg, ...);
    void error(const std::string& msg, ...);
    void fatal(const std::string& msg, ...);
    void setMinLevel(const std::string& level);
};

```

The various log levels should be used as follows;

- **debug:** should be used to output messages that are relevant only to a programmer that is debugging the plugin.
- **info:** should be used for information that is meaningful to the end users, but should not normally be logged.

- **warn:** should be used for warning messages that will normally be logged but reflect a condition that does not prevent the plugin from operating.
- **error:** should be used for conditions that cause a temporary failure in processing within the plugin.
- **fatal:** should be used for conditions that cause the plugin to fail processing permanently, possibly requiring a restart of the microservice in order to resolve.

13.7 Hybrid Plugins

In addition to plugins written in Python and C/C++ it is possible to have a hybrid plugin that is a combination of an existing plugin and configuration for that plugin. This is useful in a situation whereby there are multiple sensors or devices that you connect to FogLAMP that have common configuration. It allows devices to be added without repeating the common configuration.

Using our example of a *DHT11* sensor connected to a GPIO pin, if we wanted to create a new plugin for a *DHT11* that was always connected to pin 4 then we could do this by creating a JSON file as below that supplies a fixed default value for the GPIO pin.

```
{
  "description" : "A DHT11 sensor connected to GPIO pin 4",
  "name" : "DHT11-4",
  "connection" : "DHT11",
  "defaults" : {
    "pin" : {
      "default" : "4"
    }
  }
}
```

This creates a new hybrid plugin called DHT11-4 that is installed by copying this file into the plugins/south/DHT11-4 directory of your installation. Once installed it can be treated as any other south plugin within FogLAMP. The effect of this hybrid plugin is to load the *DHT11* plugin and always set the configuration parameter called “pin” to the value “4”. The item “pin” will be hidden from the user in the FogLAMP GUI when they create the instance of the plugin. This allows for a simpler and more streamlined user experience when adding plugins with common configuration.

The items in the JSON file are;

Name	Description
description	A description of the hybrid plugin. This will appear to the right of the selection list in the FogLAMP user interface when the plugin is selected.
name	The name of the plugin itself. This must match the filename of the JSON file and also the name of the directory the file is placed in.
connection	The name of the underlying plugin that will be used as the basis for this hybrid plugin. This must be a C/C++ or Python plugin, it can not be another hybrid plugin.
defaults	The set of values to default in this hybrid plugin. These are configuration parameters of the underlying plugin that will be fixed in the hybrid plugin. Each hybrid plugin can have one or more values here.

It may not be difficult to enter the GPIO pin in each case in this example, where it becomes more useful is for plugins such as *Modbus* where a complex map is required to be entered in a JSON document. By using a hybrid plugin we can define the map we need once and then add new sensors of the same type without having to repeat the map. An example of this would be the Flir AX8 camera that requires a total of 176 Modbus registers to be mapped into 88 different values in an asset. A hybrid plugin *foglamp-south-FlirAX8* defines that mapping once and as a result adding a new Flir AX8 camera is as simple as selecting the FlirAX8 hybrid plugin and entering the IP address of the camera.

13.8 North Plugins

North plugins are used in North tasks and micro services to extract data buffered in FogLAMP and send it Northbound, i.e. to a server or a service in the Cloud or in an Enterprise data center. North plugins may be written in Python or C/C++, a number of different north plugins are available as examples that may be used when creating new plugins.

A north plugin has a limited number of entry points that it much support, these entry points are the same for both Python and C/C++ north plugins.

Entry Point	Description
plu- gin_info	Return information about the plugin including the configuration for the plugin. This is the same as <code>plugin_info</code> in all other types of plugin and is part of the standard plugin interface.
plu- gin_init	Also part of the standard plugin interface. This call is passed the request configuration of the plugin and should be used to do any initialization of the plugin.
plu- gin_send	This entry point is the north plugin specific entry point that is used to send data from FogLAMP. This will be called repeatedly with blocks of readings.
plu- gin_shutdown	Part of the standard plugin interface, this will be called when the plugin is no longer required and will be the final call to the plugin.
plu- gin_register	Register the callback function used for control writes and operations.

The life cycle of a plugin is very similar regardless of if it is written in Python or C/C++, the `plugin_info` call is made first to determine data about the plugin. The plugin is then initialized by calling the `plugin_init` entry point. The `plugin_send` entry point will be called multiple times to send the actual data and finally the `plugin_shutdown` entry point will be called.

In the following sections each of these calls will be described in detail and samples given in both C/C++ and Python.

13.8.1 Python Plugins

Python plugins are loaded dynamically and executed either within a task, known as the *sending_task* or *north* task. This code is implemented in C++ and embedded a Python interpreter that is used to run the Python plugin.

The `plugin_info` call

The `plugin_info` call is the first call that will be made to a plugin and is called only once. It is part of the standard plugin interface that is implemented by north, south, filter, notification rule and notification delivery plugins. No arguments are passed to this call and it should return a *plugin information structure* as a Python dict.

A typical implementation for a simple north plugin simply returns a DICT as follows

```
def plugin_info():
    """ Used only once when call will be made to a plugin.

    Args:
    Returns:
        Information about the plugin including the configuration for the plugin
    """
    return {
        'name': 'http',
        'version': '1.9.1',
        'type': 'north',
```

(continues on next page)

(continued from previous page)

```

    'interface': '1.0',
    'config': _DEFAULT_CONFIG
}

```

The items in the structure returned by *plugin_info* are

Name	Description
name	The name of the plugin
version	The version of the plugin. Typically this is the same as the version of FogLAMP it is designed to work with but is not constrained to be the same.
type	The type of the plugin, in this case the type will always be <i>north</i>
interface	The version of the plugin interface that the plugin supports. In this case the version is 1.0
config	The DICT that defines the configuration that the plugin has as default.

In the case above *_DEFAULT_CONFIG* is another Python DICT that contains the defaults for the plugin configuration and will be covered in the Configuration section.

Configuration

Configuration within FogLAMP is represented in a JSON structure that defines a name, value, default, type and a number of other optional parameters. The configuration process works by the plugins having a default configuration that they return from the *plugin_init* call. The FogLAMP configuration code will then combine this with a copy of that configuration that it holds. On the first time a service is created, with no previously held configuration, the configuration manager will take the default values and make those the actual values. The user may then update these to set non-default values. In subsequent executions of the plugin these values will be combined with the defaults to create the in use configuration that is passed to the *plugin_init* entry point. The mechanism is designed to allow initial execution of a plugin, but also to allow upgrade of a plugin to create new configuration items for the plugins whilst preserving previous configuration values set by the user.

A sample default configuration of http north python based plugin is shown below.

```

{
  "plugin": {
    "description": "HTTP North Plugin",
    "type": "string",
    "default": "http_north",
    "readonly": "true"
  },
  "url": {
    "description": "Destination URL",
    "type": "string",
    "default": "http://localhost:6683/sensor-reading",
    "order": "1",
    "displayName": "URL"
  },
  "source": {
    "description": "Source of data to be sent on the stream. May be either_
↪ readings or statistics.",
    "type": "enumeration",
    "default": "readings",
    "options": ["readings", "statistics"],
    "order": "2",

```

(continues on next page)

(continued from previous page)

```

        "displayName": "Source"
    },
    "verifySSL": {
        "description": "Verify SSL certificate",
        "type": "boolean",
        "default": "false",
        "order": "3",
        "displayName": "Verify SSL"
    }
}

```

Items marked as “readonly” : “true” will not be presented to the user. The *displayName* and *order* properties are only used by the user interface to display the configuration item. The description, type and default are used by the API to verify the input and also set the initial values when a new configuration item is created.

Rules can also be given to the user interface to define the validity of configuration items based upon the values of others, or example

```

{
    "applyFilter": {
        "description": "Should filter be applied before processing data",
        "type": "boolean",
        "default": "false",
        "order": "4",
        "displayName": "Apply Filter"
    },
    "filterRule": {
        "description": "JQ formatted filter to apply (only applicable if applyFilter_
↪is True)",
        "type": "string",
        "default": ".[]",
        "order": "5",
        "displayName": "Filter Rule",
        "validity": "applyFilter == \"true\""
    }
}

```

This will only allow entry to the *filterRule* configuration item if the *applyFilter* item has been set to true.

The plugin_init call

The *plugin_init* call will be invoked after the *plugin_info* call has been called to obtain the information regarding the plugin. This call is designed to allow the plugin to do any initialization that is required and also creates the handle will be used in all subsequent calls to identify the instance of the plugin.

The *plugin_init* is passed a Python DICT as the only argument, this DICT contains the modified configuration for the plugin that is created by taking the default plugin configuration returned by *plugin_info* and adding to that the values the user has configured previously. This is the working configuration that the plugin should use.

The typical implementation of the *plugin_init* call will create an instance of a Python class which is the main body of the plugin. An object will then be returned which is the handle that will be passed into subsequent calls. This handle in a simple plugin, is commonly a Python DICT that is the configuration of the plugin, however any values may be returned. The caller treats the handle as opaque data that it stores and passed to further calls to the plugin, it will never look inside that object or have any expectations as to what is stored within that object.

The *foglamp-north-http* plugin implementation of *plugin_init* is shown below as an example

```
def plugin_init(data):
    """ Used for initialization of a plugin.

    Args:
        data - Plugin configuration
    Returns:
        Dictionary of a Plugin configuration
    """
    global http_north, config
    http_north = HttpNorthPlugin()
    config = data
    return config
```

In this case the plugin creates an object that implements the functionality and stores that object in a global variable. This can be done as only one instance of the north plugin exists within a single process. It is however perhaps better practice to return the instance of the class in the handle rather than use a global variable. Using a global is not recommended for filter plugins as multiple instances of a filter may exist within a single process. In this case the plugin uses the configuration as the handle it returns.

The plugin_send call

The *plugin_send* call is the main entry point of a north plugin, it is used to send set of readings north to the destination system. It is responsible for both the communication to that system and the translation of the internal representation of the reading data to the representation required by the external system.

The communication performed by the *plugin_send* routine should use the Python 3 asynchronous I/O primitives, the definition of the *plugin_send* entry point must also use the *async* keyword.

The *plugin_send* entry point is passed 3 arguments, the plugin handle, the data to send and a *stream_id*.

```
async def plugin_send(handle, payload, stream_id):
```

The handle is the opaque data returned by the call to *plugin_init* and may be used by the plugin to store data between invocations. The *payload* is a set of readings that should be sent, see below for more details on payload handling. The *stream_id* is an integer that uniquely identifies the connection from this FogLAMP instance to the destination system. This id can be used if the plugin needs to have a unique identifier but in most cases can be ignored.

The *plugin_send* call returns three values, a boolean that indicates if any data has been sent, the object id of the last reading sent and the number of readings sent.

The code below is the *plugin_send* entry point for the http north plugin.

```
async def plugin_send(handle, payload, stream_id):
    """ Used to send the readings block from north to the configured destination.

    Args:
        handle - An object which is returned by plugin_init
        payload - A List of readings block
        stream_id - An Integer that uniquely identifies the connection from FogLAMP_
        ↪ instance to the destination system
    Returns:
        Tuple which consists of
        - A Boolean that indicates if any data has been sent
        - The object id of the last reading which has been sent
        - Total number of readings which has been sent to the configured destination
    """
    try:
```

(continues on next page)

(continued from previous page)

```

        is_data_sent, new_last_object_id, num_sent = await http_north.send_
↪payloads(payload)
    except asyncio.CancelledError:
        pass
    else:
        return is_data_sent, new_last_object_id, num_sent

```

The plugin_shutdown call

The *plugin_shutdown* call is the final entry that is required for Python north plugin, it is called by the north service or task just prior to the task terminating or in a north service if the configuration is allowed, see reconfiguration below. The *plugin_shutdown* call is passed the plugin handle and should perform any cleanup required by the plugin.

```

def plugin_shutdown(handle):
    """ Used when plugin is no longer required and will be final call to shutdown the
↪plugin. It should do any necessary cleanup if required.

    Args:
        handle - Plugin handle which is returned by plugin_init
    Returns:
        """

```

The call should not return any data. Once called the handle should no longer be regarded as valid and no further calls will be made to the plugin using this handle.

Reconfiguration

Unlike other plugins within FogLAMP the north plugins do not have a reconfiguration entry point, this is due to the original nature of the north implementation in FogLAMP which used short lived tasks in order to send data out the north. Each new execution created a new task with new configuration, it was therefore felt that reconfiguration added a complexity to the north plugins that could be avoided.

Since the introduction of the feature that allows the north to be run as an always on service however this has become an issue. It is resolved by closing down the plugin, calling *plugin_shutdown* and then restarting by called *plugin_init* to pass new configuration and retrieve a new plugin handle with that new configuration.

Payload Handling

The payload that is passed to the *plugin_send* routine is a Python list of readings, each reading is encoded as a Python DICT. The properties of the reading dict are;

Key	Description
id	The ID of the reading. Each reading is given an integer id that is an increasing value, it is these id values that are used to track how much data is sent via north plugin. One of the returns from the <i>plugin_send</i> routine is the id of the last reading that was successfully sent.
asset_code	The asset code of the reading. Typical a south service will generate reading for one or more asset codes. These asset codes are used to identify the source of the data. Multiple asset codes may appear in a single block of readings passed to the <i>plugin_send</i> routine.
reading	A nested Python DICT that stores the actual data points associated to the reading. These reading DICT's will contain a key/value pair for each data point within the asset. The value of this pair is the value of the data point and may be numeric, string, an array, or a nested object.
ts	The timestamp when the reading was first seen by the system.
user_ts	The timestamp of the data in the reading. This may be the same as <i>ts</i> above or in some cases may be a timestamp that has been received from the source of the data itself. This timestamp is the one that should be considered the most accurately represents the timestamp of the data.

A sample payload is shown below.

```
[{'reading': {'sinusoid': 0.0}, 'asset_code': 'sinusoid', 'id': 1, 'ts': '2021-09-27_
↪06:55:52.692000+00:00', 'user_ts': '2021-09-27 06:55:49.947058+00:00'},
{'reading': {'sinusoid': 0.104528463}, 'asset_code': 'sinusoid', 'id': 2, 'ts': '2021-
↪09-27 06:55:52.692000+00:00', 'user_ts': '2021-09-27 06:55:50.947110+00:00'}]
```

13.8.2 C/C++ Plugins

The flow of a C/C++ plugin is very similar to that of a Python plugin, the entry points vary slightly compared to Python, mostly for language reasons.

The plugin_info entry point

The *plugin_info* is again the first entry point that will be called, in the case a C/C++ plugin it will return a pointer to a *PLUGIN_INFORMATION* structure, this structure contains the same elements there are seen in the Python DICT that is returned by Python plugins.

```
static PLUGIN_INFORMATION info = {
    PLUGIN_NAME,                // Name
    VERSION,                    // Version
    0,                          // Flags
    PLUGIN_TYPE_NORTH,          // Type
    "1.0.0",                    // Interface version
    default_config               // Configuration
}
```

It should be noted that the *PLUGIN_INFORMATION* structure instance is declared as static. All global variables declared with a C/C++ plugin should be declared as static as the mechanism for loading the plugins will share global variables between plugins. Using true global variables can create unexpected interactions between plugins.

The items are

Name	Description
name	The name of the plugin.
version	The version of the plugin expressed as a string. This usually but not always matches the current version of FogLAMP.
flags	A bitmap of flags that give extra information about the plugin.
inter-face	The interface version, currently north plugins are at interface version 1.0.0.
config	The default configuration for the plugin. In C/C++ plugins this is returned as a string containing the JSON structure.

A number of flags are supported by the plugins, however a small subset are supported in north plugins, this subset consists of

Name	Description
SP_PERSIST_DATA	The plugin persists data and uses the data persistence API extensions.
SP_BUILTIN	The plugin is builtin with the FogLAMP core package. This should not be used for any user added plugins.

A typical implementation of the *plugin_info* entry would merely return the *PLUGIN_INFORMATION* structure for the plugin.

```

PLUGIN_INFORMATION *plugin_info()
{
    return &info;
}

```

More complex implementations may tailor the content of the information returned based upon some criteria determined at run time. An example of such a scenario might be to tailor the default configuration based upon some element of discovery that occurs at run time. For example if the plugin is designed to send data to another service the *plugin_info* entry point could perform some service discovery and update a set of options for an enumerated type in the default configuration. This would allow the user interface to give the user a selection list of all the service instances that it found when the plugin was run.

The plugin_init entry point

The *plugin_init* entry point is called once the configuration of the plugin has been constructed by combining the default configuration with any stored configuration that the user has set for the plugin. The configuration is passed as a pointer to a C++ object of class *ConfigCategory*. This object may then be used to extract data from the configuration.

The *plugin_init* call should be used to initialize the plugin itself and to extract the configuration for the *ConfigCategory* instance and store within the instance of the plugin. Details regarding the use of the *ConfigCategory* class can be found in the C++ Support Class section of the Plugin Developers Guide. Typically the north plugin will create an instance of a class that implements the functionality required, store the configuration in that class and return a pointer to that instance as the handle for the plugin. This will ensure that subsequent calls can access that class instance and the associated state, since all future calls will be passed the handle as an argument.

The following is perhaps the most generic form of the *plugin_init* call.

```

PLUGIN_HANDLE plugin_init(ConfigCategory *configData)
{
    return (PLUGIN_HANDLE) (new myNorthPlugin(configData));
}

```

In this case it assumes we have a class, *myNorthPlugin* that implements the functionality of the plugin. The constructor takes the *ConfigCategory* pointer as an argument and performs all required initialization from that configuration category.

The *plugin_send* entry point

The *plugin_send* entry point, as with Python plugins already describe, is the heart of a north plugin. It is called with the plugin handle and a block of readings data to be sent north. Typically the *plugin_send* will extract the object created in the *plugin_init* call from the handle and then call the functionality within that object to perform whatever translation and communication logic is required to send the reading data.

```
uint32_t plugin_send(PLUGIN_HANDLE handle, std::vector<Reading *>& readings)
{
    myNorthPlugin *plugin = (myNorthPlugin *)handle;
    return plugin->send(readings);
}
```

The block of readings is sent as a C++ standard template library vector of pointers to instance of the *Reading* class, also covered above in the section on C++ Support Classes.

The return from the *plugin_send* function should be a count of the number of readings sent by the plugin.

The *plugin_shutdown* entry point

The *plugin_shutdown* entry point is called when the plugin is no longer required. It should do any necessary cleanup required. As with other entry points, it is called with the handle that was returned by *plugin_init*. In the case of our simple plugin that might simple be to delete the C++ object that implements the plugin functionality.

```
uint32_t plugin_shutdown(PLUGIN_HANDLE handle)
{
    myNorthPlugin *plugin = (myNorthPlugin *)handle;
    delete plugin;
}
```

The *plugin_register* entry point

The *plugin_register* entry point is used to pass two function pointers to the plugin. These functions pointers are the functions that should be called when either a set point write or a set point operation is required. The plugin should store these function pointers for later use.

```
void plugin_register(PLUGIN_HANDLE handle, (bool (*write)(char *name, char *value,
↳ControlDestination destination, ...), int (* operation)(char *operation, int_
↳paramCount, char *parameters[], ControlDestination destination, ...))
{
    myNorthPlugin *plugin = (myNorthPlugin *)handle;
    plugin->setpointCallbacks(write, operation);
}
```

This call will only be made if the plugin included the *SP_CONTROL* option in the flags field of the *PLUGIN_INFORMATION* structure.

13.8.3 Set Point Control

FogLAMP supports multiple paths for set point control, one of these paths allows for a north service to be bi-directional, with the north plugin receiving a trigger from the system north of FogLAMP to perform a set point control. This trigger may be the north plugin polling the system or a protocol response from the north.

Set point control is only available for north services, it is not supported for north tasks and will be ignored.

When the north plugin requires a set point write operation to be performed it calls the *write* callback that was passed to the plugin in the *plugin_register* entry point. This callback takes a number of arguments;

- The name of the set point to be written.
- The value to write to the set point. This is expressed as a string always.
- The destination of the write operation. This is passed using the *ControlDestination* enumerated type. Currently this may be one of
 - **DestinationBroadcast**: send the write operation to all south services that support control.
 - **DestinationAsset**: send the write request to the south service responsible for ingesting the given asset. The asset is passed as the next argument in the *write* call.
 - **DestinationService**: send the write request to the named south service.

For example if the north plugin wishes to write the set point called *speed* with the value 28 in the south service called *Motor Control* it would make a call as follows.

```
(*m_write)("speed", "28", DestinationService, "Motor Control");
```

Assuming the member variable *m_write* was used to store the function pointer of the *write* callback.

If the north plugin requires an operation to be performed, rather than a write, then it should call the *operation* called which was passed to it in the *plugin_register* call. This callback takes a set of arguments;

- The name of the operation to execute.
- The number of parameters the operation should be passed.
- An array of parameters, as strings, to pass to the operation
- The destination of the operation, this is the same set of destinations as per the write call.

13.9 Storage Plugins

Storage plugins are used to interact with the Storage Microservice and provide the persistent storage of information for FogLAMP.

The current version of FogLAMP comes with three storage plugins:

- The **SQLite plugin**: this is the default plugin and it is used for general purpose storage on constrained devices.
- The **SQLite In Memory plugin**: this plugin can be used in conjunction with one of the other storage plugins and will provide an in memory storage system for reading data only. Configuration data is stored using the *SQLite* or *PostgreSQL* plugins.
- The **PostgreSQL plugin**: this plugin can be set on request (or it can be built as a default plugin from source) and it is used for a more significant demand of storage on relatively larger systems.

13.9.1 Data and Metadata

Persistency is split in two blocks:

- **Metadata persistency:** it refers to the storage of metadata for FogLAMP, such as the configuration of the plugins, the scheduling of jobs and tasks and the the storage of statistical information.
- **Data persistency:** it refers to the storage of data collected from sensors and devices by the South microservices. The *SQLite In Memory* plugin is an example of a storage plugin designed to store only the data.

In the current implementation of FogLAMP, metadata and data use the same Storage plugin by default. Administrators can select different plugins for these two categories of data, with the most common configuration of this type to use the SQLite In Memory storage service for data and SQLite for the metadata. This is set by editing the storage configuration file. Currently there is no interface within FogLAMP to change the storage configuration.

The storage configuration file is stored in the FogLAMP data directory as `etc/storage.json`, the default storage configuration file is

```
{
  "plugin": {
    "value": "sqlite",
    "description": "The main storage plugin to load"
  },
  "readingPlugin": {
    "value": "",
    "description": "The storage plugin to load for readings data. If blank the main_
↪storage plugin is used."
  },
  "threads": {
    "value": "1",
    "description": "The number of threads to run"
  },
  "managedStatus": {
    "value": "false",
    "description": "Control if FogLAMP should manage the storage provider"
  },
  "port": {
    "value": "0",
    "description": "The port to listen on"
  },
  "managementPort": {
    "value": "0",
    "description": "The management port to listen on."
  }
}
```

This sets the storage plugin to use as the *SQLite* plugin and leaves the *readingPlugin* blank. If the *readingPlugin* is blank then readings will be stored via the main plugin, if it is populated then a separate plugin will be used to store the readings. As an example, to store the readings in the *SQLite In Memory* plugin the `storage.json` file would be

```
{
  "plugin": {
    "value": "sqlite",
    "description": "The main storage plugin to load"
  },
  "readingPlugin": {
    "value": "sqlitememory",
    "description": "The storage plugin to load for readings data. If blank the main_
↪storage plugin is used."
  }
}
```

(continues on next page)

(continued from previous page)

```

},
"threads": {
  "value": "1",
  "description": "The number of threads to run"
},
"managedStatus": {
  "value": "false",
  "description": "Control if FogLAMP should manage the storage provider"
},
"port": {
  "value": "0",
  "description": "The port to listen on"
},
"managementPort": {
  "value": "0",
  "description": "The management port to listen on."
}
}

```

FogLAMP must be restarted for changes to the storage.json file to take effect.

In addition to the definition of the plugins to use, the storage.json file also has a number of other configuration options for the storage service.

- **threads:** The number of threads to use to accept incoming REST requests. This is normally set to 1, increasing the number of threads has minimal impact on performance in normal circumstances.
- **managedStatus:** This configuration option allows FogLAMP to manage the underlying storage system. If, for example you used a database server and you wished FogLAMP to start and stop that server as part of the FogLAMP start up and shut down procedure you would set this option to “true”.
- **port:** This option can be used to make the storage service listen on a fixed port. This is normally not required, but can be used for diagnostic purposes.
- **managementPort:** As with *port* above this can be used for diagnostic purposes to fix the management API port for the storage service.

13.9.2 Common Elements for Storage Plugins

In designing the Storage API and plugins, we have first of all considered that there may be a large number of use cases for data and metadata persistence, therefore we have designed a flexible architecture that poses very few limitations. In practice, this means that developers can build their own Storage plugin and they can rely on anything they want to use as persistent storage. They can use a memory structure, or even a pass-through library, a file, a message queue system, a time series database, a relational database, NoSQL or something else.

After having praised the flexibility of the Storage plugins, let’s provide guidelines about the basic functionality they should provide, bearing in mind that such functionality may not be relevant for some use cases.

- **Metadata persistency:** As mentioned before, one of the main reasons to use a Storage plugin is to safely store the configuration of the FogLAMP components. Since the configuration must survive to a system crash or reboot, it is fair to say that such information should be stored in one or more files or in a database system.
- **Data buffering:** The second most important feature of a Storage plugin is the ability to buffer (or store) data coming from the outside world, typically from the South microservices. In some cases this feature may not be necessary, since administrators may want to send data to other systems as soon as possible, using a North task of microservice. Even in situations where data can be sent up North instantaneously, you should consider these scenarios:

- FogLAMP may be installed in areas where the network is unreliable. The North plugins will provide the logic of retrying to gain connectivity and resending data when the connection has been lost in the middle of the transfer operations.
- North services may rely on the use of networks that provide time windows to operate.
- Historians and other systems may work better when data is transferred in blocks instead of a constant streaming.
- **Data purging:** Data may persist for the time needed by any specific use case, but it is pretty common that after a while (it can be seconds or minutes, but also day or months) data is no longer needed in FogLAMP. For this reason, the Storage plugin is able to purge data. Purging may be by time or by space usage, in conjunction with the fact that data may have been already transferred to other systems.
- **Data backup/restore:** Data, but especially metadata (i.e. configuration), can be backed up and stored safely on other systems. In case of crash and recovery, the same data may be restored into FogLAMP. FogLAMP provides a set of generic API to execute backup and restore operations.

13.10 Filter Plugins

Filter plugins provide a mechanism to alter the data stream as it flows through a foglamp instance, filters may be applied in south or north micro-services and may form a pipeline of multiple processing elements through which the data flows. Filters applied in a south service will only process data that is received by the south service, whilst filters placed in the north will process all data that flows out of that north interface.

Filters may;

- augment data by adding static metadata or calculated values to the data
- remove data from the stream
- add data to the stream
- modify data in the stream

It should be noted that there are some alternatives to creating a filter if you wish to make simple changes to the data stream. There are a number of existing filters that provide a degree of programmability. These include the which allows an arbitrary mathematical formula to be applied to the data or the which allows a small include Python script to be applied to the data.

Filter plugins may be written in C++ or Python and have a very simple interface. The plugin mechanism and a subset of the API is common between all types of plugins including filters.

13.10.1 Configuration

Filters use the same configuration mechanism as the rest of FogLAMP, using a JSON document to describe the configuration parameters. As with any other plugin the structure is defined by the plugin and retrieve by the `plugin_info` entry point. This is then matched with the database content to pass the configured values to the `plugin_init` entry point.

13.10.2 C++ Filter Plugin API

The filter API consists of a small number of C function entry points, these are called in a strict order and based on the same set of common API entry points for all FogLAMP plugins.

Plugin Information

The *plugin_info* entry point is the first entry point that is called in a filter plugin and returns the plugin information structure. This is the exact same call that every FogLAMP plugin must support and is used to determine the type of the plugin and the configuration category defaults for the plugin.

A typical implementation of *plugin_info* would merely return a pointer to a static `PLUGIN_INFORMATION` structure.

```
PLUGIN_INFORMATION *plugin_info()
{
    return &info;
}
```

Plugin Initialise

The *plugin_init* entry point is called after *plugin_info* has been called and before any data is passed to the filter. It is called at the phase where the service is setting up the filter pipeline and provides the filter with its configuration category that now contains the user supplied values and the destination to which the filter will send the output of the filter.

```
PLUGIN_HANDLE plugin_init(ConfigCategory* config,
                          OUTPUT_HANDLE *outHandle,
                          OUTPUT_STREAM output)
{
}
```

The *config* parameter is the configuration category with the user supplied values inserted, the *outHandle* is a handle for the next filter in the chain and the *output* is a function pointer to call to send the data to the next filter in the chain. The *outHandle* and *output* arguments should be stored for future use in the *plugin_ingest* when data is to be forwarded within the pipeline.

The *plugin_init* function returns a handle that will be passed to all subsequent plugin calls. This handle can be used to store state that needs to be passed between calls. Typically the *plugin_init* call will create a C++ class that implements the filter and return a point to the instance as the handle. The instance can then be used to store the state of the filter, including the output handle and callback that needs to be used.

Filter classes can also be used to buffer data between calls to the *plugin_ingest* entry point, allowing a filter to defer the processing of the data until it has a sufficient quantity of buffered data available to it.

Plugin Ingest

The *plugin_ingest* entry point is the workhorse of the filter, it is called with sets of readings to process and then passes on the new set of readings to the next filter in the pipeline. The process of passing on the data to the next filter is via the *OUTPUT_STREAM* function pointer. A filter does not have to output data each time it ingests data, it is free to output no data or to output more or less data than it was called with.

```
void plugin_ingest(PLUGIN_HANDLE *handle,
                  READINGSET *readingSet)
```

(continues on next page)

(continued from previous page)

```
{
}
```

The number of readings that a filter is called with will depend on the environment it is run in and what any filters earlier in the filter pipeline have produced. A filter that requires a particular sample size in order to process a result should therefore be prepared to buffer data across multiple calls to *plugin_ingest*. Several examples of filters that so this are available for reference.

The *plugin_ingest* call may send data onwards in the filter pipeline by using the stored *output* and *outHandle* parameters passed to *plugin_init*.

```
(*output)(outHandle, readings);
```

Plugin Reconfigure

As with other plugin types the filter may be reconfigured during its lifetime. When a reconfiguration operation occurs the *plugin_reconfigure* method will be called with the new configuration for the filter.

```
void plugin_reconfigure(PLUGIN_HANDLE *handle, const std::string& newConfig)
{
}
```

Plugin Shutdown

As with other plugins a shutdown call exists which may be used by the plugin to perform any cleanup that is required when the filter is shut down.

```
void plugin_shutdown(PLUGIN_HANDLE *handle)
{
}
```

C++ Helper Class

It is expected that filters will be written as C++ classes, with the plugin handle being used as a mechanism to store and pass the pointer to the instance of the filter class. In order to make it easier to write filters a base *FogLAMPFilter* class has been provided, it is recommended that you derive your specific filter class from this base class in order to simplify the implementation

```
class FogLAMPFilter {
public:
    FogLAMPFilter(const std::string& filterName,
                  ConfigCategory& filterConfig,
                  OUTPUT_HANDLE *outHandle,
                  OUTPUT_STREAM output);
    ~FogLAMPFilter() {};
    const std::string&
        getName() const { return m_name; };
    bool
        isEnabled() const { return m_enabled; };
    ConfigCategory&
        getConfig() { return m_config; };
    void
        disableFilter() { m_enabled = false; };
    void
        setConfig(const std::string& newConfig);
```

(continues on next page)

(continued from previous page)

```

public:
    OUTPUT_HANDLE*    m_data;
    OUTPUT_STREAM     m_func;
protected:
    std::string       m_name;
    ConfigCategory    m_config;
    bool              m_enabled;
};

```

13.10.3 C++ Filter Example

The following example is a simple data processing example. It applies the `log()` function to numeric data in the data stream

Plugin Interface

Most plugins written in C++ have a source file that encapsulates the C API to the plugin, this is traditionally called `plugin.cpp`. The example plugin follows this model with the content of `plugin.cpp` shown below.

The first section includes the filter class that is the actual implementation of the filter logic and defines the JSON configuration category. This uses the `QUOTE` macro in order to make the JSON definition more readable.

```

/*
 * FogLAMP "log" filter plugin.
 *
 * Copyright (c) 2020 Dianomic Systems
 *
 * Released under the Apache 2.0 Licence
 *
 * Author: Mark Riddoch
 */

#include <logFilter.h>
#include <version.h>

#define FILTER_NAME "log"
const static char *default_config = QUOTE({
    "plugin" : {
        "description" : "Log filter plugin",
        "type" : "string",
        "default" : FILTER_NAME,
        "readonly": "true"
    },
    "enable": {
        "description": "A switch that can be used to enable or_
↳disable execution of the log filter.",
        "type": "boolean",
        "displayName": "Enabled",
        "default": "false"
    },
    "match" : {
        "description" : "An optional regular expression to match in_
↳the asset name.",
        "type": "string",

```

(continues on next page)

(continued from previous page)

```

        "default": "",
        "order": "1",
        "displayName": "Asset filter"}
    });

using namespace std;

```

We then define the plugin information contents that will be returned by the *plugin_info* call.

```

/**
 * The Filter plugin interface
 */
extern "C" {

/**
 * The plugin information structure
 */
static PLUGIN_INFORMATION info = {
    FILTER_NAME,           // Name
    VERSION,               // Version
    0,                     // Flags
    PLUGIN_TYPE_FILTER,    // Type
    "1.0.0",               // Interface version
    default_config          // Default plugin configuration
};

```

The final section of this file consists of the entry points themselves and the implementation. The majority of this consist of calls to the LogFilter class that in this case implements the logic of the filter.

```

/**
 * Return the information about this plugin
 */
PLUGIN_INFORMATION *plugin_info()
{
    return &info;
}

/**
 * Initialise the plugin, called to get the plugin handle.
 * We merely create an instance of our LogFilter class
 *
 * @param config      The configuration category for the filter
 * @param outHandle   A handle that will be passed to the output stream
 * @param output      The output stream (function pointer) to which data is passed
 * @return            An opaque handle that is used in all subsequent calls to the_
 * plugin
 */
PLUGIN_HANDLE plugin_init(ConfigCategory* config,
                          OUTPUT_HANDLE *outHandle,
                          OUTPUT_STREAM output)
{
    LogFilter *log = new LogFilter(FILTER_NAME,
                                    *config,
                                    outHandle,
                                    output);
}

```

(continues on next page)

(continued from previous page)

```

        return (PLUGIN_HANDLE) log;
    }

    /**
     * Ingest a set of readings into the plugin for processing
     *
     * @param handle      The plugin handle returned from plugin_init
     * @param readingSet  The readings to process
     */
    void plugin_ingest (PLUGIN_HANDLE *handle,
                       READINGSET *readingSet)
    {
        LogFilter *log = (LogFilter *) handle;
        log->ingest (readingSet);
    }

    /**
     * Plugin reconfiguration method
     *
     * @param handle      The plugin handle
     * @param newConfig   The updated configuration
     */
    void plugin_reconfigure (PLUGIN_HANDLE *handle, const std::string& newConfig)
    {
        LogFilter *log = (LogFilter *) handle;
        log->reconfigure (newConfig);
    }

    /**
     * Call the shutdown method in the plugin
     */
    void plugin_shutdown (PLUGIN_HANDLE *handle)
    {
        LogFilter *log = (LogFilter *) handle;
        delete log;
    }

    // End of extern "C"
};

```

Filter Class

Although it is not mandatory it is good practice to encapsulate the filter logic in a class, these classes are derived from the FogLAMPFilter class

```

#ifndef _LOG_FILTER_H
#define _LOG_FILTER_H
/*
 * FogLAMP "Log" filter plugin.
 *
 * Copyright (c) 2020 Dianomic Systems
 *
 * Released under the Apache 2.0 Licence
 *
 * Author: Mark Riddoch

```

(continues on next page)

(continued from previous page)

```

*/
#include <filter.h>
#include <reading_set.h>
#include <config_category.h>
#include <string>
#include <logger.h>
#include <mutex>
#include <regex>
#include <math.h>

/**
 * Convert the incoming data to use a logarithmic scale
 */
class LogFilter : public FogLAMPFilter {
public:
    LogFilter(const std::string& filterName,
              ConfigCategory& filterConfig,
              OUTPUT_HANDLE *outHandle,
              OUTPUT_STREAM output);
    ~LogFilter();
    void ingest(READINGSET *readingSet);
    void reconfigure(const std::string& newConfig);
private:
    void handleConfig(ConfigCategory& config);
    std::string m_match;
    std::regex *m_regex;
    std::mutex m_configMutex;
};

#endif

```

Filter Class Implementation

The following is the code that implements the filter logic

```

/*
 * FogLAMP "Log" filter plugin.
 *
 * Copyright (c) 2020 Dianomic Systems
 *
 * Released under the Apache 2.0 Licence
 *
 * Author: Mark Riddoch
 */
#include <logFilter.h>

using namespace std;

/**
 * Constructor for the LogFilter.
 *
 * We call the constructor of the base class and handle the initial
 * configuration of the filter.
 */

```

(continues on next page)

(continued from previous page)

```

*
* @param   filterName       The name of the filter
* @param   filterConfig     The configuration category for this filter
* @param   outHandle        The handle of the next filter in the chain
* @param   output           A function pointer to call to output data to the next_
↪filter
*/
LogFilter::LogFilter(const std::string& filterName,
                    ConfigCategory& filterConfig,
                    OUTPUT_HANDLE *outHandle,
                    OUTPUT_STREAM output) : m_regex(NULL),
                    FogLAMPFilter(filterName, filterConfig, outHandle,
↪output)
{
    handleConfig(filterConfig);
}

/**
 * Destructor for this filter class
 */
LogFilter::~LogFilter()
{
    if (m_regex)
        delete m_regex;
}

/**
 * The actual filtering code
 *
 * @param readingSet The reading data to filter
 */
void
LogFilter::ingest(READINGSET *readingSet)
{
    lock_guard<mutex> guard(m_configMutex);

    if (isEnabled()) // Filter enable, process the readings
    {
        const vector<Reading *>& readings = ((ReadingSet *)readingSet)->
↪getAllReadings();
        for (vector<Reading *>::const_iterator elem = readings.begin();
            elem != readings.end(); ++elem)
        {
            // If we set a matching regex then compare to the name of_
↪this asset

            if (!m_match.empty())
            {
                string asset = (*elem)->getAssetName();
                if (!regex_match(asset, *m_regex))
                {
                    continue;
                }
            }

            // We are modifying this asset so put an entry in the asset_
↪tracker

            AssetTracker::getAssetTracker()->
↪addAssetTrackingTuple(getName(), (*elem)->getAssetName(), string("Filter"),
(continues on next page)

```

(continued from previous page)

```

        // Get a reading DataPoints
        const vector<Datapoint *>& dataPoints = (*elem)->
getReadingData();

        // Iterate over the datapoints
        for (vector<Datapoint *>::const_iterator it = dataPoints.
begin(); it != dataPoints.end(); ++it)
        {
            // Get the reference to a DataPointValue
            DatapointValue& value = (*it)->getData();

            /*
             * Deal with the T_INTEGER and T_FLOAT types.
             * Try to preserve the type if possible but
             * if a floating point log function is applied
             * then T_INTEGER values will turn into T_FLOAT.
             * If the value is zero we do not apply the log.
            */
            if (value.getType() == DatapointValue::T_INTEGER)
            {
                long ival = value.toInt();
                if (ival != 0)
                {
                    double newValue = log((double)ival);
                    value.setValue(newValue);
                }
            }
            else if (value.getType() == DatapointValue::T_FLOAT)
            {
                double dval = value.toDouble();
                if (dval != 0.0)
                {
                    value.setValue(log(dval));
                }
            }
            else
            {
                // do nothing for other types
            }
        }
    }

    // Pass on all readings in this case
    (*m_func)(m_data, readingSet);
}

/**
 * Reconfiguration entry point to the filter.
 *
 * This method runs holding the configMutex to prevent
 * ingest using the regex class that may be destroyed by this
 * call.
 *
 * Pass the configuration to the base FilterPlugin class and

```

(continues on next page)

(continued from previous page)

```

* then call the private method to handle the filter specific
* configuration.
*
* @param newConfig The JSON of the new configuration
*/
void
LogFilter::reconfigure(const std::string& newConfig)
{
    lock_guard<mutex> guard(m_configMutex);
    setConfig(newConfig);           // Pass the configuration to the base_
→class
    handleConfig(m_config);
}

/**
 * Handle the filter specific configuration. In this case
 * it is just the single item "match" that is a regex
 * expression
 *
 * @param config The configuration category
 */
void
LogFilter::handleConfig(ConfigCategory& config)
{
    if (config.itemExists("match"))
    {
        m_match = config.getValue("match");
        if (m_regex)
            delete m_regex;
        m_regex = new regex(m_match);
    }
}

```

13.10.4 Python Filter API

Filters may also be written in Python, the API is very similar to that of a C++ filter and consists of the same set of entry points.

Plugin Information

As with C++ filters this is the first entry point called, it returns a Python dictionary that describes the filter.

```

def plugin_info():
    """ Returns information about the plugin
    Args:
    Returns:
        dict: plugin information
    Raises:
    """

```

Plugin Initialisation

The *plugin_init* call is used to pass the resolved configuration to the plugin and also pass in the handle of the next filter in the pipeline and a callback that should be called with the output data of the filter.

```
def plugin_init(config, ingest_ref, callback):  
    """ Initialise the plugin  
    Args:  
        config: JSON configuration document for the Filter plugin configuration_  
    category  
        ingest_ref: filter ingest reference  
        callback: filter callback  
    Returns:  
        data: JSON object to be used in future calls to the plugin  
    Raises:  
    """
```

Plugin Ingestion

The *plugin_ingest* method is used to pass data into the plugin, the plugin will then process that data and call the callback that was passed into the *plugin_init* entry point with the *ingest_ref* handle and the data to send along the filter pipeline.

```
def plugin_ingest(handle, data):  
    """ Modify readings data and pass it onward  
  
    Args:  
        handle: handle returned by the plugin initialisation call  
        data: readings data  
    """
```

The *data* is arranged as an array of Python dictionaries, each of which is a *Reading*. Typically the data can be processed by traversing the array

```
for elem in data:  
    process(elem)
```

Plugin Reconfigure

The *plugin_reconfigure* entry point is called whenever a configuration change occurs for the filters configuration category.

```
def plugin_reconfigure(handle, new_config):  
    """ Reconfigures the plugin  
  
    Args:  
        handle: handle returned by the plugin initialisation call  
        new_config: JSON object representing the new configuration category for the_  
    category  
    Returns:  
        new_handle: new handle to be used in the future calls  
    """
```


Plugin Shutdown

Called when the plugin is to be shutdown to allow it to perform any cleanup operations.

```
def plugin_shutdown(handle):
    """ Shutdowns the plugin doing required cleanup.

    Args:
        handle: handle returned by the plugin initialisation call
    Returns:
        plugin shutdown
    """
```

13.10.5 Python Filter Example

The following is an example of a Python filter that calculates an exponential moving average.

```
# -*- coding: utf-8 -*-

# FogLAMP_BEGIN
# See: http://foglamp.readthedocs.io/
# FogLAMP_END

""" Module for EMA filter plugin

Generate Exponential Moving Average
The rate value (x) allows to include x% of current value
and (100-x)% of history
A datapoint called 'ema' is added to each reading being filtered
"""

import time
import copy
import logging

from foglamp.common import logger
import filter_ingest

__author__ = "Massimiliano Pinto"
__copyright__ = "Copyright (c) 2022 Dianomic Systems Inc."
__license__ = "Apache 2.0"
__version__ = "${VERSION}"

_LOGGER = logger.setup(__name__, level = logging.INFO)

PLUGIN_NAME = 'ema'

_DEFAULT_CONFIG = {
    'plugin': {
        'description': 'Exponential Moving Average filter plugin',
        'type': 'string',
        'default': PLUGIN_NAME,
        'readonly': 'true'
    },
    'enable': {
        'description': 'Enable ema plugin',
```

(continues on next page)

(continued from previous page)

```

        'type': 'boolean',
        'default': 'false',
        'displayName': 'Enabled',
        'order': "3"
    },
    'rate': {
        'description': 'Rate value: include % of current value',
        'type': 'float',
        'default': '0.07',
        'displayName': 'Rate',
        'order': "2"
    },
    'datapoint': {
        'description': 'Datapoint name for calculated ema value',
        'type': 'string',
        'default': PLUGIN_NAME,
        'displayName': 'EMA datapoint',
        'order': "1"
    }
}

def compute_ema(handle, reading):
    """ Compute EMA

    Args:
        A reading data
    """
    rate = float(handle['rate']['value'])
    for attribute in list(reading):
        if not handle['latest']:
            handle['latest'] = reading[attribute]
        handle['latest'] = reading[attribute] * rate + handle['latest'] * (1 - rate)
    reading[handle['datapoint']['value']] = handle['latest']

def plugin_info():
    """ Returns information about the plugin

    Args:
    Returns:
        dict: plugin information
    Raises:
    """
    return {
        'name': PLUGIN_NAME,
        'version': '1.9.2',
        'mode': 'none',
        'type': 'filter',
        'interface': '1.0',
        'config': _DEFAULT_CONFIG
    }

def plugin_init(config, ingest_ref, callback):
    """ Initialise the plugin

    Args:
        config: JSON configuration document for the Filter plugin configuration_

```

↪category

(continues on next page)

(continued from previous page)

```

        ingest_ref: filter ingest reference
        callback: filter callback
    Returns:
        data: JSON object to be used in future calls to the plugin
    Raises:
        """
    _config = copy.deepcopy(config)
    _config['ingestRef'] = ingest_ref
    _config['callback'] = callback
    _config['latest'] = None
    _config['shutdownInProgress'] = False
    return _config

def plugin_reconfigure(handle, new_config):
    """ Reconfigures the plugin

    Args:
        handle: handle returned by the plugin initialisation call
        new_config: JSON object representing the new configuration category for the
        ↪category
    Returns:
        new_handle: new handle to be used in the future calls
    """
    _LOGGER.info("Old config for ema plugin {} \n new config {}".format(handle, new_
    ↪config))

    new_handle = copy.deepcopy(new_config)
    new_handle['shutdownInProgress'] = False
    new_handle['latest'] = None
    new_handle['ingestRef'] = handle['ingestRef']
    new_handle['callback'] = handle['callback']
    return new_handle

def plugin_shutdown(handle):
    """ Shutdowns the plugin doing required cleanup.

    Args:
        handle: handle returned by the plugin initialisation call
    Returns:
        plugin shutdown
    """
    handle['shutdownInProgress'] = True
    time.sleep(1)
    handle['callback'] = None
    handle['ingestRef'] = None
    handle['latest'] = None

    _LOGGER.info('{} filter plugin shutdown.'.format(PLUGIN_NAME))

def plugin_ingest(handle, data):
    """ Modify readings data and pass it onward

    Args:
        handle: handle returned by the plugin initialisation call

```

(continues on next page)

(continued from previous page)

```

    data: readings data
    """
    if handle['shutdownInProgress']:
        return

    if handle['enable']['value'] == 'false':
        # Filter not enabled, just pass data onwards
        filter_ingest.filter_ingest_callback(handle['callback'], handle['ingestRef'],
↪data)
        return

    # Filter is enabled: compute EMA for each reading
    for elem in data:
        compute_ema(handle, elem['readings'])

    # Pass data onwards
    filter_ingest.filter_ingest_callback(handle['callback'], handle['ingestRef'],
↪data)

_LOGGER.debug("{} filter_ingest done.".format(PLUGIN_NAME))

```

13.11 Notification Delivery Plugins

Notification delivery plugins are used by the notification system to send a notification to some other system or device. They are the transport that allows the event to be notified to that other system or device.

Notification delivery plugins may be written in C or C++ and have a very simple interface. The plugin mechanism and a subset of the API is common between all types of plugins including filters. This documentation is based on the . The sends MQTT messages to a configurable MQTT topic when a notification is triggered and cleared.

13.11.1 Configuration

Notification Delivery plugins use the same configuration mechanism as the rest of FogLAMP, using a JSON document to describe the configuration parameters. As with any other plugin the structure is defined by the plugin and retrieved by the *plugin_info* entry point. This is then matched with the database content to pass the configured values to the *plugin_init* entry point.

13.11.2 Notification Delivery Plugin API

The notification delivery plugin API consists of a small number of C function entry points, these are called in a strict order and based on the same set of common API entry points for all FogLAMP plugins.

Plugin Information

The *plugin_info* entry point is the first entry point that is called in a notification delivery plugin and returns the plugin information structure. This is the exact same call that every FogLAMP plugin must support and is used to determine the type of the plugin and the configuration category defaults for the plugin.

A typical implementation of *plugin_info* would merely return a pointer to a static PLUGIN_INFORMATION structure.

```
PLUGIN_INFORMATION *plugin_info()
{
    return &info;
}
```

Plugin Initialise

The second call that is made to the plugin is the *plugin_init* call, that is used to retrieve a handle on the plugin instance and to configure the plugin.

```
PLUGIN_HANDLE plugin_init(ConfigCategory* config)
{
    MQTT *mqtt = new MQTT(config);
    return (PLUGIN_HANDLE)mqtt;
}
```

The *config* parameter is the configuration category with the user supplied values inserted, these values are used to configure the behavior of the plugin. In the case of our MQTT example we use this to call the constructor of our MQTT class.

```
/**
 * Construct a MQTT notification plugin
 *
 * @param category The configuration of the plugin
 */
MQTT::MQTT(ConfigCategory *category)
{
    if (category->itemExists("broker"))
        m_broker = category->getValue("broker");
    if (category->itemExists("topic"))
        m_topic = category->getValue("topic");
    if (category->itemExists("trigger_payload"))
        m_trigger = category->getValue("trigger_payload");
    if (category->itemExists("clear_payload"))
        m_clear = category->getValue("clear_payload");
}
```

This constructor merely stores values out of the configuration category as private member variables of the MQTT class.

We return the pointer to our MQTT class as the handle for the plugin. This allows subsequent calls to the plugin to reference the instance created by the *plugin_init* call.

Plugin Delivery

This is the API call made whenever the plugin needs to send a triggered or cleared notification state. It may be called multiple times within the lifetime of a plugin.

```
bool plugin_deliver(PLUGIN_HANDLE handle,
                    const std::string& deliveryName,
                    const std::string& notificationName,
                    const std::string& triggerReason,
                    const std::string& message)
{
    MQTT *mqtt = (MQTT *)handle;
    return mqtt->notify(notificationName, triggerReason, message);
}
```

The delivery call is passed the handle, which gives us the MQTT class instance on this case, the name of the notification, a trigger reason, which is a JSON document and a message. The trigger reason JSON document contains information about why the delivery call was made, including the triggered or cleared status, the timestamp of the reading that caused the notification to trigger and the name of the asset or assets involved in the notification rule that triggered this delivery event.

```
{
    "reason": "triggered",
    "asset": ["sinusoid"],
    "timestamp": "2020-11-18 11:52:33.960530+00:00"
}
```

The return from the *plugin_deliver* entry point is a boolean that indicates if the delivery succeeded or not.

In the case of our MQTT example we call the notify method of the class, this then interacts with the MQTT broker.

```
/**
 * Send a notification via MQTT broker
 *
 * @param notificationName The name of this notification
 * @param triggerReason Why the notification is being sent
 * @param message The message to send
 */
bool MQTT::notify(const string& notificationName, const string& triggerReason, const_
↳ string& message)
{
    string payload = m_trigger;
    MQTTClient client;

    lock_guard<mutex> guard(m_mutex);

    // Parse the JSON that represents the reason data
    Document doc;
    doc.Parse(triggerReason.c_str());
    if (!doc.HasParseError() && doc.HasMember("reason"))
    {
        if (!strcmp(doc["reason"].GetString(), "cleared"))
            payload = m_clear;
    }

    // Connect to the MQTT broker
    MQTTClient_connectOptions conn_opts = MQTTClient_connectOptions_initializer;
    MQTTClient_message pubmsg = MQTTClient_message_initializer;
```

(continues on next page)

(continued from previous page)

```

MQTTClient_deliveryToken token;
int rc;

if ((rc = MQTTClient_create(&client, m_broker.c_str(), CLIENTID,
MQTTCLIENT_PERSISTENCE_NONE, NULL)) != MQTTCLIENT_SUCCESS)
{
    Logger::getLogger()->error("Failed to create client, return code %d\n", rc);
    return false;
}

conn_opts.keepAliveInterval = 20;
conn_opts.cleansession = 1;
if ((rc = MQTTClient_connect(client, &conn_opts)) != MQTTCLIENT_SUCCESS)
{
    Logger::getLogger()->error("Failed to connect, return code %d\n", rc);
    return false;
}

// Construct the payload
pubmsg.payload = (void *)payload.c_str();
pubmsg.payloadlen = payload.length();
pubmsg.qos = 1;
pubmsg.retained = 0;

// Publish the message
if ((rc = MQTTClient_publishMessage(client, m_topic.c_str(), &pubmsg, &token)) != MQTTCLIENT_SUCCESS)
{
    Logger::getLogger()->error("Failed to publish message, return code %d\n", rc);
    return false;
}

// Wait for completion and disconnect
rc = MQTTClient_waitForCompletion(client, token, TIMEOUT);
if ((rc = MQTTClient_disconnect(client, 10000)) != MQTTCLIENT_SUCCESS)
    Logger::getLogger()->error("Failed to disconnect, return code %d\n", rc);
MQTTClient_destroy(&client);
return true;
}

```

Plugin Reconfigure

As with other plugin types the notification delivery plugin may be reconfigured during its lifetime. When a reconfiguration operation occurs the *plugin_reconfigure* method will be called with the new configuration for the plugin.

```

void plugin_reconfigure(PLUGIN_HANDLE *handle, const std::string& newConfig)
{
    MQTT *mqtt = (MQTT *)handle;
    mqtt->reconfigure(newConfig);
    return;
}

```

In the case of our MQTT example we call the reconfigure method of our MQTT class. In this method the new values

are copied into the local member variables of the instance.

```
/**
 * Reconfigure the MQTT delivery plugin
 *
 * @param newConfig The new configuration
 */
void MQTT::reconfigure(const string& newConfig)
{
    ConfigCategory category("new", newConfig);
    lock_guard<mutex> guard(m_mutex);
    m_broker = category.getValue("broker");
    m_topic = category.getValue("topic");
    m_trigger = category.getValue("trigger_payload");
    m_clear = category.getValue("clear_payload");
}
```

The mutex is used here to prevent the plugin reconfiguration occurring when we are delivering a notification. The same mutex is held in the notify method of the MQTT class.

Plugin Shutdown

As with other plugins a shutdown call exists which may be used by the plugin to perform any cleanup that is required when the plugin is shut down.

```
void plugin_shutdown(PLUGIN_HANDLE *handle)
{
    MQTT *mqtt = (MQTT *)handle;
    delete mqtt;
}
```

In the case of our MQTT example we merely destroy the instance of the MQTT class and allow the destructor of that class to do any cleanup that is required. In the case of this example there is no cleanup required.

13.12 Plugin Packaging

There are a set of files that must exist within the repository of a plugin that are used to create the package for that plugin on the various supported platforms. The following documents what those files are and what they should contain.

13.12.1 Common files

- **Description** - It should contain a brief description of the plugin and will be used as the description for the package that is created. Also make sure description of plugin must be in a single line as of now we do not have support multi lines yet.
- **Package** - This is the main file where we define set of variables.
 - **plugin_name** - Name of the Plugin.
 - **plugin_type** - Type of the Plugin.
 - **plugin_install_dirname** - Installed Directory name.
 - **plugin_package_name (Optional)** - Name of the Package. If it is not given then the package name should be same as plugin name.

- **requirements** - Runtime Architecture specific packages list and should have comma separated values without any space.

Note: For C-based plugins if a plugin requires some additional libraries to install with then set `additional_libs` variable inside Package file. And the value must be with following contract:

`additional_libs="DIRECTORY_NAME:FILE_NAME"` - in case of single additional
`additional_libs="DIRECTORY_NAME:FILE_NAME1,DIRECTORY_NAME:FILE_NAME2"` - in case of multiple use comma separated with both directory & file name

- **service_notification.version** - It is only required if the plugin is a notification rule or notification delivery plugin. It contains the minimum version of the notification service which the plugin requires.

13.12.2 C based Plugins

- **VERSION** - It contains the version number of the plugin and is used by the build process to include the version number within the code and also within the name of the package file created.
- **foglamp.version** - It contains the minimum version number of FogLAMP required by the plugin.
- **requirements.sh (Optional)** - It is used to install any additional libraries or other artifacts that are need to build the plugin. It takes the form of a shell script. This script, if it exists, will be run as a part of the process of building the plugin before the `cmake` command is issued in the build process.
- **extras_install.sh (Optional)** - It is a shell script that is added to the package to allow for extra commands to be executed as part of the package installation. Not all plugins will require this file to be present and it can be omitted if there are no extra steps required on the installation.

Examples of filename along with content

1. VERSION

```
$ cat VERSION
1.9.2
```

2. foglamp.version

```
$ cat foglamp.version
foglamp_version>=1.9
```

3. requirements.sh

```
$ cat requirements.sh
#!/usr/bin/env bash
which apt >/dev/null 2>&1
if [ $? -eq 0 ]; then
    sudo apt install -y libmodbus-dev
else
    which yum >/dev/null 2>&1
    if [ $? -eq 0 ]; then
        sudo yum -y install epel-release libmodbus libmodbus-devel
    fi
fi
```

4. Description

```
$ cat Description
FogLAMP modbus plugin. Supports modbus RTU and modbus TCP.
```

5. Package

```
$ cat Package
# A set of variables that define how we package this repository
#
plugin_name=modbus
plugin_type=south
plugin_install_dirname=ModbusC
plugin_package_name=foglamp-south-modbus
additional_libs="usr/local/lib:/usr/local/lib/libsmmod.so*"

# Now build up the runtime requirements list. This has 3 components
# 1. Generic packages we depend on in all architectures and package managers
# 2. Architecture specific packages we depend on
# 3. Package manager specific packages we depend on
requirements="foglamp"

case "$arch" in
    x84_64)
        ;;
    armv7l)
        ;;
    aarch64)
        ;;
esac
case "$package_manager" in
    deb)
        requirements="${requirements}, libmodbus-dev"
        ;;
esac
```

Note: If your package is not supported for a specific platform then you must exit with exitcode 1.

6. service_notification.version

```
$ cat service_notification.version
service_notification_version>=1.9.2
```

Common Additional Libraries Package

Below are the packages which created a part of the process of building FogLAMP that are commonly used in plugins.

- **foglamp-mqtt** which is a packaged version of the libpaho-mqtt library.
- **foglamp-gcp** which is a packaged version of the libjwt and libjansson libraries.
- **foglamp-iec** which is a packaged version of the IEC 60870 and IEC 61850 libraries.
- **foglamp-s2opcua** which is a packaged version of libexpat and libs2opc libraries.

If your plugin depends on any of these libraries they should be added to the *requirements* variable in the **Package** file rather than adding them as *additional_libs* since the version of these is managed by the FogLAMP build and packaging process. Below is the example

```
requirements="foglamp,foglamp-s2opcua"
```

13.12.3 Python based Plugins

- **VERSION.{PLUGIN_TYPE}.{PLUGIN_NAME}** - It contains the packaged version of the plugin and also the minimum foglamp version that the plugin requires.
- **install_notes.txt (Optional)** - It is a simple text file that can be included if there are specific instructions required to be given during the installation of the plugin. These notes will be displayed at the end of the installation process for the package.
- **extras_install.sh (Optional)** - It is a shell script that is added to the package to allow for extra commands to be executed as part of the package installation. Not all plugins will require this file to be present and it can be omitted if there are no extra steps required on the installation.
- **requirements-{PLUGIN_NAME}.txt (Optional)** - It is a simple text file that can be included if there are pip dependencies required to be given during the installation of the plugin. Also make sure file should be placed inside *python* directory.

Examples of filename along with content

1. Description

```
$ cat Description
FogLAMP South Sinusoid plugin
```

2. Package

```
$ cat Package
# A set of variables that define how we package this repository
#
plugin_name=sinusoid
plugin_type=south
plugin_install_dirname=sinusoid

# Now build up the runtime requirements list. This has 3 components
# 1. Generic packages we depend on in all architectures and package managers
# 2. Architecture specific packages we depend on
# 3. Package manager specific packages we depend on
requirements="foglamp"

case "$arch" in
    x86_64)
        ;;
    armv7l)
        ;;
    aarch64)
        ;;
esac
case "$package_manager" in
    deb)
        ;;
esac
```

Note: If your package is not supported for a specific platform then you must exit with exitcode 1.

3. VERSION.{PLUGIN_TYPE}.{PLUGIN_NAME}

```
$ cat VERSION.south.sinusoid
foglamp_south_sinusoid_version=1.9.2
foglamp_version>=1.9
```

4. install_notes.txt

```
$ cat install_notes.txt
It is required to reboot the RPi, please do the following steps:
1) sudo reboot
```

5. extras_install.sh

```
#!/usr/bin/env bash

os_name=$(grep -o '^NAME=.*' /etc/os-release | cut -f2 -d\" | sed 's/\"//g')
os_version=$(grep -o '^VERSION_ID=.*' /etc/os-release | cut -f2 -d\" | sed 's/\"//g')
echo "Platform is ${os_name}, Version: ${os_version}"
arch=`arch`
ID=$(cat /etc/os-release | grep -w ID | cut -f2 -d"=")
if [ ${ID} != "mendel" ]; then
case $os_name in
  *Ubuntu*)
    if [ ${arch} = "aarch64" ]; then
      python3 -m pip install --upgrade pip
    fi
    ;;
  esac
fi
```

6. requirements-{PLUGIN_NAME}.txt

```
$ cat python/requirements-modbuscp.txt
pymodbus3==1.0.0
```

13.12.4 Building A Package

Firstly you need to clone the repository [foglamp-pkg](#). Now do the following steps

```
$ cd plugins
$ ./make_deb -b <BRANCH_NAME> <REPOSITORY_NAME>

if everything goes well with above command then you can find your package inside_
↪archive directory.

$ ls archive
```

13.13 Testing Your Plugin

The first step in testing your new plugin is to put the plugin in the location in which your FogLAMP system will be loading it from. The exact location depends on the way your installed you FogLAMP system and the type of plugin.

If your FogLAMP system was installed from a package and you used the default installation path, then your plugin must be stored under the directory `/usr/local/foglamp`. If you installed FogLAMP in a nonstandard location or you have built it from the source code, then the plugin should be stored under the directory `$FOGLAMP_ROOT`.

A C/C++ plugin or a hybrid plugin should be placed in the directory `plugins/<type>/<plugin name>` under the installed directory described above. Where `<type>` is one of `south`, `filter`, `north`, `notificationRule` or `notificationDelivery`. And `<plugin name>` is the name you gave your plugin.

A south plugin written in C/C++ and called DHT11, for a system installed from a package, would be installed in a directory called `/usr/local/foglamp/plugins/south/DHT11`. Within that directory FogLAMP would expect to find a file called `libDHT11.so`.

A south hybrid plugin called MD1421, for a development system built from source would be installed in `$(FOGLAMP_ROOT)/plugins/south/MD1421`. In this directory a JSON file called `MD1421.json` should exist, this is what the system will read to create the plugin.

A Python plugin should be installed in the directory `python/foglamp/plugins/<plugin type>/<plugin name>` under the installed directory described above. Where `<type>` is one of `south`, `filter`, `north`, `notificationRule` or `notificationDelivery`. And `<plugin name>` is the name you gave your plugin.

A Python filter plugin called normalise, on a system installed from a package in the default location should be copied into a directory `/usr/local/foglamp/python/foglamp/plugins/filter/normalise`. Within this directory should be a file called `normalise.py` and an empty file called `__init__.py`.

13.13.1 Initial Testing

After you have copied your plugin into the correct location you can test if FogLAMP is able to see it by running the API call `/foglamp/plugins/installed`. This will list all the installed plugins and their versions.

```
$ curl http://localhost:8081/foglamp/plugins/installed | jq
{
  "plugins": [
    {
      "name": "http_north",
      "type": "north",
      "description": "HTTP North Plugin",
      "version": "1.8.1",
      "installedDirectory": "north/http_north",
      "packageName": "foglamp-north-http-north"
    },
    {
      "name": "GCP",
      "type": "north",
      "description": "Google Cloud Platform IoT-Core",
      "version": "1.8.1",
      "installedDirectory": "north/GCP",
      "packageName": "foglamp-north-gcp"
    },
    ...
  ]
}
```

Note, in the above example the *jq* program has been used to format the returned JSON and the output has been truncated for brevity.

If your plugin does not appear it may be because there was a problem loading it or because the *plugin_info* call returned a bad value. Examine the syslog file to see if there are any errors recorded during the above API call.

13.13.2 C/C++ Common Faults

Common faults for C/C++ plugins are that a symbol could not be resolved when the plugin was loaded or the JSON for the default configuration is malformed.

There is a utility called *get_plugin_info* that is used by Python code to call the *C plugin_info* call, this can be used to ascertain the cause of some problems. It should return the default configuration of your plugin and will verify that your plugin has no undefined symbols.

The location of *get_plugin_info* will depend on the type of installation you have. If you have built from source then it can be found in *./cmake_build/C/plugins/utis/get_plugin_info*. If you have installed a package, or run *make install*, you can find it in */usr/local/foglamp/extras/C/get_plugin_info*.

The utility is passed the library file of your plugin as its first argument and the function to call, usually *plugin_info*.

```
$ get_plugin_info plugins/north/GCP/libGCP.so plugin_info
{"name": "GCP", "version": "1.8.1", "type": "north", "interface": "1.0.0", "flag": 0,
  "config": { "plugin": { "description": "Google Cloud Platform IoT-Core", "type":
    "string", "default": "GCP", "readonly": "true" }, "project_id": { "description":
    "The GCP IoT Core Project ID", "type": "string", "default": "", "order": "1",
    "displayName": "Project ID" }, "region": { "description": "The GCP Region", "type":
    "enumeration", "options": [ "us-centrall", "europe-west1", "asia-east1" ],
    "default": "us-centrall", "order": "2", "displayName": "The GCP Region" },
    "registry_id": { "description": "The Registry ID of the GCP Project", "type":
    "string", "default": "", "order": "3", "displayName": "Registry ID" }, "device_id":
    { "description": "Device ID within GCP IoT Core", "type": "string", "default":
    "", "order": "4", "displayName": "Device ID" }, "key": { "description": "Name
    of the key file to use", "type": "string", "default": "", "order": "5",
    "displayName": "Key Name" }, "algorithm": { "description": "JWT algorithm", "type":
    "enumeration", "options": [ "ES256", "RS256" ], "default": "RS256", "order":
    "6", "displayName": "JWT Algorithm" }, "source": { "description": "The source of
    data to send", "type": "enumeration", "default": "readings", "order": "8",
    "displayName": "Data Source", "options": [ "readings", "statistics" ] } } }
```

If there is an undefined symbol you will get an error from this utility. You can also check the validity of your JSON configuration by piping the output to a program such as *jq*.

```
$ get_plugin_info plugins/south/Random/libRandom.so plugin_info | jq
{
  "name": "Random",
  "version": "1.9.2",
  "type": "south",
  "interface": "1.0.0",
  "flag": 4096,
  "config": {
    "plugin": {
      "description": "Random data generation plugin",
      "type": "string",
      "default": "Random",
      "readonly": "true"
    }
  },
}
```

(continues on next page)

(continued from previous page)

```

    "asset": {
      "description": "Asset name",
      "type": "string",
      "default": "Random",
      "displayName": "Asset Name",
      "mandatory": "true"
    }
  }
}

```

13.13.3 Running Under a Debugger

If you have a C/C++ plugin that crashes you may want to run the plugin under a debugger. To build with debug symbols use the CMake option `-DCMAKE_BUILD_TYPE=Debug` when you create the *Makefile*.

Running a Service Under the Debugger

```
$ cmake -DCMAKE_BUILD_TYPE=Debug ..
```

The easiest approach to run under a debugger is

- Create the service that uses your plugin, say a south service and name that service as you normally would.
- Disable that service from being started by FogLAMP
- Use the foglamp status script to find the arguments to pass the service

```

$ scripts/foglamp status
FogLAMP v1.8.2 running.
FogLAMP Uptime: 1451 seconds.
FogLAMP records: 200889 read, 200740 sent, 120962 purged.
FogLAMP does not require authentication.
=== FogLAMP services:
foglamp.services.core
foglamp.services.storage --address=0.0.0.0 --port=39821
foglamp.services.south --port=39821 --address=127.0.0.1 --name=AX8
foglamp.services.south --port=39821 --address=127.0.0.1 --name=Sine
=== FogLAMP tasks:

```

- Note the `--port=` and `--address=` arguments
- Set your `LD_LIBRARY_PATH`. This is normally done in the script that launches FogLAMP but will need to be run as a manual step when running under the debugger.

```
export LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:/usr/local/foglamp/lib
```

If you built from source rather than installing a package you will need to include the libraries you built

```
export LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:${FOGLAMP_ROOT}/cmake_
↳ build/C/lib
```

- Get a startup token by calling the FogLAMP API endpoint

Note: the caller must be authenticated as the *admin* user using either the username and password authentication or the certificate authentication mechanism in order to call the API endpoint. You must first set FogLAMP to require authentication. To do this, launch the FogLAMP GUI, navigate to Configuration and then Admin API. Set Authentication to *mandatory*. Authentication Method can be left as *any*.

In order to authenticate as the *admin* user one of the two following methods should be used, the choice of which is dependant on the authentication mechanism configured in your FogLAMP installation.

– User and password login

```
curl -d '{"username": "admin", "some_pass": "foglamp"}' -X POST http://localhost:8081/foglamp/login
```

Successful authentication will produce a response as shown below.

```
{ "message": "Logged in successfully", "uid": 1, "token":
  ↪ "eyJ0eXAiOiJKV1QiLCJhbGciOiJIUzI1NiJ9.
  ↪ eyJ1aWQiOiJEsImV4cCI6MTY1NDU5NTIyMn0.1lhIgQ93LbCP-
  ↪ ztGlIuJVd6AJrBlbNBvCv7SeuMfAs", "admin": true }
```

– Certificate login

```
curl -T /some_path/admin.cert -X POST http://
  ↪ localhost:8081/foglamp/login
```

Successful authentication will produce a response as shown below.

```
{ "message": "Logged in successfully", "uid": 1, "token":
  ↪ "eyJ0eXAiOiJKV1QiLCJhbGciOiJIUzI1NiJ9.
  ↪ eyJ1aWQiOiJEsImV4cCI6MTY1NDU5NTkzN30.6VVD_
  ↪ 5RwmpLga2A7ri2bXhlo3x_CLqOYiefAAmLP63Y", "admin": true }
```

It is now possible to call the API endpoint to retrieve a startup token by passing the authentication token given in the authentication request.

```
curl -X POST 127.0.0.1:8081/foglamp/service/ServiceName/otp -H
  ↪ 'authorization: Token'
```

Where **ServiceName** is the name you gave your service when you created it and **Token** received by the authentication request above.

This call will respond with a startup token that can be used to start the service you are debugging. An example response is shown below.

```
.. code-block:: console
```

```
{ "startupToken": "WvFTYeGUvSEFMndePGbyvOsVYUzbnJdi" }
```

```
*startupToken* will be passed as service start argument: --
  ↪ token=*startupToken*
```

- Load the service you wish to use to run your plugin, e.g. a south service, under the debugger. This should be run from the FOGAMP_ROOT directory


```
$ cd $FOGLAMP_ROOT
$ gdb services/foglamp.services.south
```

- Run the service passing the `--port=` and `--address=` arguments you noted above and add `-d` and `--name=` with the name of your service and `--token=startupToken`

```
(gdb) run --port=39821 --address=127.0.0.1 --name=ServiceName -d -
↳ --token=startupToken
```

Where *ServiceName* is the name you gave your service when you created it and *startupToken* is the token issued using the method described above. Note, this token may only be used once, each time the service is restarted using the debugger a new startup token must be obtained.

- You can now use the debugger in the way you normally would to find any issues.

Note: At this stage the plugins have not been loaded into the address space. If you try to set a break point in the plugin code you will get a warning that the break point can not currently be set. However when the plugin is later loaded the break point will be set and behave as expected.

Only the plugin has been built with debug information, if you wish to be able to single step into the library code that supports the plugin, and the services you must rebuild FogLAMP itself with debug symbols. There are multiple ways this can be done, but perhaps the simplest approach is to modify the *Makefile* in the route of the FogLAMP source.

When building FogLAMP the *cmake* command is executed by the make process, hence rather than manually running *cmake* and rebuilding you can simple alter the line

```
CMAKE := cmake
```

in the *Makefile* to read

```
CMAKE := cmake -DCMAKE_BUILD_TYPE=Debug
```

After making this change you should run a *make clean* followed by a *make* command

```
$ make clean
$ make
```

One side effect of this, caused by running *make clean* is that the plugins you have previously built have been removed from the `$FOGLAMP_ROOT/plugins` directory and this must be rebuilt.

Alternatively you can manually build a debug version by running the following commands

```
$ cd $FOGLAMP_ROOT/cmake_build
$ cmake -DCMAKE_BUILD_TYPE=Debug ..
$ make
```

This has the advantage that *make clean* is not run so your plugins will be preserved.

Running a Task Under the Debugger

Running a task under the debugger is much the same as running a service, you will first need to find the management port and address of the core management service. Create the task, e.g. a north sending process in the same way as you normally would and disable it. You will also need to set your `LD_LIBRARY_PATH` as with running a service under the debugger.

If you are using a plugin with a task, such as the north sending process task, then the command to use to start the debugger is

```
$ gdb tasks/sending_process
```

Running the Storage Service Under the Debugger

Running the storage service under the debugger is more difficult as you can not start the storage service after FogLAMP has started, the startup of the storage service is coordinated by the core due to the nature of how configuration is stored. It is possible however to attach a debugger to a running storage service.

- Run a command to find the process ID of the storage service

```
$ ps aux | grep foglamp.services.storage
foglamp 23318 0.0 0.3 270848 12388 ? Ssl 10:00 0:01 /usr/
↳local/foglamp/services/foglamp.services.storage --address=0.0.0.0 --
↳port=33761
foglamp 31033 0.0 0.0 13136 1084 pts/1 S+ 10:37 0:00 grep --
↳color=auto foglamp.services.storage
```

- Use the process ID of the foglamp service as an argument to gdb. Note you will need to run gdb as root on some systems

```
$ sudo gdb /usr/local/foglamp/services/foglamp.services.storage 23318
GNU gdb (Ubuntu 8.1-0ubuntu3) 8.1.0.20180409-git
Copyright (C) 2018 Free Software Foundation, Inc.
License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/
↳gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law. Type "show_
↳copying"
and "show warranty" for details.
This GDB was configured as "x86_64-linux-gnu".
Type "show configuration" for configuration details.
For bug reporting instructions, please see:
<http://www.gnu.org/software/gdb/bugs/>.
Find the GDB manual and other documentation resources online at:
<http://www.gnu.org/software/gdb/documentation/>.
For help, type "help".
Type "apropos word" to search for commands related to "word"...
Reading symbols from services/foglamp.services.storage...done.
Attaching to program: /usr/local/foglamp/services/foglamp.services.
↳storage, process 23318
[New LWP 23320]
[New LWP 23321]
[New LWP 23322]
[New LWP 23330]
[Thread debugging using libthread_db enabled]
Using host libthread_db library "/lib/x86_64-linux-gnu/libthread_db.
↳so.1".
```

(continues on next page)

(continued from previous page)

```
0x00007f47a3e05d2d in __GI___pthread_timedjoin_ex_
↳ (threadid=139945627997952, thread_return=0x0, abstime=0x0,
    block=<optimized out>) at pthread_join_common.c:89
89   pthread_join_common.c: No such file or directory.
(gdb)
```

- You can now use gdb to set break points etc and debug the storage service and plugins.

If you are debugging a plugin that crashes the system when readings are processed you should disable the south services until you have connected the debugger to the storage system. If you have a system that is setup and crashes, use the `--safe-mode` flag to the startup of FogLAMP in order to disable all processes and services. This will allow you to disable services or to run a particular service manually.

13.13.4 Using strace

You can also use a similar approach to that of running gdb to use the *strace* command to trace system calls and signals

- Create the service that uses your plugin, say a south service and name that service as you normally would.
- Disable that service from being started by FogLAMP
- Use the foglamp status script to find the arguments to pass the service

```
$ scripts/foglamp status
FogLAMP v1.8.2 running.
FogLAMP Uptime: 1451 seconds.
FogLAMP records: 200889 read, 200740 sent, 120962 purged.
FogLAMP does not require authentication.
=== FogLAMP services:
foglamp.services.core
foglamp.services.storage --address=0.0.0.0 --port=39821
foglamp.services.south --port=39821 --address=127.0.0.1 --name=AX8
foglamp.services.south --port=39821 --address=127.0.0.1 --name=Sine
=== FogLAMP tasks:
```

- Note the `--port=` and `--address=` arguments
- Run *strace* with the service adding the same set of arguments you used in gdb when running the service

```
$ strace services/foglamp.services.south --port=39821 --address=127.0.0.1
↳ --name=ServiceName --token=StartupToken -d
```

Where *ServiceName* is the name you gave your service and *startupToken* as issued following above steps.

13.13.5 Memory Leaks and Corruptions

The same approach can be used to make use of the *valgrind* command to find memory corruption and leak issues in your plugin

- Create the service that uses your plugin, say a south service and name that service as you normally would.
- Disable that service from being started by FogLAMP
- Use the foglamp status script to find the arguments to pass the service

```
$ scripts/foglamp status
FogLAMP v1.8.2 running.
FogLAMP Uptime: 1451 seconds.
FogLAMP records: 200889 read, 200740 sent, 120962 purged.
FogLAMP does not require authentication.
=== FogLAMP services:
foglamp.services.core
foglamp.services.storage --address=0.0.0.0 --port=39821
foglamp.services.south --port=39821 --address=127.0.0.1 --name=AX8
foglamp.services.south --port=39821 --address=127.0.0.1 --name=Sine
=== FogLAMP tasks:
```

- Note the *--port=* and *--address=* arguments
- Run *valgrind* with the service adding the same set of arguments you used in *gdb* when running the service.

Add any arguments you wish to pass to *valgrind* itself before the service executable name, in this case we are passing *--leak-check=full*.

```
$ valgrind --leak-check=full services/foglamp.services.south --
↪port=39821 --address=127.0.0.1 --name=ServiceName --token=StartupToken_
↪-d
```

Where *ServiceName* is the name you gave your service and *startupToken* is a one time use token obtained following the steps shown above.

- Once the service has run for a while shut it down to trigger *valgrind* to print a summary of memory leaks found during the execution.

13.13.6 Python Plugin Info

It is also possible to test the loading and validity of the *plugin_info* call in a Python plugin.

- From the */usr/include/foglamp* or *\${FOGLAMP_ROOT}* directory run the command

```
python3 -c 'from foglamp.plugins.south.<name>.<name> import plugin_info;
↪print(plugin_info())'
```

Where *<name>* is the name of your plugin.

```
python3 -c 'from foglamp.plugins.south.sinusoid.sinusoid import plugin_info;
↪print(plugin_info())'
{'name': 'Sinusoid Poll plugin', 'version': '1.8.1', 'mode': 'poll', 'type':
↪'south', 'interface': '1.0', 'config': {'plugin': {'description': 'Sinusoid_
↪Poll Plugin which implements sine wave with data points', 'type': 'string',
↪'default': 'sinusoid', 'readonly': 'true'}, 'assetName': {'description': 'Name_
↪of Asset', 'type': 'string', 'default': 'sinusoid', 'displayName': '(continues on next page)
↪'mandatory': 'true'}}}
```

(continued from previous page)

This allows you to confirm the plugin can be loaded and the *plugin_info* entry point can be called.

You can also check your default configuration. Although in Python this is usually harder to get wrong.

```
$ python3 -c 'from foglamp.plugins.south.sinusoid.sinusoid import plugin_info;
↳ print(plugin_info()["config"])'
{'plugin': {'description': 'Sinusoid Poll Plugin which implements sine wave with data',
↳ points', 'type': 'string', 'default': 'sinusoid', 'readonly': 'true'}, 'assetName':
↳ {'description': 'Name of Asset', 'type': 'string', 'default': 'sinusoid',
↳ 'displayName': 'Asset name', 'mandatory': 'true'}}
```

13.14 Developing with Windows Subsystem for Linux (WSL2)

Windows Subsystem for Linux (WSL2) allows you to run a Linux environment directly on Windows without the overhead of Hyper-V on Windows 10 or a dual-boot setup. You can run many Linux command-line tools, utilities, and applications on a special lightweight virtual machine running on Windows. It is possible to run a complete FogLAMP system on WSL2. This includes the FogLAMP GUI which can be accessed from a browser running on the host Windows environment.

Microsoft's Visual Studio Code is a cross-platform editor that supports extensions for building and debugging software in a variety of languages and environments. This article describes how to configure Visual Studio Code to edit, build and debug FogLAMP plugins written in C++ running in Linux under WSL2.

Note: It is possible to configure Visual Studio Code to build and test Python code in WSL2 but this is not covered in this article.

13.14.1 Preparing the Development Environment

This section outlines the steps to configure WSL2 and the Linux environment.

Installing Windows Subsystem for Linux (WSL2)

You must be running Windows 10 version 2004 and higher (Build 19041 and higher) or Windows 11 to install WSL2. The easiest way to install is to open a Windows Command Prompt as Administrator and run this command:

```
wsl --install
```

Windows will perform all the necessary steps for you. It will install the default Linux distribution which is the latest version of Ubuntu. If you wish to perform the steps manually, or install a Linux distribution other than the default, see the Microsoft documentation on [Installing WSL](#).

When the installation completes, the Linux distribution will launch in a new window. It will prompt you for a username to serve as the root account and password. This username has nothing to do with your Windows environment so it can be any name you choose.

You can start the Linux distribution at any time by finding it in the Windows Start Menu. If you hit the Windows key and type the name of your Linux distribution (default: "Ubuntu"), you should see it immediately.

Some Useful Features of WSL2

A Linux distribution running in WSL2 is a lightweight virtual machine but is well integrated with the Windows environment. Here are some useful features:

- *Cut and paste text into and out of the Linux window:* The Linux window behaves just like a Command Prompt window or a Powershell window. You can copy text from any window and paste it into any other.
- *Access the Linux file system from Windows:* The Linux file system appears as a Network drive in Windows. Open the Windows File Explorer and navigate to “\\wsl\$.” You will see your Linux distributions appear as network folders.
- *Access the Windows file system from Linux:* From the *bash* command line, navigate to the mount point “/mnt.” You will see your Windows drive letters in this directory.
- *Access the Linux environment from the Windows host through the network:* From the *bash* command line, run the command *hostname -I*. The external IP address returned by this command can be used in the Windows host to reach Linux.
- *Access the Windows host from the Linux environment through the network:* From the *bash* command line, run the command *cat /etc/resolv.conf*. The IP address after the label *nameserver* can be used in the Linux environment to reach the Windows host.

Preparing the Linux Distribution for FogLAMP

The *systemd* service manager is not configured by default in an Ubuntu distribution running in WSL2. Since FogLAMP relies on *systemd*, you must run a script to enable it. From your home directory in the Ubuntu window, enter the commands:

```
git clone https://github.com/DamionGans/ubuntu-wsl2-systemd-script.git
cd ubuntu-wsl2-systemd-script
bash ubuntu-wsl2-systemd-script.sh
```

Restart the Ubuntu distribution using *sudo reboot* or *sudo systemctl reboot*. When the distribution has restarted, run the command *systemctl*. You should see no error and a list of units. The script must be run *one time only*. Whenever you start up your Ubuntu distribution, *systemd* should be ready.

Installing FogLAMP

Following the normal instructions for [Installing FogLAMP on Ubuntu](#). Make sure the package repository matches your version of Ubuntu. You can check the operating system version in your distribution with the command *hostnamectl* or *cat /etc/os-release*.

Installing Visual Studio Code

Navigate to the [Visual Studio Code](#) webpage in your Windows browser. Click the *Download for Windows* button. Run the installer to install Visual Studio Code.

Visual Studio Code is available for Microsoft Windows, Apple MacOS, and several Linux distributions. **Do not install the Linux build of Visual Studio Code in your Linux distribution in WSL2.** You will actually be launching Visual Studio Code for Windows from your Linux distribution!

13.14.2 Starting the Linux Distribution

Perform these steps every time you start your Linux distribution if you plan to run FogLAMP:

Starting syslog

The system log `/var/log/syslog` is not configured to run automatically in a Linux distribution in WSL2. Start *syslog* with the command:

```
sudo service rsyslog start
```

You must do this at every startup.

Starting Nginx

FogLAMP uses [Nginx](#) as a web server to host the FogLAMP GUI. If you plan to run FogLAMP GUI during your Linux distribution session, enter the command:

```
sudo service nginx start
```

You must do this at every startup if you plan to run the FogLAMP GUI.

Starting FogLAMP

Start FogLAMP normally. You can start it from the normal run directory, or from your build directory by following the directions on the webpage [Testing Your Plugin](#).

Starting FogLAMP GUI

If *Nginx* is running, you can run the FogLAMP GUI in a browser in your host Windows environment. Find the external IP address for your Linux distribution using the command:

```
hostname -I
```

This address is reachable from your Windows environment. Copy the IP address to a new tab in your browser and hit Enter. You should see the FogLAMP GUI Dashboard page.

Note: The Linux distribution's external IP address is (usually) different every time you start it. You will need to run the *hostname -I* command every time to obtain the current IP address.

13.14.3 Configuring Visual Studio Code

This section describes how to configure Visual Studio Code to edit, build and debug your C++ Linux projects. These instructions are summarized from the Visual Studio Code tutorial [Using C++ and WSL in VS Code](#).

Installing Extensions

Navigate to a directory containing your C++ source code files and issue the command:

```
code .
```

This will launch Visual Studio Code in your Windows environment but it will be looking at the current directory in your Linux distribution. Since you are launching Visual Studio Code from your Linux distribution, Code should prompt you to install two Extensions:

- [Remote-WSL](#)
- [C/C++](#)

If you are not prompted, follow these links to install the extensions and restart Visual Studio Code. If the extensions are installed and working, you should see a green label in the lower left-hand corner of the Visual Studio Code window with the text *WSL:* followed by the name of your Linux distribution.

Configuring your Workspace

Visual Studio Code refers to your directory of source code files as the *Workspace*. In order to edit, build and debug your code, you must create 3 Json files in a Workspace subdirectory called *.vscode*:

- **c_cpp_properties.json**: compiler path, IntelliSense settings, and include file paths
- **tasks.json**: build instructions
- **launch.json**: debugger settings

You can create these files manually or use Visual Studio Code's configuration wizards. These subsections describe creation and required contents of each of these three files.

Code Editor Configuration: c_cpp_properties.json

- Open the Command Palette using the key sequence *Ctrl+Shift+P*.
- Choose the command *C/C++: Edit Configurations (JSON)*.
- This will create the *.vscode* subdirectory (if it doesn't already exist) and the *c_cpp_properties.json* file.
- This Json file will be opened for editing.
- You will see a new array called *configurations* with a single configuration object defined.
- This configuration will have a string array called *includePath*.
- Add the paths to your own include files, and those required by the FogLAMP API to the *includePath* array.
- You can use Linux environment variables in your paths. For example:

```
"${FOGLAMP_ROOT}/C/common/include"
```

- You can find the list of include files by running your *make* command:

```
make --just-print
```

which will list all commands defined by *make* without executing them. You will see the include file list in every instance of the *gcc* compiler command.

Build Configuration: tasks.json

- From the Visual Studio Code main menu, choose *Terminal -> Configure Default Build Task*.
- A dropdown will display of available tasks for C++ projects.
- Choose *g++ build active file*.
- This will create the `.vscode` subdirectory (if it doesn't already exist) and the `tasks.json` file.
- Open the Json file for editing.

Building the project will be done using the `make` file rather than the `gcc` compiler. To make this change, edit the `command` and `args` entries as follows:

```
"command": "make",
"args": [
  "-C",
  "${workspaceFolder}/build"
],
```

The “-C” argument for `make` will move into the specified directory before doing anything.

You can invoke a build from Visual Studio Code at any time with the key sequence `Ctrl+Shift+B`.

Debugger Configuration: launch.json

- From the Visual Studio Code main menu, choose *Run -> Add Configuration...*
- Choose *C++ (GDB/LLDB)*.
- This will create the `.vscode` subdirectory (if it doesn't already exist) and the `launch.json` file.
- Edit the `launch.json` file so it looks like this:

```
{
  "version": "0.2.0",
  "configurations": [
    {
      "name": "Debug Plugin",
      "type": "cppdbg",
      "request": "launch",
      "targetArchitecture": "x86_64",
      "cwd": "${fileDirname}",
      "program": "/full/path/to/foglamp.services.north",
      "externalConsole": false,
      "stopAtEntry": true,
      "MIMode": "gdb",
      "avoidWindowsConsoleRedirection": false,
      "args": [
        "--port=42467",
        "--address=0.0.0.0",
        "--name=MyPluginInstance",
        "-d"
      ]
    }
  ]
}
```

Note:

- The *program* attribute holds the program that the *gdb* debugger should launch. For FogLAMP plugin development, this is either *foglamp.services.north* or *foglamp.services.south* depending on which one you are building. These service executables will dynamically load your plugin library when they run.
 - The *args* attribute has the arguments normally passed to the service executable. Since the TCP/IP *port* changes every time FogLAMP starts up, you must edit this file to update the *port* number before starting your debug session.
-

Start your debug session from the Visual Studio Code main menu. Choose *Run -> Start Debugging* or by hitting the F5 key.

13.14.4 Known Problems

- *Environment variables in launch.json*: Support for environment variables in the *program* attribute is inconsistent. Variables created by Visual Studio Code itself will work but user-defined environment variables like `FOGLAMP_ROOT` will not.
- *gdb startup errors*: It can occur that *gdb* stops with error 42 and exits immediately when you start a debugging session. To fix this, shut down your Linux distributions and reinstall Visual Studio Code in Windows. You will not lose your configuration settings or your installed extensions.
- *Inconsistent breakpoint lists*: Visual Studio Code shows a list of breakpoints in the lower left corner of the window. The *gdb* debugger maintains its own list of breakpoints. It can occur that the two lists fall out of sync. You can still create, view and delete breakpoints from the *Debug Console* tab at the bottom of the screen which gives you access to the *gdb* command line. When using the *Debug Console*, you must precede all *gdb* commands with “-exec.”

To manipulate breakpoints:

- Set a breakpoint: *-exec b functionName*.
- View breakpoints: *-exec info b*. This will display an ordinal number for each breakpoint.
- Delete breakpoints: *-exec del ##*. Use the original number returned by *-exec info b* as “##.”

13.14.5 References

- [Visual Studio Code](#)
- [Using C++ and WSL in VS Code](#)
- [Remote development in WSL](#)
- [Debug C++ in Visual Studio Code](#)
- [Predefined Variables Reference](#)
- [C_cpp_properties.json reference](#)
- [Schema for tasks.json](#)
- [Configuring C/C++ Debugging \(launch.json\)](#)

REST API DEVELOPERS GUIDE

14.1 The FogLAMP REST API

Users, administrators and applications interact with FogLAMP via a REST API. This section presents a full reference of the API.

Note: The FogLAMP REST API should not be confused with the internal REST API used by FogLAMP tasks and microservices to communicate with each other.

14.1.1 Introducing the FogLAMP REST API

The REST API is the route into the FogLAMP appliance, it provides all user and program interaction to configure, monitor and manage the FogLAMP system. A separate specification will define the contents of the API, in summary however it is designed to allow for:

- The complete configuration of the FogLAMP appliance
- Access to monitoring statistics for the FogLAMP appliance
- User and role management for access to the API
- Access to the data buffer contents

Port Usage

In general FogLAMP components use dynamic port allocation to determine which port to use, the admin API is however an exception to this rule. The Admin API port has to be known to end-users and any user interface or management system that uses it, therefore the port on which the admin API listens must be consistent and fixed between invocations. This does not mean however that it can not be changed by the user. The user must have the option to define the port to use by the admin API to listen on. To achieve this the port will be stored in the configuration data for the admin API, using the configuration category *AdminAPI*, see Configuration. Administrators who have access to the appliance can find information regarding the port and the protocol to used (i.e. HTTP or HTTPS) in the *pid* file stored in *\$FOGLAMP_DATA/var/run/*:

```
$ cat data/var/run/foglamp.core.pid
{ "adminAPI" : { "protocol" : "HTTP",
                 "port"      : 8081,
                 "addresses" : [ "0.0.0.0" ] },
  "processID" : 3585 }
$
```

FogLAMP is shipped with a default port for the admin API to use, however the user is free to change this after installation. This can be done by first connecting to the port defined as the default and then modifying the port using the admin API. FogLAMP should then be restarted to make use of this new port.

Infrastructure

There are two REST API's that allow external access to FogLAMP, the **Administration API** and the **User API**. The User API is intended to allow access to the data in the FogLAMP storage layer which buffers sensor readings, and it is not part of this current version.

The Administration API is the first API is concerned with all aspects of managing and monitoring the FogLAMP appliance. This API is used for all configuration operations that occur beyond basic installation.

14.2 Administration API Reference

This section presents the list of administrative API methods in alphabetical order.

14.2.1 Audit Trail

The audit trail API is used to interact with the audit trail log tables in the storage microservice. In FogLAMP, log information is stored in the system log where the microservice is hosted. All the relevant information used for auditing are instead stored inside FogLAMP and they are accessible through the Admin REST API. The API allows the reading but also the addition of extra audit logs, as if such logs are created within the system.

audit

The *audit* methods implement the audit trail, they are used to create and retrieve audit logs.

GET Audit Entries

GET /foglamp/audit - return a list of audit trail entries sorted with most recent first.

Request Parameters

- **limit** - limit the number of audit entries returned to the number specified
- **skip** - skip the first n entries in the audit table, used with limit to implement paged interfaces
- **source** - filter the audit entries to be only those from the specified source
- **severity** - filter the audit entries to only those of the specified severity

Response Payload

The response payload is an array of JSON objects with the audit trail entries.

Name	Type	Description	Example
timestamp	timestamp	The timestamp when the audit trail item was written.	2018-04-16 14:33:18.215
source	string	The source of the audit trail entry.	CoAP
severity	string	The severity of the event that triggered the audit trail entry to be written. This will be one of SUCCESS, FAILURE, WARNING or INFORMATION.	FAILURE
details	object	A JSON object that describes the detail of the audit trail event.	{ "message": "Sensor readings discarded due to malformed payload" }

Example

```
$ curl -s http://localhost:8081/foglamp/audit?limit=2
{ "totalCount" : 24,
  "audit"      : [ { "timestamp" : "2018-02-25 18:58:07.748",
                    "source"    : "SRVRG",
                    "details"   : { "name" : "COAP" },
                    "severity"  : "INFORMATION" },
                  { "timestamp" : "2018-02-25 18:58:07.742",
                    "source"    : "SRVRG",
                    "details"   : { "name" : "HTTP_SOUTH" },
                    "severity"  : "INFORMATION" },
                  { "timestamp" : "2018-02-25 18:58:07.390",
                    "source"    : "START",
                    "details"   : {},
                    "severity"  : "INFORMATION" }
                ]
}
$ curl -s http://localhost:8081/foglamp/audit?source=SRVUN&limit=1
{ "totalCount" : 4,
  "audit"      : [ { "timestamp" : "2018-02-25 05:22:11.053",
                    "source"    : "SRVUN",
                    "details"   : { "name": "COAP" },
                    "severity"  : "INFORMATION" }
                ]
}
$
```

POST Audit Entries

POST /foglamp/audit - create a new audit trail entry.

The purpose of the create method on an audit trail entry is to allow a user interface or an application that is using the FogLAMP API to utilize the FogLAMP audit trail and notification mechanism to raise user defined audit trail entries.

Request Payload

The request payload is a JSON object with the audit trail entry minus the timestamp.

Name	Type	Description	Example
source	string	The source of the audit trail entry.	LOGGN
severity	string	The severity of the event that triggered the audit trail entry to be written. This will be one of SUCCESS, FAILURE, WARNING or INFORMATION.	FAILURE
details	object	A JSON object that describes the detail of the audit trail event.	{ "message" : "Internal System Error" }

Response Payload

The response payload is the newly created audit trail entry.

Name	Type	Description	Example
timestamp	timestamp	The timestamp when the audit trail item was written.	2018-04-16 14:33:18.215
source	string	The source of the audit trail entry.	LOGGN
severity	string	The severity of the event that triggered the audit trail entry to be written. This will be one of SUCCESS, FAILURE, WARNING or INFORMATION.	FAILURE
details	object	A JSON object that describes the detail of the audit trail event.	{ "message" : "Internal System Error" }

Example

```
$ curl -X POST http://localhost:8081/foglamp/audit \
-d '{ "severity": "FAILURE", "details": { "message": "Internal System Error" },
↪ "source": "LOGGN" }'
{ "source": "LOGGN",
  "timestamp": "2018-04-17 11:49:55.480",
  "severity": "FAILURE",
  "details": { "message": "Internal System Error" }
}
$
$ curl -X GET http://localhost:8081/foglamp/audit?severity=FAILURE
{ "totalCount": 1,
  "audit": [ { "timestamp": "2018-04-16 18:32:28.427",
               "source" : "LOGGN",
               "details" : { "message": "Internal System Error" },
               "severity" : "FAILURE" }
            ]
}
$
```

14.2.2 Configuration Management

Configuration management is an important aspect of the REST API, however due to the discoverable form of the configuration of FogLAMP the API itself is fairly small.

The configuration REST API interacts with the configuration manager to create, retrieve, update and delete the configuration categories and values. Specifically all updates must go via the management layer as this is used to trigger the notifications to the components that have registered interest in configuration categories. This is the means by which the dynamic reconfiguration of FogLAMP is achieved.

category

The *category* interface is part of the Configuration Management for FogLAMP and it is used to create, retrieve, update and delete configuration categories and items.

GET categor(ies)

GET /foglamp/category - return the list of known categories in the configuration database

Response Payload

The response payload is a JSON object with an array of JSON objects, one per valid category.

Name	Type	Description	Example
key	string	The category key, each category has a unique textual key that defines it.	network
description	string	A description of the category that may be used for display purposes.	Network Settings
display-Name	string	Name of the category that may be used for display purposes.	Network Settings

Example

```
$ curl -X GET http://localhost:8081/foglamp/category
{
  "categories":
  [
    {
      "key": "SCHEDULER",
      "description": "Scheduler configuration",
      "displayName": "Scheduler"
    },
    {
      "key": "SMNTR",
      "description": "Service Monitor",
      "displayName": "Service Monitor"
    },
    {
      "key": "rest_api",
      "description": "FogLAMP Admin and User REST API",
      "displayName": "Admin API"
    },
    {
      "key": "service",
```

(continues on next page)

(continued from previous page)

```

    "description": "FogLAMP Service",
    "displayName": "FogLAMP Service"
  },
  {
    "key": "Installation",
    "description": "Installation",
    "displayName": "Installation"
  },
  {
    "key": "General",
    "description": "General",
    "displayName": "General"
  },
  {
    "key": "Advanced",
    "description": "Advanced",
    "displayName": "Advanced"
  },
  {
    "key": "Utilities",
    "description": "Utilities",
    "displayName": "Utilities"
  }
]
$

```

GET category

GET /foglamp/category/{name} - return the configuration items in the given category.

Path Parameters

- **name** is the name of one of the categories returned from the GET /foglamp/category call.

Response Payload

The response payload is a set of configuration items within the category, each item is a JSON object with the following set of properties.

Name	Type	Description	Example
description	string	A description of the configuration item that may be used in a user interface.	The IPv4 network address of the FogLAMP server
type	string	A type that may be used by a user interface to know how to display an item.	IPv4
default	string	An optional default value for the configuration item.	127.0.0.1
displayName	string	Name of the category that may be used for display purposes.	IPv4 address
order	integer	Order at which category name will be displayed.	1
value	string	The current configured value of the configuration item. This may be empty if no value has been set.	192.168.0.27

Example

```
$ curl -X GET http://localhost:8081/foglamp/category/rest_api
{
  "enableHttp": {
    "description": "Enable HTTP (disable to use HTTPS)",
    "type": "boolean",
    "default": "true",
    "displayName": "Enable HTTP",
    "order": "1",
    "value": "true"
  },
  "httpPort": {
    "description": "Port to accept HTTP connections on",
    "type": "integer",
    "default": "8081",
    "displayName": "HTTP Port",
    "order": "2",
    "value": "8081"
  },
  "httpsPort": {
    "description": "Port to accept HTTPS connections on",
    "type": "integer",
    "default": "1995",
    "displayName": "HTTPS Port",
    "order": "3",
    "validity": "enableHttp==\"false\"",
    "value": "1995"
  },
  "certificateName": {
    "description": "Certificate file name",
    "type": "string",
    "default": "foglamp",
    "displayName": "Certificate Name",
    "order": "4",
    "validity": "enableHttp==\"false\"",
    "value": "foglamp"
  },
  "authentication": {
    "description": "API Call Authentication",
    "type": "enumeration",
    "options": [
      "mandatory",
      "optional"
    ],
    "default": "optional",
    "displayName": "Authentication",
    "order": "5",
    "value": "optional"
  },
  "authMethod": {
    "description": "Authentication method",
    "type": "enumeration",
    "options": [
      "any",
      "password",
      "certificate"
    ],
    "value": "any"
  }
}
```

(continues on next page)

(continued from previous page)

```

    "default": "any",
    "displayName": "Authentication method",
    "order": "6",
    "value": "any"
  },
  "authCertificateName": {
    "description": "Auth Certificate name",
    "type": "string",
    "default": "ca",
    "displayName": "Auth Certificate",
    "order": "7",
    "value": "ca"
  },
  "allowPing": {
    "description": "Allow access to ping, regardless of the authentication required_
↪and authentication header",
    "type": "boolean",
    "default": "true",
    "displayName": "Allow Ping",
    "order": "8",
    "value": "true"
  },
  "passwordChange": {
    "description": "Number of days after which passwords must be changed",
    "type": "integer",
    "default": "0",
    "displayName": "Password Expiry Days",
    "order": "9",
    "value": "0"
  },
  "authProviders": {
    "description": "Authentication providers to use for the interface (JSON array_
↪object)",
    "type": "JSON",
    "default": "{\\"providers\\": [\\"username\\", \\"ldap\\"] }",
    "displayName": "Auth Providers",
    "order": "10",
    "value": "{\\"providers\\": [\\"username\\", \\"ldap\\"] }"
  }
}
$

```

GET category item

GET /foglamp/category/{name}/{item} - return the configuration item in the given category.

Path Parameters

- **name** - the name of one of the categories returned from the GET /foglamp/category call.
- **item** - the item within the category to return.

Response Payload

The response payload is a configuration item within the category, each item is a JSON object with the following set of properties.

Name	Type	Description	Example
description	string	A description of the configuration item that may be used in a user interface.	The IPv4 network address of the FogLAMP server
type	string	A type that may be used by a user interface to know how to display an item.	IPv4
default	string	An optional default value for the configuration item.	127.0.0.1
displayName	string	Name of the category that may be used for display purposes.	IPv4 address
order	integer	Order at which category name will be displayed.	1
value	string	The current configured value of the configuration item. This may be empty if no value has been set.	192.168.0.27

Example

```
$ curl -X GET http://localhost:8081/foglamp/category/rest_api/httpsPort
{
  "description": "Port to accept HTTPS connections on",
  "type": "integer",
  "default": "1995",
  "displayName": "HTTPS Port",
  "order": "3",
  "validity": "enableHttp==\"false\"",
  "value": "1995"
}
$
```

PUT category item

PUT /foglamp/category/{name}/{item} - set the configuration item value in the given category.

Path Parameters

- **name** - the name of one of the categories returned from the GET /foglamp/category call.
- **item** - the item within the category to set.

Request Payload

A JSON object with the new value to assign to the configuration item.

Name	Type	Description	Example
value	string	The new value of the configuration item.	192.168.0.27

Response Payload

The response payload is the newly updated configuration item within the category, the item is a JSON object object with the following set of properties.

Name	Type	Description	Example
description	string	A description of the configuration item that may be used in a user interface.	The IPv4 network address of the FogLAMP server
type	string	A type that may be used by a user interface to know how to display an item.	IPv4
default	string	An optional default value for the configuration item.	127.0.0.1
displayName	string	Name of the category that may be used for display purposes.	IPv4 address
order	integer	Order at which category name will be displayed.	1
value	string	The current configured value of the configuration item. This may be empty if no value has been set.	192.168.0.27

Example

```
$ curl -X PUT http://localhost:8081/foglamp/category/rest_api/httpsPort \
-d '{ "value" : "1996" }'
{
  "description": "Port to accept HTTPS connections on",
  "type": "integer",
  "default": "1995",
  "displayName": "HTTPS Port",
  "order": "3",
  "validity": "enableHttp=="false"",
  "value": "1996"
}
$
```

DELETE category item

DELETE /foglamp/category/{name}/{item}/value - unset the value of the configuration item in the given category.

This will result in the value being returned to the default value if one is defined. If not the value will be blank, i.e. the value property of the JSON object will exist with an empty value.

Path Parameters

- **name** - the name of one of the categories returned from the GET /foglamp/category call.
- **item** - the the item within the category to return.

Response Payload

The response payload is the newly updated configuration item within the category, the item is a JSON object object with the following set of properties.

Name	Type	Description	Example
description	string	A description of the configuration item that may be used in a user interface.	The IPv4 network address of the FogLAMP server
type	string	A type that may be used by a user interface to know how to display an item.	IPv4
default	string	An optional default value for the configuration item.	127.0.0.1
displayName	string	Name of the category that may be used for display purposes.	IPv4 address
order	integer	Order at which category name will be displayed.	1
value	string	The current configured value of the configuration item. This may be empty if no value has been set.	127.0.0.1

Example

```
$ curl -X DELETE http://localhost:8081/foglamp/category/rest_api/httpsPort/value
{
  "description": "Port to accept HTTPS connections on",
  "type": "integer",
  "default": "1995",
  "displayName": "HTTPS Port",
  "order": "3",
  "validity": "enableHttp==\"false\"",
  "value": "1995"
}
$
```

POST category

POST /foglamp/category - creates a new category

Request Payload

A JSON object that defines the category.

Name	Type	Description	Example
key	string	The key that identifies the category. If the key already exists as a category then the contents of this request is merged with the data stored.	backup
description	string	A description of the configuration category	Backup configuration
items	list	An optional list of items to create in this category	
name	string	The name of a configuration item	destination
description	string	A description of the configuration item	The destination to which the backup will be written
type	string	The type of the configuration item	string
default	string	An optional default value for the configuration item	/backup

NOTE: with list we mean a list of JSON objects in the form of { obj1,obj2,etc. }, to differ from the concept of *array*, i.e. [obj1,obj2,etc.]

Example

```
$ curl -X POST http://localhost:8081/foglamp/category
-d '{ "key": "My Configuration", "description": "This is my new configuration",
    "value": { "item one": { "description": "The first item", "type": "string",
↪ "default": "one" },
    "item two": { "description": "The second item", "type": "string",
↪ "default": "two" },
    "item three": { "description": "The third item", "type": "string",
↪ "default": "three" } } }'
{ "description": "This is my new configuration", "key": "My Configuration", "value": {
    "item one": { "default": "one", "type": "string", "description": "The first item
↪ ", "value": "one" },
    "item two": { "default": "two", "type": "string", "description": "The second
↪ item", "value": "two" },
    "item three": { "default": "three", "type": "string", "description": "The third
↪ item", "value": "three" } }
}
$
```

14.2.3 Task Management

The task management API's allow an administrative user to monitor and control the tasks that are started by the task scheduler either from a schedule or as a result of an API request.

task

The *task* interface allows an administrative user to monitor and control FogLAMP tasks.

GET task

GET /foglamp/task - return the list of all known task running or completed

Request Parameters

- **name** - an optional task name to filter on, only executions of the particular task will be reported.
- **state** - an optional query parameter that will return only those tasks in the given state.

Response Payload

The response payload is a JSON object with an array of task objects.

Name	Type	Description	Example
id	string	A unique identifier for the task. This takes the form of a uuid and not a Linux process id as the ID's must survive restarts and failovers	0a787bf3-4f48-4235-ae9a-2816f8ac76cc
name	string	The name of the task	purge
state	string	The current state of the task	Running
start-Time	times-tamp	The date and time the task started	2018-04-17 08:32:15.071
end-Time	times-tamp	The date and time the task ended This may not exist if the task is not completed.	2018-04-17 08:32:14.872
exit-Code	integer	Exit Code of the task.	0
reason	string	An optional reason string that describes why the task failed.	No destination available to write backup

Example

```
$ curl -X GET http://localhost:8081/foglamp/task
{
  "tasks": [
    {
      "id": "a9967d61-8bec-4d0b-8aa1-8b4dfb1d9855",
      "name": "stats collection",
      "processName": "stats collector",
      "state": "Complete",
      "startTime": "2020-05-28 09:21:58.650",
      "endTime": "2020-05-28 09:21:59.155",
      "exitCode": 0,
      "reason": ""
    },
    {
      "id": "7706b23c-71a4-410a-a03a-9b517dcd8c93",
      "name": "stats collection",
      "processName": "stats collector",
      "state": "Complete",
      "startTime": "2020-05-28 09:22:13.654",
      "endTime": "2020-05-28 09:22:14.160",
      "exitCode": 0,
      "reason": ""
    },
    ... ] }

$ curl -X GET http://localhost:8081/foglamp/task?name=purge
{
  "tasks": [
    {
      "id": "c24e006d-22f2-4c52-9f3a-391a9b17b6d6",
      "name": "purge",
      "processName": "purge",
      "state": "Complete",
      "startTime": "2020-05-28 09:44:00.175",
      "endTime": "2020-05-28 09:44:13.915",
      "exitCode": 0,
      "reason": ""
    },
    {
```

(continues on next page)

(continued from previous page)

```

    "id": "609f35e6-4e89-4749-ac17-841ae3ee2b31",
    "name": "purge",
    "processName": "purge",
    "state": "Complete",
    "startTime": "2020-05-28 09:44:15.165",
    "endTime": "2020-05-28 09:44:28.154",
    "exitCode": 0,
    "reason": ""
  },
  ... ] }
$
$ curl -X GET http://localhost:8081/foglamp/task?state=complete
{
"tasks": [
  {
    "id": "a9967d61-8bec-4d0b-8aa1-8b4dfb1d9855",
    "name": "stats collection",
    "processName": "stats collector",
    "state": "Complete",
    "startTime": "2020-05-28 09:21:58.650",
    "endTime": "2020-05-28 09:21:59.155",
    "exitCode": 0,
    "reason": ""
  },
  {
    "id": "7706b23c-71a4-410a-a03a-9b517dcd8c93",
    "name": "stats collection",
    "processName": "stats collector",
    "state": "Complete",
    "startTime": "2020-05-28 09:22:13.654",
    "endTime": "2020-05-28 09:22:14.160",
    "exitCode": 0,
    "reason": ""
  },
  ... ] }
$

```

GET task latest

GET /foglamp/task/latest - return the list of most recent task execution for each name.

This call is designed to allow a monitoring interface to show when each task was last run and what the status of that task was.

Request Parameters

- **name** - an optional task name to filter on, only executions of the particular task will be reported.
- **state** - an optional query parameter that will return only those tasks in the given state.

Response Payload

The response payload is a JSON object with an array of task objects.

Name	Type	Description	Example
id	string	A unique identifier for the task. This takes the form of a uuid and not a Linux process id as the ID's must survive restarts and failovers	0a787bf3-4f48-4235-ae9a-2816f8ac76cc
name	string	The name of the task	purge
state	string	The current state of the task	Running
start-Time	times-tamp	The date and time the task started	2018-04-17 08:32:15.071
end-Time	times-tamp	The date and time the task ended This may not exist if the task is not completed.	2018-04-17 08:32:14.872
exit-Code	integer	Exit Code of the task.	0
reason	string	An optional reason string that describes why the task failed.	No destination available to write backup
pid	integer	Process ID of the task.	17481

Example

```
$ curl -X GET http://localhost:8081/foglamp/task/latest
{
  "tasks": [
    {
      "id": "ea334d3b-8a33-4a29-845c-8be50efd44a4",
      "name": "certificate checker",
      "processName": "certificate checker",
      "state": "Complete",
      "startTime": "2020-05-28 09:35:00.009",
      "endTime": "2020-05-28 09:35:00.057",
      "exitCode": 0,
      "reason": "",
      "pid": 17481
    },
    {
      "id": "794707da-dd32-471e-8537-5d20dc0f401a",
      "name": "stats collection",
      "processName": "stats collector",
      "state": "Complete",
      "startTime": "2020-05-28 09:37:28.650",
      "endTime": "2020-05-28 09:37:29.138",
      "exitCode": 0,
      "reason": "",
      "pid": 17926
    }
  ]
}

$ curl -X GET http://localhost:8081/foglamp/task/latest?name=purge
{
  "tasks": [
    {
      "id": "609f35e6-4e89-4749-ac17-841ae3ee2b31",
      "name": "purge",
      "processName": "purge",
      "state": "Complete",
      "startTime": "2020-05-28 09:44:15.165",
      "endTime": "2020-05-28 09:44:28.154",
```

(continues on next page)

(continued from previous page)

```

    "exitCode": 0,
    "reason": "",
    "pid": 20914
  }
]
}
$

```

GET task by ID

GET /foglamp/task/{id} - return the task information for the given task

Path Parameters

- **id** - the uuid of the task whose data should be returned.

Response Payload

The response payload is a JSON object containing the task details.

Name	Type	Description	Example
id	string	A unique identifier for the task. This takes the form of a uuid and not a Linux process id as the ID's must survive restarts and failovers	0a787bf3-4f48-4235-ae9a-2816f8ac76cc
name	string	The name of the task	purge
state	string	The current state of the task	Running
start-Time	times-tamp	The date and time the task started	2018-04-17 08:32:15.071
end-Time	times-tamp	The date and time the task ended This may not exist if the task is not completed.	2018-04-17 08:32:14.872
exit-Code	integer	Exit Code of the task.	0
reason	string	An optional reason string that describes why the task failed.	No destination available to write backup

Example

```

$ curl -X GET http://localhost:8081/foglamp/task/ea334d3b-8a33-4a29-845c-8be50efd44a4
{
  "id": "ea334d3b-8a33-4a29-845c-8be50efd44a4",
  "name": "certificate checker",
  "processName": "certificate checker",
  "state": "Complete",
  "startTime": "2020-05-28 09:35:00.009",
  "endTime": "2020-05-28 09:35:00.057",
  "exitCode": 0,
  "reason": ""
}
$

```

Cancel task by ID

PUT /foglamp/task/{id}/cancel - cancel a task

Path Parameters

- **id** - the uuid of the task to cancel.

Response Payload

The response payload is a JSON object with the details of the cancelled task.

Name	Type	Description	Example
id	string	A unique identifier for the task. This takes the form of a uuid and not a Linux process id as the ID's must survive restarts and failovers	0a787bf3-4f48-4235-ae9a-2816f8ac76cc
name	string	The name of the task	purge
state	string	The current state of the task	Running
start-Time	times-tamp	The date and time the task started	2018-04-17 08:32:15.071
end-Time	times-tamp	The date and time the task ended This may not exist if the task is not completed.	2018-04-17 08:32:14.872
reason	string	An optional reason string that describes why the task failed.	No destination available to write backup

Example

```
$ curl -X PUT http://localhost:8081/foglamp/task/ea334d3b-8a33-4a29-845c-8be50efd44a4/
↪cancel
{"id": "ea334d3b-8a33-4a29-845c-8be50efd44a4", "message": "Task cancelled successfully
↪"}
$
```

14.2.4 Other Administrative API calls

ping

The *ping* interface gives a basic confidence check that the FogLAMP appliance is running and the API aspect of the appliance is functional. It is designed to be a simple test that can be applied by a user or by an HA monitoring system to test the liveness and responsiveness of the system.

GET ping

GET /foglamp/ping - return liveness of FogLAMP

NOTE: the GET method can be executed without authentication even when authentication is required. This behaviour is configurable via a configuration option.

Response Payload

The response payload is some basic health information in a JSON object.

Name	Type	Description	Example
uptime	numeric	Time in seconds since FogLAMP started	2113.076449394226
dataRead	numeric	A count of the number of sensor readings	1452
dataSent	numeric	A count of the number of readings sent to PI	347
dataPurged	numeric	A count of the number of readings purged	226
authenticationOptional	boolean	When true, the REST API does not require authentication. When false, users must successfully login in order to call the rest API. Default is <i>true</i>	true
serviceName	string	Name of service	FogLAMP
hostName	string	Name of host machine	foglamp
ipAddresses	list	IPv4 and IPv6 address of host machine	["10.0.0.0","123:234:345:456:567:789:101:102"]
health	string	Health of FogLAMP services	"green"
safeMode	boolean	True if FogLAMP is started in safe mode (only core and storage services will be started)	2113.076449394226

Example

```
$ curl -s http://localhost:8081/foglamp/ping
{
  "uptime": 276818,
  "dataRead": 0,
  "dataSent": 0,
  "dataPurged": 0,
  "authenticationOptional": true,
  "serviceName": "FogLAMP",
  "hostName": "foglamp",
  "ipAddresses": [
    "x.x.x.x",
    "x:x:x:x:x:x:x:x"
  ],
  "health": "green",
  "safeMode": false
}
```

statistics

The *statistics* interface allows the retrieval of live statistics and statistical history for the FogLAMP device.

GET statistics

GET /foglamp/statistics - return a general set of statistics

Response Payload

The response payload is a JSON document with statistical information (all numerical), these statistics are absolute counts since FogLAMP started.

Key	Description
BUFFERED	Readings currently in the FogLAMP buffer
DISCARDED	Readings discarded by the South Service before being placed in the buffer. This may be due to an error in the readings themselves.
PURGED	Readings removed from the buffer by the purge process
READINGS	Readings received by FogLAMP
UNSENT	Readings filtered out in the send process
UNSNPURGED	Readings that were purged from the buffer before being sent

Example

```
$ curl -s http://localhost:8081/foglamp/statistics
[ {
  "key": "BUFFERED",
  "description": "Readings currently in the FogLAMP buffer",
  "value": 0
},
...
{
  "key": "UNSNPURGED",
  "description": "Readings that were purged from the buffer before being sent",
  "value": 0
},
... ]
$
```

GET statistics/history

GET /foglamp/statistics/history - return a historical set of statistics. This interface is normally used to check if a set of sensors or devices are sending data to FogLAMP, by comparing the recent statistics and the number of readings received for an asset.

Request Parameters

- **limit** - limit the result set to the *N* most recent entries.

Response Payload

A JSON document containing an array of statistical information, these statistics are delta counts since the previous entry in the array. The time interval between values is a constant defined that runs the gathering process which populates the history statistics in the storage layer.

Key	Description
interval	The interval in seconds between successive statistics values
statistics[].BUFFERED	Readings currently in the FogLAMP buffer
statistics[].DISCARDED	Readings discarded by the South Service before being placed in the buffer. This may be due to an error in the readings themselves.
statistics[].PURGED	Readings removed from the buffer by the purge process
statistics[].READINGS	Readings received by FogLAMP
statistics[*NORTH_TASK_NAME]	The number of readings sent to the PI system via the OMF plugin with north instance name
statistics[].UNSENT	Readings filtered out in the send process
statistics[].UNSNPURGED	Readings that were purged from the buffer before being sent
statistics[*ASSET-CODE]	The number of readings received by FogLAMP since startup with name <i>asset-code</i>

Example

```
$ curl -s http://localhost:8081/foglamp/statistics/history?limit=2
{
  "interval": 15,
  "statistics": [
    {
      "history_ts": "2020-06-01 11:21:04.357",
      "READINGS": 0,
      "BUFFERED": 0,
      "UNSENT": 0,
      "PURGED": 0,
      "UNSNPURGED": 0,
      "DISCARDED": 0,
      "Readings Sent": 0
    },
    {
      "history_ts": "2020-06-01 11:20:48.740",
      "READINGS": 0,
      "BUFFERED": 0,
      "UNSENT": 0,
      "PURGED": 0,
      "UNSNPURGED": 0,
      "DISCARDED": 0,
      "Readings Sent": 0
    }
  ]
}
```

14.3 User API Reference

The user API provides a mechanism to access the data that is buffered within FogLAMP. It is designed to allow users and applications to get a view of the data that is available in the buffer and do analysis and possibly trigger actions based on recently received sensor readings.

In order to use the entry points in the user API, with the exception of the `/foglamp/authenticate` entry point, there must be an authenticated client calling the API. The client must provide a header field in each request, `authtoken`, the value of which is the token that was retrieved via a call to `/foglamp/authenticate`. This token must be checked for validity and also that the authenticated entity has user or admin permissions.

14.3.1 Browsing Assets

asset

The asset method is used to browse all or some assets, based on search and filtering.

GET all assets

GET /foglamp/asset - Return an array of asset codes buffered in FogLAMP and a count of assets by code.

Response Payload

An array of JSON objects, one per asset.

Name	Type	Description	Example
[].assetCode	string	The code of the asset	fogbench/accelerometer
[].count	number	The number of recorded readings for the asset code	22359

Example

```
$ curl -s http://localhost:8081/foglamp/asset
[ { "count": 18, "assetCode": "fogbench/accelerometer" },
  { "count": 18, "assetCode": "fogbench/gyroscope" },
  { "count": 18, "assetCode": "fogbench/humidity" },
  { "count": 18, "assetCode": "fogbench/luxometer" },
  { "count": 18, "assetCode": "fogbench/magnetometer" },
  { "count": 18, "assetCode": "fogbench/mouse" },
  { "count": 18, "assetCode": "fogbench/pressure" },
  { "count": 18, "assetCode": "fogbench/switch" },
  { "count": 18, "assetCode": "fogbench/temperature" },
  { "count": 18, "assetCode": "fogbench/wall clock" } ]
$
```

GET asset readings

GET /foglamp/asset/{code} - Return an array of readings for a given asset code.

Path Parameters

- **code** - the asset code to retrieve.

Request Parameters

- **limit** - set the limit of the number of readings to return. If not specified, the defaults is 20 readings.

Response Payload

An array of JSON objects with the readings data for a series of readings sorted in reverse chronological order.

Name	Type	Description	Example
[].timestamp	timestamp	The time at which the reading was received.	2018-04-16 14:33:18.215
[].reading	JSON object	The JSON reading object received from the sensor.	{ "reading": { "x":0, "y":0, "z":1 } }

Example

```
$ curl -s http://localhost:8081/foglamp/asset/fogbench%2Faccelerometer
[ { "reading": { "x": 0, "y": -2, "z": 0 }, "timestamp": "2018-04-19 14:20:59.692" },
  { "reading": { "x": 0, "y": 0, "z": -1 }, "timestamp": "2018-04-19 14:20:54.643" },
  { "reading": { "x": -1, "y": 2, "z": 1 }, "timestamp": "2018-04-19 14:20:49.899" },
  { "reading": { "x": -1, "y": -1, "z": 1 }, "timestamp": "2018-04-19 14:20:47.026" },
  { "reading": { "x": -1, "y": -2, "z": -2 }, "timestamp": "2018-04-19 14:20:42.746" }
]
$
$ curl -s http://localhost:8081/foglamp/asset/fogbench%2Faccelerometer?limit=5
[ { "reading": { "x": 0, "y": -2, "z": 0 }, "timestamp": "2018-04-19 14:20:59.692" },
  { "reading": { "x": 0, "y": 0, "z": -1 }, "timestamp": "2018-04-19 14:20:54.643" },
  { "reading": { "x": -1, "y": 2, "z": 1 }, "timestamp": "2018-04-19 14:20:49.899" },
  { "reading": { "x": -1, "y": -1, "z": 1 }, "timestamp": "2018-04-19 14:20:47.026" },
  { "reading": { "x": -1, "y": -2, "z": -2 }, "timestamp": "2018-04-19 14:20:42.746" }
]
```

GET asset reading

GET /foglamp/asset/{code}/{reading} - Return an array of single readings for a given asset code.

Path Parameters

- **code** - the asset code to retrieve.
- **reading** - the sensor from the assets JSON formatted reading.

Request Parameters

- **limit** - set the limit of the number of readings to return. If not specified, the defaults is 20 single readings.

Response Payload

An array of JSON objects with a series of readings sorted in reverse chronological order.

Name	Type	Description	Example
timestamp	timestamp	The time at which the reading was received.	2018-04-16 14:33:18.215
{reading}	JSON object	The value of the specified reading.	"temperature": 20

Example

```
$ curl -s http://localhost:8081/foglamp/asset/fogbench%2Fhumidity/temperature
[ { "temperature": 20, "timestamp": "2018-04-19 14:20:59.692" },
```

(continues on next page)

(continued from previous page)

```

{ "temperature": 33, "timestamp": "2018-04-19 14:20:54.643" },
{ "temperature": 35, "timestamp": "2018-04-19 14:20:49.899" },
{ "temperature": 0, "timestamp": "2018-04-19 14:20:47.026" },
{ "temperature": 37, "timestamp": "2018-04-19 14:20:42.746" },
{ "temperature": 47, "timestamp": "2018-04-19 14:20:37.418" },
{ "temperature": 26, "timestamp": "2018-04-19 14:20:32.650" },
{ "temperature": 12, "timestamp": "2018-04-19 14:06:05.870" },
{ "temperature": 38, "timestamp": "2018-04-19 14:06:05.869" },
{ "temperature": 7, "timestamp": "2018-04-19 14:06:05.869" },
{ "temperature": 21, "timestamp": "2018-04-19 14:06:05.868" },
{ "temperature": 5, "timestamp": "2018-04-19 14:06:05.867" },
{ "temperature": 40, "timestamp": "2018-04-19 14:06:05.867" },
{ "temperature": 39, "timestamp": "2018-04-19 14:06:05.866" },
{ "temperature": 29, "timestamp": "2018-04-19 14:06:05.865" },
{ "temperature": 41, "timestamp": "2018-04-19 14:06:05.865" },
{ "temperature": 46, "timestamp": "2018-04-19 14:06:05.864" },
{ "temperature": 10, "timestamp": "2018-04-19 13:45:15.881" } ]
$
$ curl -s http://localhost:8081/foglamp/asset/fogbench%2Faccelerometer?limit=5
[ { "temperature": 20, "timestamp": "2018-04-19 14:20:59.692" },
  { "temperature": 33, "timestamp": "2018-04-19 14:20:54.643" },
  { "temperature": 35, "timestamp": "2018-04-19 14:20:49.899" },
  { "temperature": 0, "timestamp": "2018-04-19 14:20:47.026" },
  { "temperature": 37, "timestamp": "2018-04-19 14:20:42.746" } ]
$

```

GET asset reading summary

GET /foglamp/asset/{code}/{reading}/summary - Return minimum, maximum and average values of a reading by asset code.

Path Parameters

- **code** - the asset code to retrieve.
- **reading** - the sensor from the assets JSON formatted reading.

Response Payload

An array of JSON objects with a series of readings sorted in reverse chronological order.

Name	Type	Description	Example
{reading}.average	number	The average value of the set of sensor values selected in the query string	27
{reading}.min	number	The minimum value of the set of sensor values selected in the query string	0
{reading}.max	number	The maximum value of the set of sensor values selected in the query string	47

Example

```

$ curl -s http://localhost:8081/foglamp/asset/fogbench%2Fhumidity/temperature/summary
{ "temperature": { "max": 47, "min": 0, "average": 27 } }
$

```

GET timed average asset reading

GET /foglamp/asset/{code}/{reading}/series - Return minimum, maximum and average values of a reading by asset code in a time series. The default interval in the series is one second.

Path Parameters

- **code** - the asset code to retrieve.
- **reading** - the sensor from the assets JSON formatted reading.

Request Parameters

- **limit** - set the limit of the number of readings to return. If not specified, the defaults is 20 single readings.

Response Payload

An array of JSON objects with a series of readings sorted in reverse chronological order.

Name	Type	Description	Example
times-tamp	times-tamp	The time the reading represents.	2018-04-16 14:33:18
average	number	The average value of the set of sensor values selected in the query string	27
min	number	The minimum value of the set of sensor values selected in the query string	0
max	number	The maximum value of the set of sensor values selected in the query string	47

Example

```
$ curl -s http://localhost:8081/foglamp/asset/fogbench%2Fhumidity/temperature/series
[ { "timestamp": "2018-04-19 14:20:59", "max": 20, "min": 20, "average": 20 },
  { "timestamp": "2018-04-19 14:20:54", "max": 33, "min": 33, "average": 33 },
  { "timestamp": "2018-04-19 14:20:49", "max": 35, "min": 35, "average": 35 },
  { "timestamp": "2018-04-19 14:20:47", "max": 0, "min": 0, "average": 0 },
  { "timestamp": "2018-04-19 14:20:42", "max": 37, "min": 37, "average": 37 },
  { "timestamp": "2018-04-19 14:20:37", "max": 47, "min": 47, "average": 47 },
  { "timestamp": "2018-04-19 14:20:32", "max": 26, "min": 26, "average": 26 },
  { "timestamp": "2018-04-19 14:06:05", "max": 46, "min": 5, "average": 27.8 },
  { "timestamp": "2018-04-19 13:45:15", "max": 10, "min": 10, "average": 10 } ]

$
$ curl -s http://localhost:8081/foglamp/asset/fogbench%2Fhumidity/temperature/series
[ { "timestamp": "2018-04-19 14:20:59", "max": 20, "min": 20, "average": 20 },
  { "timestamp": "2018-04-19 14:20:54", "max": 33, "min": 33, "average": 33 },
  { "timestamp": "2018-04-19 14:20:49", "max": 35, "min": 35, "average": 35 },
  { "timestamp": "2018-04-19 14:20:47", "max": 0, "min": 0, "average": 0 },
  { "timestamp": "2018-04-19 14:20:42", "max": 37, "min": 37, "average": 37 } ]
```

BUILDING FOGLAMP

15.1 Building Developers Guide

15.1.1 Introduction

What Is FogLAMP?


FogLAMP is an open source platform for the **Internet of Things** and an essential component in **Fog Computing**. It uses a modular **microservices architecture** including sensor data collection, storage, processing and forwarding to historians, Enterprise systems and Cloud-based services. FogLAMP can run in highly available, stand alone, unattended environments that assume unreliable network connectivity.

By providing a modular and distributable framework under an open source Apache v2 license, FogLAMP is the best platform to manage the data infrastructure for IoT. The modules can be distributed in any layer - Edge, Fog and Cloud - and they act together to provide scalability, elasticity and resilience.

FogLAMP offers an “all-round” solution for data management, combining a bi-directional **Northbound/Southbound** data and metadata communication with a **Eastbound/Westbound** service and object distribution.

FogLAMP Positioning in an IoT and IIoT Infrastructure

FogLAMP can be used in IoT and IIoT infrastructure at Edge and in the Fog. It stretches bi-directionally South-North/North-South and it is distributed East-West/West-East (see figure below).



Note: In this scenario we refer to “Cloud” as the layer above the Fog. “Fog” is where historians, gateways and middle servers coexist. In practice, the Cloud may also represent internal Enterprise systems, concentrated in regional or global corporate data centers, where larger historians, Big Data and analytical systems reside.

In practical terms, this means that:

- Intra-layer communication and data exchange:
 - At the **Edge**, microservices are installed on devices, sensors and actuators.
 - In the **Fog**, data is collected and aggregated in gateways and regional servers.
 - In the **Cloud**, data is distributed and analysed on multiple servers, such as Big Data Systems and Data Historians.
- Inter-layer communication and data exchange:
 - From **Edge to Fog**, data is retrieved from multiple sensors and devices and it is aggregated on resilient and highly available middle servers and gateways, either in traditional Data Historians and in the new edge of Machine Learning systems.
 - From **Fog to Edge**, configuration information, metadata and other valuable data is transferred to sensors and devices.
 - From **Fog to Cloud**, the data collected and optionally transformed is transferred to more powerful distributed Cloud and Enterprise systems.
 - From **Cloud to Fog**, results of complex analysis and other valuable information are sent to the designated gateways and middle servers that will interact with the Edge.
- Intra-layer service distribution:

- A microservice architecture based on secure communication allows lightweight service distribution and information exchange among **Edge to Edge** devices.
- FogLAMP provides high availability, scalability and data distribution among **Fog-to-Fog** systems. Due to its portability and modularity, FogLAMP can be installed on a large number of intermediate servers and gateways, as application instances, appliances, containers or virtualized environments.
- **Cloud to Cloud FogLAMP server** capabilities provide scalability and elasticity in data storage, retrieval and analytics. The data collected at the Edge and Fog, also combined with external data, can be distributed to multiple systems within a Data Center and replicated to multiple Data Centers to guarantee local and faster access.

All these operations are **scheduled, automated and executed securely, unattended** and in a **transactional** fashion (i.e. the system can always revert to a previous state in case of failures or unexpected events).

FogLAMP Features

In a nutshell, these are main features of FogLAMP:

- Transactional, always on, server platform designed to work unattended and with zero maintenance.
- Microservice architecture with secured inter-communication:
 - Core System
 - Storage Layer
 - South side, sensors and device communication
 - North side, Cloud and Enterprise communication
 - Application Modules, internal application logic
- Pluggable modules for:
 - South side: multiple, data and metadata bi-directional communication
 - North side: multiple, data and metadata bi-directional communication
 - East/West side: IN/OUT Communicator with external applications
 - Plus:
 - * Data and communication authentication
 - * Data and status monitoring and alerting
 - * Data transformation
 - * Data storage and retrieval
- Small memory and processing footprint. FogLAMP can be installed and executed on inexpensive Edge devices; microservices can be distributed on sensors and actuator boards.
- Resilient and optionally highly available.
- Discoverable and cluster-based.
- Based on APIs (RESTful and non-RESTful) to communicate with sensors and other devices, to interact with user applications, to manage the platform and to be integrated with a Cloud or Data Center-based data infrastructure.
- Hardened with default secure communication that can be optionally relaxed.

15.1.2 Developers Toolkit

Development of plugins and extensions for FogLAMP does not always require the user to have access to the source code or to have built a FogLAMP from source code. A developers toolkit exists that includes all of the header files and libraries required to develop plugins for a FogLAMP distribution. The developers toolkit is installed using the package manager in the same way that binary packages are installed.

Choosing Developers Toolkit Version

In general it is best to choose the develop toolkit version which is the same as the version of FogLAMP you are running. However if you wish to develop a plugin that supports multiple versions of FogLAMP you should choose the toolkit for the minimum version you wish to support. FogLAMP offers compatibility with newer versions but does not guarantee that a plugin developed on newer version of FogLAMP will run on an older version.

Installation

The first step in installing the developers toolkit via the package manager is to configuration the repository archive from which to obtain the package. If you have already installed FogLAMP from a package you can skip this section as it uses the same package repository.

Ubuntu or Debian

On a Ubuntu or Debian system, including the Raspberry Pi, the package manager that is supported is *apt*. You will need to add the Dianomic Systems archive server into the configuration of apt on your system. The first thing that must be done is to add the key that is used to verify the package repository. To do this run the command

```
wget -q -O - http://archives.dianomic.com/KEY.gpg | sudo apt-key add -
```

Once complete you can add the repository itself into the apt configuration file `/etc/apt/sources.list`. The simplest way to do this is the use the *add-apt-repository* command. The exact command will vary between systems;

- Raspberry Pi does not have an *apt-add-repository* command, the user must edit the apt sources file manually

```
sudo vi /etc/apt/sources.list
```

and add the line

```
deb http://archives.dianomic.com/foglamp/latest/buster/armv7l/ /
```

to the end of the file.

Note: Replace *buster* with *stretch* or *bullseye* based on the OS image used.

- Users with an Intel or AMD system with Ubuntu 18.04 should run

```
sudo add-apt-repository "deb http://archives.dianomic.com/foglamp/latest/  
↳ubuntu1804/x86_64/ / "
```

- Users with an Intel or AMD system with Ubuntu 20.04 should run

```
sudo add-apt-repository "deb http://archives.dianomic.com/foglamp/latest/  
↳ubuntu2004/x86_64/ / "
```

Note: We do not support the *aarch64* architecture with Ubuntu 20.04 yet.

- Users with an Arm system with Ubuntu 18.04, such as the Odroid board, should run

```
sudo add-apt-repository "deb http://archives.dianomic.com/foglamp/latest/
↳ubuntu1804/aarch64/ / "
```

- Users of the Mendel operating system on a Google Coral create the file `/etc/apt/sources.list.d/foglamp.list` and insert the following content

```
deb http://archives.dianomic.com/foglamp/latest/mendel/aarch64/ /
```

Once the repository has been added you must inform the package manager to go and fetch a list of the packages it supports. To do this run the command

```
sudo apt -y update
```

You are now ready to install the FogLAMP packages. You do this by running the command

```
sudo apt -y install *package*
```

Developer Toolkit Package

You are now ready to install the developer toolkit package

```
$ sudo apt install -y foglamp-dev
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following additional packages will be installed:
  autoconf automake autotools-dev binutils binutils-common binutils-x86-64-linux-gnu_
↳build-essential cpp cpp-7
  dh-python dpkg-dev fakeroot g++ g++-7 gcc gcc-4.8-base gcc-7 gcc-7-base_
↳libalgorithm-diff-perl
  libalgorithm-diff-xs-perl libalgorithm-merge-perl libasan0 libasan4 libatomic1_
↳libbinutils
  libboost-atomic1.65-dev libboost-atomic1.65.1 libboost-chrono1.65-dev libboost-
↳chrono1.65.1
  libboost-date-time1.65-dev libboost-date-time1.65.1 libboost-dev libboost-
↳serialization1.65-dev
  libboost-serialization1.65.1 libboost-system-dev libboost-system1.65-dev libboost-
↳system1.65.1
  libboost-thread-dev libboost-thread1.65-dev libboost-thread1.65.1 libboost1.65-dev_
↳libc-dev-bin libc6 libc6-dev
  libcc1-0 libcilkrts5 libdpkg-perl libexpat1-dev libfakeroot libfile-fcntllock-perl_
↳libgcc-4.8-dev libgcc-7-dev
  libgomp1 libisl19 libitm1 liblsan0 libltdl-dev libltdl7 libmpc3 libmpx2 libpq-dev_
↳libpq5 libpython3-dev
  libpython3.6-dev libquadmath0 libsensors4 libstdc++-4.8-dev libstdc++-7-dev libtool_
↳libtsan0 libubsan0
  linux-libc-dev m4 make manpages-dev postgresql postgresql-10 postgresql-client-10_
↳postgresql-client-common
  postgresql-common python-pip-whl python3-crypto python3-dev python3-distutils_
↳python3-keyring
```

(continues on next page)

(continued from previous page)

```

python3-keyrings.alt python3-lib2to3 python3-pip python3-secretstorage python3-
↪setuptools python3-wheel
python3-xdg python3.6-dev sqlite3 ssl-cert sysstat
Suggested packages:
  autoconf-archive gnu-standards autoconf-doc gettext binutils-doc cpp-doc gcc-7-
↪locales debian-keyring
  g++-multilib g++-7-multilib gcc-7-doc libstdc++6-7-dbg gcc-multilib flex bison gdb_
↪gcc-doc gcc-7-multilib
  libgcc1-dbg libgomp1-dbg libitm1-dbg libatomic1-dbg libasan4-dbg liblsan0-dbg_
↪libtsan0-dbg libubsan0-dbg
  libcilkrts5-dbg libmpx2-dbg libquadmath0-dbg libboost-doc libboost1.65-doc libboost-
↪container1.65-dev
  libboost-context1.65-dev libboost-coroutine1.65-dev libboost-exception1.65-dev_
↪libboost-fiber1.65-dev
  libboost-filesystem1.65-dev libboost-graph1.65-dev libboost-graph-parallel1.65-dev_
↪libboost-iostreams1.65-dev
  libboost-locale1.65-dev libboost-log1.65-dev libboost-math1.65-dev libboost-mpi1.65-
↪dev
  libboost-mpi-python1.65-dev libboost-numpy1.65-dev libboost-program-options1.65-dev_
↪libboost-python1.65-dev
  libboost-random1.65-dev libboost-regex1.65-dev libboost-signals1.65-dev libboost-
↪stacktrace1.65-dev
  libboost-test1.65-dev libboost-timer1.65-dev libboost-type-erasure1.65-dev libboost-
↪wave1.65-dev
  libboost1.65-tools-dev libmpfrc++-dev libntl-dev glibc-doc bzip libtool-doc_
↪postgresql-doc-10 lm-sensors
  libstdc++-4.8-doc libstdc++-7-doc gfortran | fortran95-compiler gcj-jdk m4-doc make-
↪doc postgresql-doc
  locales-all libjson-perl python-crypto-doc gnome-keyring libkf5wallet-bin gir1.2-
↪gnomekeyring-1.0
  python-secretstorage-doc python-setuptools-doc sqlite3-doc openssl-blacklist isag
...
Setting up foglamp-dev (1.9.2) ...
Processing triggers for ureadahead (0.100.0-21) ...
Processing triggers for install-info (6.5.0.dfsg.1-2) ...
Processing triggers for libc-bin (2.27-3ubuntu1.5) ...
Processing triggers for systemd (237-3ubuntu10.52) ...
Processing triggers for man-db (2.8.3-2ubuntu0.1) ...
$

```

At the end of this command all of the packages needed to build plugins for FogLAMP will have been installed and the FogLAMP header files and libraries will also be installed. FogLAMP itself will not be installed.

The FogLAMP header files can be found in the directory `/usr/include/foglamp` and the libraries in `/usr/lib/foglamp`.

Building Plugins

FogLAMP uses the *cmake* system to build plugins and other components. In order to find the various libraries and header files there is a standard *FindFoglamp.cmake* file that is used. This will allow the *cmake* system to find the libraries and header files when FogLAMP has either been installed as source code or via the developer package. The content of the *FindFoglamp.cmake* file is reproduced below.

```

# This CMake file locates the Foglamp header files and libraries
#
# The following variables are set:

```

(continues on next page)

(continued from previous page)

```

# FOGLAMP_INCLUDE_DIRS - Path(s) to Foglamp headers files found
# FOGLAMP_LIB_DIRS - Path to Foglamp shared libraries
# FOGLAMP_SUCCESS - Set on succes
#
# In case of error use SEND_ERROR and return()
#

# Set defaults paths of installed Foglamp SDK package
set(FOGLAMP_DEFAULT_INCLUDE_DIR "/usr/include/foglamp" CACHE INTERNAL "")
set(FOGLAMP_DEFAULT_LIB_DIR "/usr/lib/foglamp" CACHE INTERNAL "")

# CMakeLists.txt options
set(FOGLAMP_SRC "" CACHE INTERNAL "")
set(FOGLAMP_INCLUDE "" CACHE INTERNAL "")
set(FOGLAMP_LIB "" CACHE INTERNAL "")

# Return variables
set(FOGLAMP_INCLUDE_DIRS "" CACHE INTERNAL "")
set(FOGLAMP_LIB_DIRS "" CACHE INTERNAL "")
set(FOGLAMP_FOUND "" CACHE INTERNAL "")

# No options set
# If FOGLAMP_ROOT env var is set, use it
if (NOT FOGLAMP_SRC AND NOT FOGLAMP_INCLUDE AND NOT FOGLAMP_LIB)
    if (DEFINED ENV{FOGLAMP_ROOT})
        message(STATUS "No options set.\n"
            "    +Using found FOGLAMP_ROOT $ENV{FOGLAMP_ROOT}")
        set(FOGLAMP_SRC $ENV{FOGLAMP_ROOT})
    endif()
endif()

# -DFOGLAMP_SRC=/some_path or FOGLAMP_ROOT path
# Set return variable FOGLAMP_INCLUDE_DIRS
if (FOGLAMP_SRC)
    unset(_INCLUDE_LIST CACHE)
    file(GLOB_RECURSE _INCLUDE_COMMON "${FOGLAMP_SRC}/C/common/*.h")
    file(GLOB_RECURSE _INCLUDE_SERVICES "${FOGLAMP_SRC}/C/services/common/*.h")
    list(APPEND _INCLUDE_LIST ${_INCLUDE_COMMON} ${_INCLUDE_SERVICES})
    foreach(_ITEM ${_INCLUDE_LIST})
        get_filename_component(_ITEM_PATH ${_ITEM} DIRECTORY)
        list(APPEND FOGLAMP_INCLUDE_DIRS ${_ITEM_PATH})
    endforeach()
    unset(INCLUDE_LIST CACHE)

    list(REMOVE_DUPLICATES FOGLAMP_INCLUDE_DIRS)

    string(REPLACE ";" "\n    +" DISPLAY_PATHS "${FOGLAMP_INCLUDE_DIRS}")
    if (NOT DEFINED ENV{FOGLAMP_ROOT})
        message(STATUS "Using -DFOGLAMP_SRC option for includes\n    +" "${
→{DISPLAY_PATHS}")
    else()
        message(STATUS "Using FOGLAMP_ROOT for includes\n    +" "${DISPLAY_
→PATHS}")
    endif()

    if (NOT FOGLAMP_INCLUDE_DIRS)
        message(SEND_ERROR "Needed Foglamp header files not found in path $
→{FOGLAMP_SRC}/C")

```

(continues on next page)

(continued from previous page)

```

        return()
    endif()
else()
    # -DFOGLAMP_INCLUDE=/some_path
    if (NOT FOGLAMP_INCLUDE)
        set(FOGLAMP_INCLUDE ${FOGLAMP_DEFAULT_INCLUDE_DIR})
        message(STATUS "Using Foglamp dev package includes " ${FOGLAMP_
↪INCLUDE})
    else()
        message(STATUS "Using -DFOGLAMP_INCLUDE option " ${FOGLAMP_INCLUDE})
    endif()
    # Remove current value from cache
    unset(_FIND_INCLUDES CACHE)
    # Get up to date var from find_path
    find_path(_FIND_INCLUDES NAMES plugin_api.h PATHS ${FOGLAMP_INCLUDE})
    if (_FIND_INCLUDES)
        list(APPEND FOGLAMP_INCLUDE_DIRS ${_FIND_INCLUDES})
    endif()
    # Remove current value from cache
    unset(_FIND_INCLUDES CACHE)

    if (NOT FOGLAMP_INCLUDE_DIRS)
        message(SEND_ERROR "Needed Foglamp header files not found in path $
↪${FOGLAMP_INCLUDE}")
        return()
    endif()
endif()

#
# Foglamp Libraries
#
# Check -DFOGLAMP_LIB=/some path is valid
# or use FOGLAMP_SRC/cmake_build/C/lib
# FOGLAMP_SRC might have been set to FOGLAMP_ROOT above
#
if (FOGLAMP_SRC)
    # Set return variable FOGLAMP_LIB_DIRS
    set(FOGLAMP_LIB "${FOGLAMP_SRC}/cmake_build/C/lib")

    if (NOT DEFINED ENV{FOGLAMP_ROOT})
        message(STATUS "Using -DFOGLAMP_SRC option for libs \n    +" "$
↪${FOGLAMP_SRC}/cmake_build/C/lib")
    else()
        message(STATUS "Using FOGLAMP_ROOT for libs \n    +" "${FOGLAMP_SRC}/
↪cmake_build/C/lib")
    endif()

    if (NOT EXISTS "${FOGLAMP_SRC}/cmake_build")
        message(SEND_ERROR "Foglamp has not been built yet in ${FOGLAMP_SRC} ↪
↪Compile it first.")
        return()
    endif()

    # Set return variable FOGLAMP_LIB_DIRS
    set(FOGLAMP_LIB_DIRS "${FOGLAMP_SRC}/cmake_build/C/lib")
else()
    if (NOT FOGLAMP_LIB)

```

(continues on next page)

(continued from previous page)

```

        set(FOGLAMP_LIB ${FOGLAMP_DEFAULT_LIB_DIR})
        message(STATUS "Using Foglamp dev package libs " ${FOGLAMP_LIB})
    else()
        message(STATUS "Using -DFOGLAMP_LIB option " ${FOGLAMP_LIB})
    endif()
    # Set return variable FOGLAMP_LIB_DIRS
    set(FOGLAMP_LIB_DIRS ${FOGLAMP_LIB})
endif()

# Check NEEDED_FOGLAMP_LIBS in libraries in FOGLAMP_LIB_DIRS
# NEEDED_FOGLAMP_LIBS variables comes from CMakeLists.txt
foreach(_LIB ${NEEDED_FOGLAMP_LIBS})
    # Remove current value from cache
    unset(_FOUND_LIB CACHE)
    # Get up to date var from find_library
    find_library(_FOUND_LIB NAME ${_LIB} PATHS ${FOGLAMP_LIB_DIRS})
    if (_FOUND_LIB)
        # Extract path form founf library file
        get_filename_component(_DIR_LIB ${_FOUND_LIB} DIRECTORY)
    else()
        message(SEND_ERROR "Needed Foglamp library ${_LIB} not found in $
→ ${FOGLAMP_LIB_DIRS}")
        return()
    endif()
    # Remove current value from cache
    unset(_FOUND_LIB CACHE)
endforeach()

# Set return variable FOGLAMP_FOUND
set(FOGLAMP_FOUND "true")

```

This should be placed in the base directory of your plugin, along with the plugins *CMakeLists.txt* file. In that file simply add the lines

```

# Find Foglamp includes and libs, by including FindFoglamp.cmake file
set(CMAKE_MODULE_PATH ${CMAKE_MODULE_PATH} ${CMAKE_CURRENT_SOURCE_DIR})
find_package(Foglamp)
include_directories(${FOGLAMP_INCLUDE_DIRS})

```

This will enable *cmake* to find the header files and libraries that you include. It will also enable you to add the FogLAMP libraries using directives such as

```
target_link_libraries(${PROJECT_NAME} common-lib)
```

The above example will add the FogLAMP *common-lib* to the build, this will give access to readings objects and other common objects needed for a FogLAMP plugin.

The convention for any plugins built for FogLAMP is to provide a script called *requirements.sh* that will install any dependencies required by the plugin. If such a script exists then this should be run as the first step in building a plugin or service.

Building then becomes a case of

- creating the build directory,
- changing to that directory,
- running *cmake*,

- followed by *make*
- and optionally *make install*.

```
$ mkdir build
$ cd build
$ cmake ..
-- The C compiler identification is GNU 7.5.0
-- The CXX compiler identification is GNU 7.5.0
-- Check for working C compiler: /usr/bin/cc
-- Check for working C compiler: /usr/bin/cc -- works
-- Detecting C compiler ABI info
-- Detecting C compiler ABI info - done
-- Detecting C compile features
-- Detecting C compile features - done
-- Check for working CXX compiler: /usr/bin/c++
-- Check for working CXX compiler: /usr/bin/c++ -- works
-- Detecting CXX compiler ABI info
-- Detecting CXX compiler ABI info - done
-- Detecting CXX compile features
-- Detecting CXX compile features - done
-- Using Foglamp dev package includes /usr/include/foglamp
-- Using Foglamp dev package libs /usr/lib/foglamp
-- Configuring done
-- Generating done
-- Build files have been written to: /home/mark/foglamp-south-etherip/build
$ make
[ 25%] Generating version header
Scanning dependencies of target etherip
[ 50%] Building CXX object CMakeFiles/etherip.dir/plctag.o
[ 50%] Building CXX object CMakeFiles/etherip.dir/plugin.o
[ 75%] Linking CXX shared library libetherip.so
[100%] Built target etherip
$
```

VERSION HISTORY

16.1 FogLAMP v2

16.1.1 v2.0.0

Release Date: 2022-09-09

- **FogLAMP Core**

- New Features:

- * Add options for choosing the FogLAMP Asset name: Browser Name, Subscription Path and Full Path. Use the OPC UA Source timestamp as the User Timestamp in FogLAMP.
- * The storage interface used to query generic configuration tables has been improved to support tests for null and non-null column values.
- * The ability for north services to support control inputs coming from systems north of FogLAMP has been introduced.
- * The handling of a failed storage service has been improved. The client now attempt to re-connect and if that fails they will down. The logging produced is now much less verbose, removing the repeated messages previously seen.
- * A new service has been added to FogLAMP to facilitate the routing of control messages within FogLAMP. This service is responsible for determining which south services to send control requests to and also for the security aspects of those requests.
- * Ensure that new FogLAMP data types not supported by OMF are not processed.
- * The storage service now supports a richer set of queries against the generic table interface. In particular, joins between tables are now supported.
- * OPC UA Security has been enhanced. This plugin now supports Security Policies Basic256 and Basic256Sha256, with Security Modes Sign and Sign & Encrypt. Authentication types are anonymous and username/password.
- * South services that have a slow poll rate can take a long time to shutdown, this sometimes resulted in those services not shutting down cleanly. The shutdown process has been modified such that these services now shutdown promptly regardless of polling rate.
- * A new configuration item type has been added for the selection of access control lists.
- * Support has been added to the Python query builder for NULL and NOT NULL columns.
- * The Python query builder has been updated to support nested database queries.

- * The third party packages on which FogLAMP is built have been updated to use the latest versions to resolve issues with vulnerabilities in these underlying packages.
- * When the data stream from a south plugin included an OMF Hint of AFLocation, performance of the OMF North plugin would degrade. In addition, process memory would grow over time. These issues have been fixed.
- * The version of the PostgreSQL database used by the Postgres storage plugin has been updated to PostgreSQL 13.
- * An enhancement has been added to the North service to allow the user to specify the block size to use when sending data to the plugin. This helps tune the north services and is described in the tuning guide within the documentation.
- * The notification server would previously output warning messages when it was starting, these were not an indication of a problem and should have been information messages. This has now been resolved.
- * The backup mechanism has been improved to include some external items in the backup and provide a more secure backup.
- * The purge option that controls if unsent assets can be purged or not has been enhanced to provide options for sent to any destination or sent to all destinations as well as sent to no destinations.
- * It is now possible to add control features to Python south plugins.
- * Certificate based authentication is now possible between services in a single instance. This allows for secure control messages to be implemented between services.
- * Performance improvements have been made such that the display of south service data when large numbers of assets are in use.
- * The new micro service, control dispatcher, is now available as a package that can be installed via the package manager.
- * New data types are now supported for data points within an asset and are encoded into various Python types when passed to Python plugins or scripts run within standard plugin. This includes numpy arrays for images and data buffers, 2 dimensional Python lists and others. Details of the type encoding can be found in the plugin developers guide of the online product documentation.
- * The mechanism for online update of configuration has been extended to allow for more configuration to be modified without the need to restart any services.
- * Support has been added for the Raspberry Pi Bullseye release.
- * A problem with a file descriptor leak in Python that could cause FogLAMP to fail has been resolved.
- * The control of logging levels has now been added to the Python code run within a service such that the advanced settings option is now honoured by the Python code.
- * Enhancements have been made to the asset tracker API to retrieve the service responsive for the ingest of a given asset.
- * A new developers toolkit has been added to FogLAMP to aid the production and deployment of machine learning models to FogLAMP plugins.
- * A new micro service has been added to FogLAMP, the bucket storage service. This is an optional service that may be installed, its purpose is to provide a mechanism for other system components to store large objects within the FogLAMP system. Each object stored is assigned a number of attributes and may be searched for and retrieved using attribute matching or via the unique ID that is associated with each stored object.
- * A new API has been added to allow external viewing and managing of the data that various plugins persist.

- * A new REST API entry point has been added that allows all instances of a specified asset to be purged from the buffer. A further entry point has also been added to purge all data from the reading buffer. These entry points should be used with care as they will cause data to be discarded.
- * A new parameter has been added to the asset retrieval API that allows image data to be returned, `images=include`. By default image type datapoints will be replaced with a message, "Image removed for brevity", in order to reduce the size of the returned payload.
- * A new API has been added to the management API that allows services to request that URL's in the public API are proxied to the service API. This is used when extending the functionality of the system with custom microservices.
- * A new set of API calls have been added to the public REST API of the product to support the control dispatcher and for the creation and management of control scripts.
- * A new API has been added to the public API that will return the latest reading for a given asset. This will return all data types including images.
- * A new API has been added that allows asset tracking records to be marked as deprecated. This allows the flushing of relationships between assets and the services that have processed them. It is useful only in development systems and should not be used in production systems.
- * A new API call has been added that allows the persisted data related to a plugin to be retrieved via the public REST API. This is intended for use by plugin writers and to allow for better tracking of data persisted between service executions.
- * A new query parameter has been added to the API used to fetch log messages from the system log, `nontotals`. This will increase the performance of the call at the expense of not returning the total number of logs that match the search criteria.
- * New API entry points have been added for the management of Python packages.
- * Major performance improvements have been made to the code for retrieving log messages from the system log. This is mainly an issue on systems with very large log files.
- * The storage service API has been extended to support the creation of private schemas for the use of optional micro services registered to a FogLAMP instance.
- * Filtering by service type has now been added to the API that retrieve service information via the public REST API.
- * A number of new features have been added to the user interface to aid developers creating data pipelines and plugins. These features allow for manual purging of data, deprecating the relationship between the services and the assets they have ingested and viewing the persisted data of the plugins. These are all documented in the section on developing pipelines within the online documentation.
- * A new section has been added to the documentation which discusses the process and best practices for building data pipelines in FogLAMP.
- * A glossary has been added to the documentation for the product.
- * The documentation that describes the writing of asynchronous Python plugins has been updated in line with the latest code changes.
- * The documentation has been updated to reflect the new tabs available in the FogLAMP user interface for editing the configuration of services and tasks.
- * A new introduction section has been added to the FogLAMP documentation that describes the new features and some typical use cases of FogLAMP.
- * A new section has been added to the FogLAMP Tuning guide that discusses the tuning of North services and tasks. Also scheduler tuning has been added to the tuning guide along with the tuning of the service monitor which is used to detect failures of services within FogLAMP.

- * The Tuning FogLAMP section of the documentation has been updated to include information on tuning the FogLAMP service monitor that is used to monitor and restart FogLAMP services. A section has also been added that describes the tuning of north services and tasks. A new section describes the different storage plugins available, when they should be used and how to tune them.
 - * Added an article on Developing with Windows Subsystem for Linux (WSL2) to the Plugin Developer Guide. WSL2 allows you to run a Linux environment directly on Windows without the overhead of Windows Hyper-V. You can run FogLAMP and develop plugins on WSL2.
 - * Documentation has been added for the purge process and the new options recently added.
 - * Documentation has been added to the plugin developer guides that explain what needs to be done to allow the packaging mechanism to be able to package a plugin.
 - * Documentation has been added to the Building Pipelines section of the documentation for the new UI feature that allows Python packages to be installed via the user interface.
 - * Documentation has been updated to show how to build FogLAMP using the requirements.sh script.
 - * The documentation ordering has been changed to make the section order more logical.
 - * The plugin developers guide has been updated to include information on the various flags that are used to communicate the options implemented by a plugin.
 - * Updated OMF North plugin documentation to include current OSIsoft (AVEVA) product names.
 - * Fixed a typo in the quick start guide.
 - * Improved north plugin developers documentation is now available.
- Bug Fix:
- * The FogLAMP control script has options for purge and reset that requires a confirmation before it will continue. The message that was produced if this confirmation was not given was unclear. This has now been improved.
 - * An issue that could cause a north service or task that had been disabled for a long period of time to fail to send data when it was re-enabled has been resolved.
 - * S2OPCUA Toolkit changes required an update in build procedures for the S2OPCUA South Plugin.
 - * Previously it has not been possible to configure the advanced configuration of a south service until it has been run at least once. This has now been resolved and it is possible to add a south service in disable mode and edit the advanced configuration.
 - * The diagnostics when a plugin fails to load have been improved.
 - * The South Plugin shutdown problem was caused by errors in the plugin startup procedure which would throw an exception for any error. The plugin startup has been fixed so errors are reported properly. The problem of plugin shutdown when adding a filter has been resolved.
 - * The S2OPCUA South Plugin would throw an exception for any error during startup. This would cause the core system to shut down the plugin permanently after a few retries. This has been fixed. Error messages has been recategorized to properly reflect informational, warning and error messages.
 - * The update process has been optimised to remove an unnecessary restart if no new version of the software are available.
 - * The OMF North plugin was unable to process configuration changes or shut down if the PI Web API hostname was not correct. This has been fixed.
 - * S2OPC South plugin builds have been updated to explicitly reference S2OPC Toolkit Version 1.2.0.
 - * An issue that could on rare occasions cause the SQLite plugin to silently discard readings has been resolved.

- * An issue with the automatic renewal of authentication certificates has been resolved.
- * Deleting a service which had a filter pipeline could cause some orphaned configuration information to be left stored. This prevented creating filters of the same name in the future. This has now been resolved.
- * The error reporting has been improved when downloading backups from the system.
- * An issue that could cause north plugins to occasionally fail to shutdown correctly has now been resolved.
- * Some fixes are made in Package update API that allows the core package to be updated.
- * Improvements have been made to the exponential moving average filter to resolve issues seen when data is heavily delayed before passing through the filter.
- * The documentation has been updated to correct a statement regarding running the south side as a task.

- **GUI**

- New Features:

- * A new *Developer* item has been added to the user interface to allow for the management of Python packages via the UI. This is enabled by turning on developer features in the user interface *Settings* page.
 - * A control has been added that allows the display of assets in the *South* screen to be collapsed or expanded. This allows for more services to be seen when services ingest multiple assets.
 - * A new feature has been added to the south page that allows the relationship between an asset and a service to be deprecated. This is a special feature enabled with the Developer Features option. See the documentation on building pipelines for a full description.
 - * A new feature has been added to the Assets and Readings page that allows for manual purging of named assets or all assets. This is a developer only feature and should not be used on production systems. The feature is enabled, along with other developer features via the Settings page.
 - * A new feature has been added to the South and North pages for each service that allows the user to view, import, export and delete the data persisted by a plugin. This is a developer only feature and should not be used on production systems. It is enabled via the Setting page.
 - * A new configuration type, Access Control List, is now supported in user interface. This allows for selection of an ACL from those already created.
 - * A new tabbed layout has been adopted for the editing of south and north services and tasks. Configuration, Advanced and Security tabs are supported as our tabs for developer features if enabled.
 - * The user interface for displaying system logs has been modified to improve the performance of log viewing.
 - * The User Interface has been updated to use the latest versions of a number of packages it depends upon, due to vulnerabilities reported in those packages.
 - * With the introduction of image data types to the readings supported by the system the user interface has been updated to add visualisation features for these images. A new feature also allows the latest reading for a given asset to be shown.
 - * A new feature has been added to the south and north pages that allows the user to view the logs for the service.
 - * The service status display now includes the Control Dispatcher service if it has been installed.
 - * The user interface now supports the new control dispatcher service. This includes the graphical creation and editing of control scripts and access control lists used by control features.

- * An option has been added to the Asset and Readings page to show just the latest values for a given asset.
 - * The notification user interface now links to the relevant sections of the online documentation allowing users to navigate to the help based on the current context.
 - * Some timezone inconsistencies in the user interface have been resolved.
 - * The user interface now has a new screen used to manage machine learning model storage within a FogLAMP instance. These models may then be used by plugins within the system.
- Bug Fix:
- * An issue that would cause the GUI to not always allow JSON data to be saved has been resolved.
 - * An issue with the auto refresh in the systems log page that made selecting the service to filter difficult has been resolved.
 - * The sorting of services and tasks in the South and North pages has been improved such that enabled services appear above disabled services.
 - * An issue the prevented gaps in the data from appearing in the graphs displayed by the GUI has now been resolved.
 - * Entering times in the GUI could sometimes be difficult and result in unexpected results. This has now been improved to ease the entry of time values.
 - * An issue that made the editing of scripts in the user interface for south plugins has been resolved.

• Plugins

– New Features:

- * A new foglamp-filter-contrast has been added that allows the contrast of image type datapoints to be altered.
- * A new foglamp-filter-mirror plugin has been added that will mirror image data points in a reading either vertically or horizontally.
- * A new foglamp-filter-rotate plugin has been added that allows for images in data points to be rotated by 90, 180 or 270 degrees.
- * A new foglamp-filter-greyscale plugin has been added that will convert 24bit RGB images in the reading data to either 8bit or 16bit greyscale images. All non-image and on-24bit images are left unaltered by the plugin.- A new foglamp-notify-control plugin has been added that allows notifications to be delivered via the control dispatcher service. This allows the full features of the control dispatcher to be used with the edge notification path.
- * A new foglamp-rule-watchdog plugin has been added that allows notifications to be send if data stops being ingress for specified assets.
- * A new foglamp-south-video4linux has been created that uses the Video4Linux interface to support the capture of image data from a variety of of video sources supported by Linux.
- * A new foglamp-south-etherip plugin has been added to retrieve data for a number of different Allan Bradley PLCs. Also it has enable control features. This allows FogLAMP to write data back to PLC's using the etherip protocol.
- * Support has been added for proxy servers in the north HTTP-C plugin.
- * The OPCUA north plugin has been updated to include the ability for systems outside of FogLAMP to write to the server that FogLAMP advertises. These write are taken as control input into the FogLAMP system.

- * The HTTPC North plugin has been enhanced to add an optional Python script that can be used to format the payload of the data sent in the HTTP REST request.
 - * The SQLite storage plugins have been updated to support service extension schemas. This is a mechanism that allows services within the FogLAMP system to add new schemas within the storage service that are exclusive to that service.
 - * The Python35 filter has been updated to use the common Python interpreter. This allows for packages such as numpy to be used. The resilience and error reporting of this plugin have also been improved.
 - * A set of developer only features designed to aid the process of developing data pipelines and plugins has been added in this release. These features are turned on and off via a toggle setting on the Settings page.
 - * A new option has been added to the Python35 filter that changes the way datapoint names are used in the JSON readings. Previously there had to be encoded and decode by use of the b'xxx' mechanism. There is now a toggle that allows for either this to be required or simple text string use to be enabled.
 - * The API of the storage service has been updated to allow for custom schemas to be created by services that extend the core functionality of the system.
 - * New image type datapoints can now be sent between instances using the http north and south plugins.
 - * The ability to define response headers in the http south plugin has been added to aid certain circumstances where CORS provided data flows.
 - * The documentation of the Python35 filter has been updated to include a fuller description of how to make use of the configuration data block supported by the plugin.
 - * The documentation describing how to run services under the debugger has been improved along with other improvements to the documentation aimed at plugin developers.
 - * Documentation has been added for the Azure north plugin.
 - * Documentation has now been added for foglamp-north-harperdb.
 - * Documentation has been added for the custom asset notification plugin.
 - * The documentation has been updated to include the new watchdog notification rule.
 - * Documentation has been added for the Suez Water south plugin
 - * Documentation has been added for the foglamp-rule-periodic plugin.
 - * Documentation has been added for the foglamp-filter-asset-split plugin.
 - * Documentation has been added for the foglamp-filter-specgram plugin.
 - * Documentation has been added for the foglamp-filter-fft2 plugin.
- Bug Fix:
- * Build procedures were updated to accommodate breaking changes in the S2OPC OPCUA Toolkit.
 - * Occasionally switching from the sqlite to the sqlitememory plugin for the storage of readings would cause a fatal error in the storage layer. This has now been fixed and it is possible to change to sqlitememory without an error.
 - * A race condition within the modbus south plugin that could cause unfair scheduling of read versus write operations has been resolved. This could cause write operations to be delayed in some circumstances. The scheduling of set point write operations is now fairly interleaved between the read operations in all cases.
 - * A problem that caused the HTTPC North plugin to fail if the path component of the URL was omitted has been resolved.

- * The modbus-c south plugin documentation has been enhanced to include details of the function codes used to read modbus data.
- * An incorrect error message in the modbus-c south plugin has been fixed and others have been improved to aid resolving configuration issues. The documentation has been updated to include descriptive text for the error messages that may occur.
- * The Python35 filter plugin has been updated such that if no data is to be passed onwards it may now simply return the None Python constant or an empty list.
- * The Python35 plugin which allows simple Python scripts to be added into filter pipelines has had a number of updates to improve the robustness of the plugin in the event of incorrect script code being provided by the user. The behaviour of the plugin has also been updated such that any errors run the script will prevent data being passed onwards the filter pipeline.
- * The Average rule has been updated to improve the user interaction during the configuration of the rule.
- * The first time a plugin that persisted data is executed erroneous errors and warnings would be written to the system log. This has now been resolved.
- * Python35 filter code that failed to return a properly formed asset in the response would previously crash rather than fail gracefully. An error explaining the exact cause of the failure is now logged in the system log.
- * An issue with the Kafka north plugin not sending data in certain circumstances has been resolved.
- * Adding some notification plugins would cause incorrect errors to be logged to the system log. The functioning of the notifications was not affected. This has now been resolved and the error logs no longer appear.
- * The Simple-REST south plugin has been enhanced to allow it to return no records using the Python None mechanism. Also its error reporting has been improved for cases where the script is missing or incorrectly named.
- * A problem with installing the csvplayback plugin on aarch64 platforms has been resolved.
- * The MQTT Scripted south plugin has been updated to give improved error messages when problems are found with the supplied convert script. The documentation has also been updated to include these messages and typical causes of the errors.
- * The documentation for the foglamp-rule-delta plugin has been corrected.
- * The documentation for the Python35 filter has been updated to discuss Python package imports and issues when removing previously used imports.

16.2 FogLAMP v1

16.2.1 v1.9.2

Release Date: 2021-09-29

- **FogLAMP Core**

- New Features:

- * The ability for south plugins to persist data between executions of south services has been added for plugins written in C/C++. This follows the same model as already available for north plugins.

- * Notification delivery plugins now also receive the data that caused the rule to trigger. This can be used to deliver values in the notification delivery plugins.
- * A new option has been added to the sqlite storage plugin only that allows assets to be excluded from consideration in the purge process.
- * A new purge process has been added to control the growth of statistics history and audit trails. This new process is known as the “System Purge” process.
- * The support bundle has been updated to include details of the packages installed.
- * The package repository API endpoint has been updated to support Ubuntu 20.04 repository end point.
- * The handling of updates from RPM package repositories has been improved.
- * The certificate store has been updated to support more formats of certificates, including DER, P12 and PFX format certificates.
- * The documentation has been updated to include an improved & detailed introduction to filters.
- * The OMF north plugin documentation has been re-organised and updated to include the latest features that have been introduced to this plugin.
- * A new section has been added to the documentation that discusses the tuning of the edge based control path.

– **Bug Fix:**

- * A rare race condition during ingestion of readings would cause the south service to terminate and restart. This has now been resolved.
- * In some circumstances it was seen that north services could send the same data more than once. This has now been corrected.
- * An issue that caused an intermittent error in the tracking of data sent north has been resolved. This only impacted north services and not north tasks.
- * An optimisation has been added to prevent north plugins being sent empty data sets when the filter chain removes all the data in a reading set.
- * An issue that prevented a north service restarting correctly when certain combinations of filters were present has been resolved.
- * The API for retrieving the list of backups on the system has been improved to honour the limit and offset parameters.
- * An issue with the restore operation always restoring the latest backup rather than the chosen backup has been resolved.
- * The support package failed to include log data if binary data had been written to syslog. This has now been resolved.
- * The configuration category for the system purge was in the incorrect location with the configuration category tree, this has now been correctly placed underneath the “Utilities” item.
- * It was not possible to set a notification to always retrigger as there was a limitation that there must always be 1 second between notification triggers. This restriction has now been removed and it is possible to set a retrigger time of zero.
- * An error in the documentation for the plugin developers guide which incorrectly documented how to build debug binaries has been corrected.

• **GUI**

- New Features:

- * The user interface has been updated to improve the filtering of logs when a large number of services have been defined within the instance.
 - * The user interface input validation for hostnames and port has been improved in the setup screen. A message is now displayed when an incorrect port or address is entered.
 - * The user interface now prompts to accept a self signed certificate if one is configured.
- Bug Fix:
- * If a south or north plugin included a script type configuration item the GUI failed to allow the service or task using this plugin to be created correctly. This has now been resolved.
 - * The ability to paste into password fields has been enabled in order to allow copy/paste of keys, tokens etc into configuration of the south and north services.
 - * An issue that could result in filters not being correctly removed from a pipeline of 2 or more filters has been resolved.
- **Plugins**
- New Features:
- * A new south plugin has been added that can be used to support a number of REST based APIs. The plugin allows processing of JSON payloads or with the addition of Python scripting other payload formats may also be supported. This plugin also supports a choice of methods to control the set of readings data that will be returned.
 - * A new OPC/UA south plugin has been created based on the Safe and Secure OPC/UA library. This plugin supports authentication and encryption mechanisms.
 - * A new plugin has been added to fetch data from the Suez Water cloud API service.
 - * Control features have now been added to the modbus south plugin that allows the writing of registers and coils via the south service control channel.
 - * The modbus south control flow has been updated to use both 0x06 and 0x10 function codes. This allows items that are split across multiple modbus registers to be written in a single write operation.
 - * The MQTT Scripted south plugin has been updated to allow multiple assets to be ingested in a single plugin.
 - * The MQTT Scripted south plugin has been enhanced to support MQTTS as well as MQTT.
 - * The MQTT scripted plugin has been updated to support the return of a specific asset as well as values.
 - * The OMF plugin has been updated to support more complex scenarios for the placement of assets with the PI Asset Framework.
 - * The OMF north plugin hinting mechanism has been extended to support asset framework hierarchy hints.
 - * The OMF north plugin now defaults to using a concise naming scheme for tags in the PI server.
 - * The Kafka north plugin has been updated to allow timestamps of higher granularity than 1 second, previously timestamps would be truncated to the previous second.
 - * The Kafka north plugin has been enhanced to give the option of sending JSON objects as strings to Kafka, as previously the default, or sending them as JSON objects.
 - * The HTTP-C north plugin has been updated to allow the inclusion of customer HTTP headers.
 - * The Python35 Filter plugin did not correctly handle string type data points. This has now been resolved.
 - * The vibration velocity filter has been updated to support multiple channel data.

- * The MQTT broker package now supports RPM platforms.
- * The OMF Hint filter documentation has been updated to describe the use of regular expressions when defining the asset name to which the hint should be applied.
- * The Beckhoff south plugin documentation has been updated to include details on how to create the AMS route in a number of different scenarios.

– Bug Fix:

- * An issue with string data that had quote characters embedded within the reading data has been resolved. This would cause data to be discarded with a bad formatting message in the log.
- * An issue that could result in the configuration for the incorrect plugin being displayed has now been resolved.
- * An issue with the modbus south plugin that could cause resource starvation in the threads used for set point write operations has been resolved.
- * A race condition in the modbus south that could cause an issue if the plugin configuration is changed during a set point operation.
- * Importing the Pandas Python library into the script within the MQTT scripted plugin previously failed due to the way Pandas uses global variables. This has now been resolved such that Pandas can be imported, however it should be noted that a filter can not import Pandas if the south plugin already imports Pandas.
- * When using the South MQTT Scripted plugin, if the Python script returned an asset name as well as a reading the asset name would be corrupted on second and subsequent calls. This has now been resolved.
- * The MQTT scripted plugin would occasionally fail to shutdown cleanly. This issue has now been resolved.
- * The MQTT Scripted plugin could not previously deal with payloads that consisted of a simple negative number. This has now been corrected.
- * An issue with the MQTT notification plugin and the MQTT scripted plugin when installing with RPM packages has been resolved.
- * The CSV playback south plugin installation on CentOS 7 platforms has now been corrected.
- * The digiducer south plugin has been updated to support the latest release of the underlying libraries that support it.
- * The error handling of the OMF north plugin has been improved such that assets that contain data types that are not supported by the OMF endpoint of the PI Server are removed and other data continues to be sent to the PI Server.
- * The Kafka north plugin was not always able to reconnect if the Kafka service was not available when it was first started. This issue has now been resolved.
- * The Kafka north plugin would on occasion duplicate data if a connection failed and was later reconnected. This has been resolved.
- * A number of fixes have been made to the Kafka north plugin, these include; fixing issues caused by quoted data in the Kafka payload, sending timestamps accurate to the millisecond, fixing an issue that caused data duplication and switching the the user timestamp.
- * A problem with the quoting of string type data points on the North HTTP-C plugin has been fixed.
- * String type variables in the OPC/UA north plugin were incorrectly having extra quotes added to them. This has now been resolved.

- * The delta filter previously did not manage calculating delta values when a datapoint changed from being an integer to a floating point value or vice versa. This has now been resolved and delta values are correctly calculated when these changes occur.
- * The vibration features plugin has been updated to run on Ubuntu 20 platforms.
- * The signal processing filter plugin now installs correctly on CentOS platforms.
- * The data frames filter plugin is now supported on RPM based platforms.
- * An issue with the vibration features filter on Ubuntu 18 has been resolved.
- * The example path shown in the DHT11 plugin in the developers guide was incorrect, this has now been fixed.

16.2.2 v1.9.1

Release Date: 2021-05-27

- **FogLAMP Core**

- New Features:

- * Support has been added for Ubuntu 20.04 LTS.
- * The core components have been ported to build and run on CentOS 8
- * A new option has been added to the command line tool that controls the system. This option, called purge, allows all readings related data to be purged from the system whilst retaining the configuration. This allows a system to be tested and then reset without losing the configuration.
- * A new service interface has been added to the south service that allows set point control and operations to be performed via the south interface. This is the first phase of the set point control feature in the product.
- * The documentation has been improved to include the new control functionality in the south plugin developers guide.
- * An improvement has been made to the documentation layout for default plugins to make the GUI able to find the plugin documentation.
- * Documentation describing the installation of PostgreSQL on CentOS has been updated.
- * The documentation has been updated to give more detail around the topic of self-signed certificates.

- Bug Fix:

- * A security flaw that allowed non-privileged users to update the certificate store has been resolved.
- * A bug that prevented users being created with certificate based authentication rather than password based authentication has been fixed.
- * Switching storage plugins from SQLite to PostgreSQL caused errors in some circumstances. This has now been resolved.
- * The HTTP code returned by the ping command has been updated to correctly report 401 errors if the option to allow ping without authentication is turned off.
- * The HTTP error code returned when the notification service is not available has been corrected.
- * Disabling and re-enabling the backup and restore task schedules sometimes caused a restart of the system. This has now been resolved.
- * The error message returned when schedules could not be enabled or disabled has been improved.

- * A problem related to readings with nested data not correctly getting copied has been resolved.
- * An issue that caused problems if a service was deleted and then a new service was recreated using the name of the previously deleted service has been resolved.

- **GUI**

- New Features:

- * Links to the online help have been added on a number of screens in the user interface.
 - * Improvements have been made to the user management screens of the GUI.

- **Plugins**

- New Features:

- * North services now support Python as well as C++ plugins.
 - * A new south plugin has been created to read data from the ABB cloud service.
 - * A new south plugin has been added for getting vibration data from a set of FLIR GW65 vibration sensors.
 - * A new delivery notification plugin has been added that uses the set point control mechanism to invoke an action in the south plugin.
 - * A new notification delivery mechanism has been implemented that uses the set point control mechanism to assert control on a south service. The plugin allows you to set the values of one or more control items on the notification triggered and set a different set of values when the notification rule clears.
 - * Support has been added in the OPC/UA north plugin for array data. This allows FFT spectrum data to be represented in the OPC/UA server.
 - * The documentation for the OPC/UA north plugin has been updated to recommend running the plugin as a service.
 - * A new storage plugin has been added that uses SQLite. This is designed for situations with low bandwidth sensors and stores all the readings within a single SQLite file.
 - * The CSV Writer filter has been updated to support writing encrypted files.
 - * Support has been added to use RTSP video streams in the person detection plugin.
 - * The delta filter has been updated to allow an optional set of asset specific tolerances to be added in addition to the global tolerance used by the plugin when deciding to forward data.
 - * The Python script run by the MQTT scripted plugin now receives the topic as well as the message.
 - * The OMF plugin has been updated in line with recommendations from the OMF group regarding the use of SCRF Defense.
 - * The OMFHint plugin has been updated to support wildcarding of asset names in the rules for the plugin.
 - * New documentation has been added to help in troubleshooting PI connection issues.
 - * The pi_server and ocs north plugins are deprecated in favour of the newer and more feature rich OMF north plugin. These deprecated plugins cannot be used in north services and are only provided for backward compatibility when run as north tasks. These plugins will be removed in a future release.

- Bug Fix:

- * The OMF plugin has been updated to better deal with nested data.

- * Some improvements to error handling have been added to the InfluxDB north plugin for version 1.x of InfluxDB.
- * The Python 35 filter stated it used the Python version 3.5 always, in reality it uses whatever Python 3 version is installed on your system. The documentation has been updated to reflect this.
- * The Asset Split filter plugin previously logged debug messages by default, this has now been resolved.
- * Fixed a bug that treated arrays of bytes as if they were strings in the OPC/UA south plugin.
- * The FFT2 filter used a single asset name for all output FFT's. If an incoming asset had multiple data points they would each have a separate FFT applied to them and then output with the same asset name. This caused confusion. Now if there are multiple data points each will have a unique asset name for the output FFT. This asset name is made up of the configured output asset name with the data point name appended. For example an inout asset having X, Y and Z data points with the output asset configured to be FFT will result in 3 assets, FFTX, FFTY and FFTZ.
- * The HTTP North C plugin would not correctly shutdown, this effected reconfiguration when run as an always on service. This issue has now been resolved.
- * The description of the statistics filter was incorrect, this has now been corrected.
- * An issue with the SQLite In Memory storage plugin that caused database locks under high load conditions has been resolved.

16.2.3 v1.9.0

Release Date: 2021-02-19

- **FogLAMP Core**

- New Features:

- * Support has been added in the Python north sending process for nested JSON reading payloads.
 - * A new section has been added to the documentation to document the process of writing a notification delivery plugin. As part of this documentation a new delivery plugin has also been written which delivers notifications via an MQTT broker.
 - * The plugin developers guide has been updated with information regarding installation and debugging of new plugins.
 - * The developer documentation has been updated to include details for writing both C++ and Python filter plugins.
 - * An always on north service has been added. This compliments the current north task and allows a choice of using scheduled windows to send data north or sending data as soon as it is available.
 - * The Python north sending process required the JQ filter information to be mandatory in north plugins. JQ filtering has been deprecated and will be removed in the next major release.
 - * Storage plugins may now have configuration options that are controllable via the API and the graphical interface.
 - * The ping API call has been enhanced to return the version of the core component of the system.
 - * The SQLite storage plugin has been enhanced to distribute readings for multiple assets across multiple databases. This improves the ingest performance and also improves the responsiveness of the system when very large numbers of readings are buffered within the instance.
 - * Documentation has been added for configuration of the storage service.

- Bug Fix:

- * The REST API for the notification service was missing the re-trigger time information for configured notification in the retrieval and update calls. This has now been added.
 - * If the SQLite storage plugin is configured to use managed storage FogLAMP fails to restart. This has been resolved, the SQLite storage service no longer uses the managed option and will ignore it if set.
 - * An upgraded version of the HTTPS library has been applied, this solves an issue with large payloads in HTTPS exchanges.
 - * A number of Python source files contained incorrect references to the readthedocs page. This has now been resolved.
 - * The retrieval of log information was incorrectly including debug log output if the requested level was information and higher. This is now correctly filtered out.
 - * If a south plugin generates bad data that can not be inserted into the storage layer, that plugin will buffer the bad data forever and continually attempt to insert it. This causes the queue to build on the south plugin and eventually will exhaust system memory. To prevent this if data can not be inserted for a number of attempts it will be discarded in the south service. This allows the bad data to be dropped and newer, good data to be handled correctly.
 - * When a statistics value becomes greater than 2,147,483,648 the storage layer would fail, this has now been fixed.
 - * During installation of plugins the user interface would occasionally flag the system as down due to congestion in the API layer. This has now been resolved and the correct status of the system should be reflected.
 - * The notification service previously logged errors if no rule/delivery notification plugins had been installed. This is no longer the case.
 - * An issue with JSON configuration options that contained escaped strings within the JSON caused the service with the associated configuration to fail to run. This has now been resolved.
 - * The Postgres storage engine limited the length of asset codes to 50 characters, this has now been increased to 255 characters.
 - * Notifications based on asset names that contain the character ‘.’ in the name would not receive any data. This has now been resolved.
- Known Issues:
 - * Known issues with Postgres storage plugins. During the final testing of the 1.9.0 release a problem has been found with switching to the PostgreSQL storage plugin via the user interface. Until this is resolved switching to PostgreSQL is only supported by manual editing the storage.json as per version 1.8.0. A patch to resolve this is likely to be released in the near future.

• GUI

- New Features:
 - * The user interface now shows the retrigger time for a notification.
 - * The user interface now supports adding a north service as well as a north task.
 - * A new help menu item has been added to the user interface which will cause the readthedocs documentation to be displayed. Also the wizard to add the south and north services has been enhanced to give an option to display the help for the plugins.
- Bug Fix:
 - * The user interface now supports the ability to filter on all severity levels when viewing the system log.

• Plugins

– New Features:

- * The OPC/UA south plugin has been updated to allow the definition of the minimum reporting time between updates. It has also been updated to support subscription to arrays and DATE_TIME type with the OPC/UA server.
- * AWS SiteWise requires the SourceTimestamp to be non-null when reading from an OPC/UA server. This was not always the case with the OPC/UA north plugin and caused issues when ingesting data into SiteWise. This has now been corrected such that SourceTimestamp is correctly set in addition to server timestamp.
- * The HTTP-C north plugin has been updated to support primary and secondary destinations. It will automatically failover to the secondary if the primary becomes unavailable. Fail back will occur either when the secondary becomes unavailable or the plugin is restarted.

– Bug Fix:

- * An issue with different versions of the libmodbus library prevented the modbus-c plugin building on Moxa gateways, this has now been resolved.
- * An issue with building the MQTT notification plugin on CentOS/RedHat platforms has been resolved. This plugin now builds correctly on those platforms.
- * The modbus plugin has been enhanced to support Modbus over IPv6, also request timeout has been added as a configuration option. There have been improvements to the error handling also.
- * The DNP3 south plugin incorrectly treated all data as strings, this meant it was not easy to process the data with generic plugins. This has now been resolved and data is treated as floating point or integer values.
- * The OMF north plugin previously reported the incorrect version information. This has now been resolved.
- * A memory issue with the python35 filter integration has been resolved.
- * Packaging conflicts between plugins that used the same additional libraries have been resolved to allow both plugins to be installed on the same machine. This issue impacted the plugins that used MQTT as a transport layer.
- * The OPC/UA north plugin did not correctly handle the types for integer data, this has now been resolved.
- * The OPCUA south plugin did not allow subscriptions to integer node ids. This has now been added.
- * A problem with reading multiple modbus input registers into a single value has been resolved in the ModbusC plugin.
- * OPC/UA north nested objects did not always generate unique node IDs in the OPC/UA server. This has now been resolved.

16.2.4 v1.8.2

Release Date: 2020-11-03

- **FogLAMP Core**

– Bug Fix:

- * Following the release of a new version of a Python package the 1.8.1 release was no longer installable. This issue is resolved by the 1.8.2 patch release of the core package. All plugins from the 1.8.1 release will continue to work with the 1.8.2 release.

16.2.5 v1.8.1

Release Date: 2020-07-08

- **FogLAMP Core**

- New Features:

- * Support has been added for the deployment on Moxa gateways running a variant of Debian 9 Stretch.
- * The purge process has been improved to also purge the statistics history and audit trail of the system. New configuration parameters have been added to manage the amount of data to be retain for each of these.
- * An issue with installing on the Mendel Day release on Google's Coral boards has been resolved.
- * The REST API has been expanded to allow an API call to be made to set the repository from which new packages will be pulled when installing plugins via the API and GUI.
- * A problem with the service discovery failing to respond correctly after it had been running for a short while has been rectified. This allows external micro services to now correctly discover the core micro service.
- * Details for making contributions to the FogLAMP project have been added to the source repository.
- * The support bundle has been improved to include more information needed to diagnose issues with sending data to PI Servers
- * The REST API has been extended to add a new call that will return statistics in terms of rates rather than absolute values.
- * The documentation has been updated to include guidance on setting up package repositories for installing the software and plugins.

- Bug Fix:

- * If JSON type configuration parameters were marked as mandatory there was an issue that prevented the update of the parameters. This has now been resolved.
- * After changing storage engine from sqlite to Postgres using the configuration option in the GUI or via the API, the new storage engine would incorrectly report itself as sqlite in the API and user interface. This has now been resolved.
- * External micro-services that restarted without a graceful shutdown would fail to register with the service registry as nothing was able to unregister the failed service. This has now been relaxed to allow the recovered service to be correctly registered.
- * The configuration of the storage system was previously not available via the GUI. This has now been resolved and the configuration can be viewed in the Advanced category of the configuration user interface. Any changes made to the storage configuration will only take effect on the next restart of FogLAMP. This allows administrators to change the storage plugins used without the need to edit the storage.json configuration file.

- **GUI**

- Bug Fix:

- * An improvement to the user experience for editing password in the GUI has been implemented that stops the issue with passwords disappearing if the input field is clicked.
- * Password validation was not correctly occurring in the GUI wizard that adds south plugins. This has now be rectified.

- **Plugins**

– New Features:

- * The Modbus plugin did not gracefully handle interrupted reads of data from modes TCP devices during the bulk transfer of data. This would result in assets missing certain data points and subsequent issues in the north systems that received those assets getting changes in the asset data type. This was a particular issue when dealign with the PI Web API and would result in excessive types being created. The Modbus plugin now detects the issues and takes action to ensure complete assets are read.
- * A new image processing plugin, south human detector, that uses the Google Tensor Flow machine learning platform has been added to the FogLAMP project.
- * A new Python plugin has been added that can send data north to a Kafka system.
- * A new south plugin has been added for the Dynamic Ratings B100 Electronic Temperature Monitor used for monitoring the condition of electricity transformers.
- * A new plugin has been contributed to the project by Nexcom that implements the SAE J1708 protocol for accessing the ECU's of heavy duty vehicles.
- * An issue with missing dependencies on the Coral Mendel platform prevent 1.8.0 packages installing correctly without manual intervention. This has now been resolved.
- * The image recognition plugin, south-human-detector, has been updated to work with the Google Coral board running the Mendel Day release of Linux.

– Bug Fix:

- * A missing dependency in v1.8.0 release for the package foglamp-south-human-detector meant that it could not be installed without manual intervention. This has now been resolved.
- * Support has been added to the south-human-detector plugin for the Coral Camera module in addition to the existing support for USB connected cameras.
- * An issue with installation of the external shared libraries required by the USB4704 plugin has been resolved.

16.2.6 v1.8.0

Release Date: 2020-05-08

• FogLAMP Core

– New Features:

- * Documentation has been added for the use of the SQLite In Memory storage plugin.
- * The support bundle functionality has been improved to include more detail in order to aid tracking down issues in installations.
- * Improvements have been made to the documentation of the OMF plugin in line with the enhancements to the code. This includes the documentation of OCS and EDS support as well as PI Web API.
- * An issue with forwarding data between two FogLAMP instances in different time zones has been resolved.
- * A new API entry point has been added to the FogLAMP REST API to allow the removal of plugin packages.
- * The notification service has been updated to allow for the delivery of multiple notifications in parallel.
- * Improvements have been made to the handling of asset codes within the buffer in order to improve the ingest performance of FogLAMP. This is transparent to all services outside of the storage service and has no impact on the public APIs.

- * Extra information has been added to the notification trigger such that trigger time and the asset that triggered the notification is included.
 - * A new configuration item type of “northTask” has been introduced. It allows the user to enter the name of a northTask in the configuration of another category within FogLAMP.
 - * Data on multiple assets may now be requested in a single call to the asset growing API within FogLAMP.
 - * An additional API has been added to the asset browser to allow time bucketed data to be returned for multiple data points of multiple assets in a single call.
 - * Support has been added for nested readings within the reading data.
 - * Messages about exceeding the configured latency of the south service may be repeated when the latency is above the configured value for a period of time. These have now been replaced with a single message when the latency is exceeded and another when the condition is cleared.
 - * The feedback provided to the user when a configuration item is set to an invalid value has been improved.
 - * Configuration items can now be marked as mandatory, this improves the user experience when configuring plugins.
 - * A new configuration item type, code, has been added to improve the user experience when adding code snippets in configuration data.
 - * Improvements have been made to the caching of configuration data within the core of FogLAMP.
 - * The logging of package installation has been improved.
 - * Additions have been added to the public API to allow multiple audit log sources to be extracted in a single API call.
 - * The audit trail has been improved to show all package additions and updates in the audit trail.
 - * A new API has been added to allow notification plugin packages to be updated.
 - * A new API has been added to allow filter code versions to be updated.
 - * A new API call has been added to allow retrieval of reading data over a period of time which is averaged into time buckets within that time period.
 - * The notification service now supports rule plugins implemented in Python as well as C++.
 - * Improvements have been made to the checking of configuration items such that minimum, maximum values and string lengths are now checked.
 - * The plugin developers documentation has been updated to include a description building C/C++ south plugins.
- Bug Fix:
- * Improvements have been made to the generation of the support bundle.
 - * An issue in the reporting of the task names in the foglamp status script has been resolved.
 - * The purge by size (number of readings) would remove all data if the number of rows to retain was less than 1000, this has now been resolved.
 - * On occasions plugins would disappear from the list of available plugins, this has now been resolved.
 - * Improvements have been made to the management of the certificate store to ensure the correct files are uploaded to the store.

- * An expensive and unnecessary test was being performed in the asset browsing API of FogLAMP. This slowed down the user interface and put load on the server. This has now been removed and has improved the performance of examining the buffered data within the FogLAMP instance.
- * The FogBench utility used to send data to FogLAMP has been updated in line with new Python packages for the CoAP protocol.
- * Configuration category relationships were not always correctly cleaned up when a filter is deleted, this has now been resolved.
- * The support bundle functionality has been updated to provide information on the Python processes.
- * The REST API incorrectly allowed configuration categories with a blank name to be created. This has now been prevented.
- * Validation of minimum and maximum configuration item values was not correctly performed in the REST API, this has now been resolved.
- * Nested objects within readings could cause the storage engine to fail and those readings to not be stored. This has now been resolved.
- * On occasion shutting down a service may fail if the filters for that service have not been activated, this has now been resolved.
- * An issue that caused notifications for asset whose names contain special characters has been resolved.
- * The asset tracker was not correctly adding entries to the asset tracker, this has now been resolved.
- * An intermittent issue that prevented the notification service being enabled on the Buster release on Raspberry Pi has been resolved.
- * An intermittent problem that would prevent the north sending process to fail has been resolved.
- * Performance improvements have been made to the installation of new packages from the package repository from within the FogLAMP API and user interface.
- * It is now possible to reuse the name of a north process after deleting one with the same name.
- * The incorrect HTTP error code is returned by the asset summary API call if an asset does not exist, this has now been resolved.
- * Deleting and recreating a south service may cause errors in the log to appear. These have now been resolved.
- * The SQLite and SQLiteInMemory storage engines have been updated to enable a purge to be defined that reduces the number of readings to a specified value rather than simply allowing a purge by the age of the data. This is designed to allow tighter controls on the size of the buffer database when high frequency data in particular is being stored within the FogLAMP buffer.

• GUI

– New Features:

- * The user interface for viewing logs has been improved to allow filtering by service and task. A search facility has also been added.
- * The requirement that a key file is uploaded with every certificate file has been removed from the graphical user interface as this is not always true.
- * The performance of adding a new notification via the graphical user interface has been improved.
- * The feedback in the graphical user interface has been improved when installation of the notification service fails.

- * Installing the FogLAMP graphical user interface on OSX platforms fails due to the new version of the brew package manager. This has now been resolved.
 - * Improve script editing has been added to the graphical user interface.
 - * Improvements have been made to the user interface for the installations and enabling of the notification service.
 - * The notification audit log user interface has been improved in the GUI to allow all the logs relating to notifications to be viewed in a single screen.
 - * The user interface has been redesigned to make better use of the screen space when editing south and north services.
 - * Support has been added to the graphical user interface to determine when configuration items are not valid based on the values of other items. These items that are not valid in the current configuration are greyed out in the interface.
 - * The user interface now shows the version of the code in the settings page.
 - * Improvements have been made to the user interface layout to force footers to stay at the bottom of the screen.
- Bug Fix:
- * Improvements have been made to the zoom and pan options within the graph displays.
 - * The wizard used for the creation of new notifications in the graphical user interface would lose values when going back and forth between pages, this has now been resolved.
 - * A memory leak that was affecting the performance of the graphical user interface has been fixed, improving performance of the interface.
 - * Incorrect category names may be displayed in the graphical user interface, this has now been resolved.
 - * Issues with the layout of the graphical user interface when viewed on an Apple iPad have been resolved.
 - * The asset graph in the graphical user interface would sometimes not resize to fit the screen correctly, this has now been resolved.
 - * The “Asset & Readings” option in the graphical user interface was initially slow to respond, this has now been improved.
 - * The pagination of audit logs has been improved when multiple sources are displayed.
 - * The counts in the user interface for notifications have been corrected.
 - * Asset data graphs are not able to handle correctly the transition between one day and the next. This is now resolved.

• Plugins

- New Features:
- * The existing set of OMF north plugins have been rationalised and replaced by a single OMF north plugin that is able to support the connector relay, PI Web API, EDS and OCS.
 - * When a Modbus TCP connection is closed by the remote end we fail to read a value, we then reconnect and move on to read the next value. On device with short timeout values, smaller than the poll interval, we fail the same reading every time and never get a value for that reading. The behaviour has been modified to allow us to retry reading the original value after re-establishing the connection.
 - * The OMF north plugin has been updated to support the released version of the OSIsoft EDS product as a destination for data.

- * New functionality has been added to the north data to PI plugin when using PI Web API that allows the location in the PI Server AF hierarchy to be defined. A default location can be set and an override based on the asset name or metadata within the reading. The data may also be placed in multiple locations within the AF hierarchy.
 - * A new notification delivery plugin has been added that allows a north task to be triggered to send data for a period of time either side of the notification trigger event. This allows conditional forwarding of large amounts of data when a trigger event occurs.
 - * The asset notification delivery plugin has been updated to allow creation of new assets both for notifications that are triggered and/or cleared.
 - * The rate filter now allows the termination of sending full rate data either by use of an expression or by specifying a time in milliseconds.
 - * A new simple Python filter has been added that calculates an exponential moving average,
 - * Some typos in the OPCUA south and north plugin configuration have been fixed.
 - * The OPCUA north plugin has been updated to support nested reading objects correctly and also to allow a name to be set for the OPCUA server. These have also been some stability fixes in the underlying OPCUA layer used by this and the south OPCUA plugin.
 - * The modbus map configuration now supports byte swapping and word swapping by use of the `{{swap}}` property of the map. This may take the values `{{bytes}}`, `{{words}}` or `{{both}}`.
 - * The people detection machine learning plugin now supports RTSP streams as input.
 - * The option list items in the OMF plugin have been updated to make them more user friendly and descriptive.
 - * The threshold notification rule has been updated such that the unused fields in the configuration now correctly grey out in the GUI dependent upon the setting of the window type or single item asset validation.
 - * The configuration of the OMF north plugin for connecting to the PI Server has been improved to give the user better feedback as to what elements are valid based on choice of connection method and security options chosen.
 - * Support has been added for simple Python code to be entered into a filter that does not require all of the support code. This is designed to allow a user to very quickly develop filters with limited programming.
 - * Support has been added for filters written entirely in Python, these are full featured filters as supported by the C++ filtering mechanism and include dynamic reconfiguration.
 - * The foglamp-filter-expression filter has been modified to better deal with streams which contain multiple assets. It is now possible to use the syntax `<assetName>.<datapointName>` in an expression in addition to the previous `<datapointName>`. The result is that if two assets in the data stream have the same data point names it is now possible to differentiate between them.
 - * A new plugin to collect variables from Beckhoff PLC's has been written. The plugin uses the TwinCAT 2 or TwinCAT 3 protocols to collect specified variable from the running PLC.
- Bug Fix:
- * An issue in the sending of data to the PI server with large values has been resolved.
 - * The playback south plugin was not correctly replaying timestamps within the file, this has now been resolved.
 - * Use of the asset filter in a north task could result in the north task terminating. This has now resolved.

- * A small memory leak in the south service statistics handling code was impacting the performance of the south service, this is now resolved.
- * An issue has been discovered in the Flir camera plugin with the validity attribute of the spot temperatures, this has now been resolved.
- * It was not possible to send data for the same asset from two different FogLAMP's into the PI Server using PI Web API, this has now been resolved.
- * The filter FogLAMP RMS Trigger was not able to be dynamically reconfigured, this has now been resolved.
- * If a filter in the north sending process increased the number of readings it was possible that the limit of the number of readings sent in a single block . The sending process will now ensure this can not happen.
- * RMS filter plugin was not able to be dynamically reconfigured, this has now been resolved.
- * The HTTP South plugin that is used to receive data from another FogLAMP instance may fail with some combinations of filters applied to the service. This issue has now been resolved.
- * The rule filter may give errors if expressions have variables not satisfied in the reading data. Under some circumstances it has been seen that the filter fails to process data after giving this error. This has been resolved by changes to make the rate filter more robust.
- * Blank values for asset names in the south service may cause the service to become unresponsive. Blank asset names have now been correctly detected, asset names are required configuration values.
- * A new version of the driver software for the USB-4704 Data Acquisition Module has been released, the plugin has been updated to use this driver version.
- * The OPCUA North plugin might report incorrect counts for sent readings on some platforms, this has now been resolved.
- * The simple Python filter plugin was not adding correct asset tracking data, this has now been updated.
- * An issue with the asset filter failing when incorrect configuration was present has been resolved.
- * The benchmark plugin now enforces a minimum number of asset of 1.
- * The OPCUA plugins are now available on the Raspberry Pi Buster platform.
- * Errors that prevented the use of the Postgres storage plugin have been resolved.

16.2.7 v1.7.0

Release Date: 2019-08-15

- **FogLAMP Core**

- New Features:

- * Added support for Raspbian Buster
 - * Additional, optional flow control has been added to the south service to prevent it from overwhelming the storage service. This is enabled via the throttling option in the south service advanced configuration.
 - * The mechanism for including JSON configuration in C++ plugins has been improved and the macros for the inline coding moved to a standard location to prevent duplication.
 - * An option has been added that allows the system to be updated to the latest version of the system packages prior to installing a new plugin or component.

- * FogLAMP now supports password type configuration items. This allows passwords to be hidden from the user in the user interface.
 - * A new feature has been added that allows the logs of plugin or other package installation to be retrieved.
 - * Installation logs for package installations are now retained and available via the REST API.
 - * A mechanism has been added that allows plugins to be marked as deprecated prior to the removal of these plugins in future releases. Running a deprecated plugin will result in a warning being logged, but otherwise the plugin will operate as normal.
 - * The FogLAMP REST API has been updated to add a new entry point that will allow a plugin to be updated from the package repository.
 - * An additional API has been added to fetch the set of installed services within a FogLAMP installation.
 - * An API has been added that allows the caller to retrieve the list of plugins that are available in the FogLAMP package repository.
 - * The /foglamp/plugins REST API has been extended to allow plugins to be installed from an APT/RPM repository.
 - * Addition of support for hybrid plugins. A hybrid plugin is a JSON file that defines another plugin to load along with some default configuration for that plugin. This gives a means to create a new plugin by customising the configuration of an existing plugin. An example might be a plugin for a specific modbus device type that uses the generic modbus plugin and a predefined modbus map.
 - * The notification service has been improved to allow the re-trigger time of a notification to be defined by the user on a per notification basis.
 - * A new environment variable, FOGlamp_PLUGIN_PATH has been added to allow plugins to be stored in multiple locations or locations outside of the usual FogLAMP installation directory.
 - * Added support for FOGlamp_PLUGIN_PATH environment variable, that would be used for searching additional directory paths for plugins/filters to use with FogLAMP.
 - * FogLAMP packages for the Google Coral Edge TPU development board have been made available.
 - * Support has been added to the OMF north plugin for the PI Web API OMF endpoint. The PI Server functionality to support this is currently in beta test.
- Bug Fix/Improvements:
- * An issue with the notification service becoming unresponsive on the Raspberry Pi Buster release has been resolved.
 - * A debug message was being incorrectly logged as an error when adding a Python south plugin. The message level has now been corrected.
 - * A problem whereby not all properties of configuration items are updated when a new version of a configuration category is installed has been fixed.
 - * The notification service was not correctly honouring the notification types for one shot, toggled and retriggered notifications. This has now be bought in line with the documentation.
 - * The system log was becoming flooded with messages from the plugin discovery utility. This utility now logs at the correct level and only logs errors and warning by default.
 - * Improvements to the REST API allow for selective sets of statistic history to be retrieved. This reduces the size of the returned result set and improves performance.

- * The order in which filters are shutdown in a pipeline of filters has been reversed to resolve an issue regarding releasing Python interpreters, under some circumstances shutdowns of later filters would fail if multiple Python filters were being used.
- * The output of the *foglamp status* command was corrupt, showing random text after the number of seconds for which foglamp has been up. This has now been resolved.

- **GUI**

- New Features:

- * A new log option has been added to the GUI to show the logs of package installations.
- * It is now possible to edit Python scripts directly in the GUI for plugins that load Python snippets.
- * A new log retrieval option has been added to the GUI that will show only notification delivery events. This makes it easier for a user to see what notifications have been sent by the system.
- * The GUI asset graphs have been improved such that multiple tabs are now available for graphing and tabular display of asset data.
- * The GUI menu has been reordered to move the Notifications entry below the South and North entries.
- * Support has been added to the FogLAMP GUI for entry of password fields. Data is obfuscated as it is entered or edited.
- * The GUI now shows plugin name and version for each north task defined.
- * The GUI now shows the plugin name and version for each south service that is configured.
- * The GUI has been updated such that it can install new plugins from the FogLAMP package repository for south services and north tasks. A list of available packages from the repository is displayed to allow the user to pick from that list. The FogLAMP instance must have connectivity to the package repository to allow this feature to succeed.
- * The GUI now supports using certificates to authenticate with the FogLAMP instance.

- Bug Fix/Improvements:

- * Improved editing of JSON configuration entities in the configuration editor.
- * Improvements have been made to the asset browser graphs in the GUI to make better use of the available space to show the graph itself.
- * The GUI was incorrectly showing FogLAMP as down in certain circumstances, this has now been resolved.
- * An issue in the edit dialog for the north plugin which sometimes prevented the enabled state from being correctly modified has been resolved.
- * Exported CSV data from the GUI would sometimes be missing column headers, these are now always present.
- * The exporting of data as a CSV file in the GUI has been improved such that it no longer outputs the readings as a block of JSON, but rather individual columns. This allows the data to be imported into a spreadsheet with ease.
- * Missing help text has been added for notification trigger and enabled elements.
- * A number of issues in the filter configuration editor have been resolved. These issues meant that sometimes new values were not honoured or when changes were made with multiple filters in a chain only one filter would be updated.
- * Under some rare circumstances the GUI asset graph may show incorrect dates, this issue has now been resolved.

- * The FogLAMP GUI build and start commands did not work on Windows platforms and preventing the running on those platforms. This has now been resolved and the FogLAMP GUI can be built and run on Windows platforms.
- * The GUI was not correctly interpreting the value of the readonly attribute of configuration items when the value was anything other than true. This has been resolved.
- * The FogLAMP GUI RPM package had an error that caused installation to fail on some systems, this is now resolved.

- **Plugins**

- New Features:

- * A new filter has been created that looks for changes in values and only sends full rate data around the time of those changes. At other times the filter can be configured to send reduced rate averages of the data.
 - * A new rule plugin has been implemented that will create notifications if the value of a data point moves more than a defined percentage from the average for that data point. A moving average for each data point is calculated by the plugin, this may be a simple average or an exponential moving average.
 - * A new south plugin has been created that supports the DNP3 protocol.
 - * A south plugin has been created based on the Google TensorFlow people detection model. It uses a live feed from a video camera and returns data regarding the number of people detected and the position within the frame.
 - * A south plugin based on the Google TensorFlow demo model for people recognition has been created. The plugin reads an image from a file and returns the people co-ordinates of the people it detects within the image.
 - * A new north plugin has been added that creates an OPCUA server based on the data ingested by the FogLAMP instance.
 - * Support has been added for a Flir Thermal Imaging Camera connected via Modbus TCP. Both a south plugin to gather the data and a filter plugin, to clean the data, have been added.
 - * A new south plugin has been created based on the Google TensorFlow demo model that accepts a live feed from a Raspberry Pi camera and classifies the images.
 - * A new south plugin has been created based on the Google TensorFlow demo model for object detection. The plugin return object count, name position and confidence data.
 - * The change filter has been made available on CentOS and RedHat 7 releases.

- Bug Fix/Improvements:

- * Support for reading floating point values in a pair of 16 bit registers has been added to the modbus plugin.
 - * Improvements have been made to the performance of the modbus plugin when large numbers of contiguous registers are read. Also the addition of support for floating point values in modbus registers.
 - * Flir south service has been modified to support the Flir camera range as currently available, i.e. a maximum of 10 areas as opposed to the 20 that were previously supported. This has improved performance, especially on low performance platforms.
 - * The python35 filter plugin did not allow the Python code to add attributes to the data. This has now been resolved.
 - * The playback south plugin did not correctly take the timestamp data from the CSV file. An option is now available that will allow this.

- * The rate filter has been enhanced to accept a list of assets that should be passed through the filter without having the rate of those assets altered.
- * The filter plugin python35 crashed on the Buster release on the Raspberry Pi, this has now been resolved.
- * The FFT filter now enforces that the number of samples must be a power of 2.
- * The ThingSpeak north plugin was not updated in line with changes to the timestamp handling in FogLAMP, this resulted in a crash when it tried to send data to ThingSpeak. This has been resolved and the cause of the crash also fixed such that now an error will be logged rather than the task crashing.
- * The configuration of the simple expression notification rule plugin has been simplified.
- * The DHT 11 plugin mistakenly had a dependency on the Wiring PI package. This has now been removed.
- * The system information plugin was missing a dependency that would cause it to fail to install on systems that did not already have the package it was depend on installed. This has been resolved.
- * The phidget south plugin reconfiguration method would crash the service on occasions, this has now been resolved.
- * The notification service would sometimes become unresponsive after calling the notify-python35 plugin, this has now been resolved.
- * The configuration options regarding notification evaluation of single items and windows has been improved to make it less confusing to end users.
- * The OverMax and UnderMin notification rules have been combined into a single threshold rule plugin.
- * The OPCUA south plugin was incorrectly reporting itself as the upcua plugin. This is now resolved.
- * The OPCUA south plugin has been updated to support subscriptions both using browse names and Node Id's. Node ID is now the default subscription mechanism as this is much higher performance than traversing the object tree looking at browse names.
- * Shutting down the OPCUA service when it has failed to connect to an OPCUA server, either because of an incorrect configuration or the OPCUA server being down resulted in the service crashing. The service now shuts down cleanly.
- * In order to install the foglamp-south-modbus package on RedHat Enterprise Linux or CentOS 7 you must have configured the epel repository by executing the command:

sudo yum install epel-release

- * A number of packages have been renamed in order to obtain better consistency in the naming and to facilitate the upgrade of packages from the API and graphical interface to FogLAMP. This will result in duplication of certain plugins after upgrading to the release. This is only an issue of the plugins had been previously installed, these old plugin should be manually removed form the system to alleviate this problem.

The plugins involved are,

- foglamp-north-http Vs foglamp-north-http-north
- foglamp-south-http Vs foglamp-south-http-south
- foglamp-south-Csv Vs foglamp-south-csv
- foglamp-south-Expression Vs foglamp-south-expression
- foglamp-south-dht Vs foglamp-south-dht11V2
- foglamp-south-modbusc Vs foglamp-south-modbus

16.2.8 v1.6.0

Release Date: 2019-05-22

- **FogLAMP Core**

- New Features:

- * The scope of the FogLAMP certificate store has been widened to allow it to store .pem certificates and keys for accessing cloud functions.
 - * The creation of a Docker container for FogLAMP has been added to the packaging options for FogLAMP in this version of FogLAMP.
 - * Red Hat Enterprise Linux packages have been made available from this release of FogLAMP onwards. These packages include all the applicable plugins and notification service for FogLAMP.
 - * The FogLAMP API now supports the creation of configuration snapshots which can be used to create configuration checkpoints and rollback configuration changes.
 - * The FogLAMP administration API has been extended to allow the installation of new plugins via API.

- Improvements/Bug Fix:

- * A bug that prevents multiple FogLAMP's on the same network being discoverable via multicast DNS lookup has been fixed.
 - * Set, unset optional configuration attributes

- **GUI**

- New Features:

- * The FogLAMP Graphical User Interface now has the ability to show sets of graphs over a time period for data such as the spectrum analysis produced but the Fast Fourier transform filter.
 - * The FogLAMP Graphical User Interface is now available as an RPM file that may be installed on Red Hat Enterprise Linux or CentOS.

- Improvements/Bug Fix:

- * Improvements have been made to the FogLAMP Graphical User Interface to allow more control of the time periods displayed in the graphs of asset values.
 - * Some improvements to screen layout in the FogLAMP Graphical User Interface have been made in order to improve the look and reduce the screen space used in some of the screens.
 - * Improvements have been made to the appearance of dropdown and other elements with the FogLAMP Graphical User Interface.

- **Plugins**

- New Features:

- * A new threshold filter has been added that can be used to block onward transmission of data until a configured expression evaluates too true.
 - * The Modbus RTU/TCP south plugin is now available on CentOS 7 and RHEL 7.
 - * A new north plugin has been added to allow data to be sent the Google Cloud Platform IoT Core interface.
 - * The FFT filter now has an option to output raw frequency spectra. Note this can not be accepted into all north bound systems.
 - * Changed the release status of the FFT filter plugin.

- * Added the ability in the modbus plugin to define multiple registers that create composite values. For example two 16 bit registers can be put together to make one 32 bit value. This is done using an array of register values in a modbus map, e.g. {"name": "rpm", "slave": 1, "register": [33, 34], "scale": 0.1, "offset": 0}. Register 33 contains the low 16 bits of the RPM and register 34 the high 16 bits of the RPM.
- * Addition of a new Notification Delivery plugin to send notifications to a Google Hangouts chat-room.
- * A new plugin has been created that uses machine learning based on Google's TensorFlow technology to classify image data and populate derived information the north side systems. The current TensorFlow model in use will recognise hand written digits and populate those digits. This plugin is currently a proof of concept for machine learning.

– Improvements/Bug Fix:

- * Removal of unnecessary include directive from Modbus-C plugin.
- * Improved error reporting for the modbus-c plugin and added documentation on the configuration of the plugin.
- * Improved the subscription handling in the OPCUA south plugin.
- * Stability improvements have been made to the notification service, these related to the handling of dynamic reconfigurations of the notifications.
- * Removed erroneous default for script configuration option in Python35 notification delivery plugin.
- * Corrected description of the enable configuration item.

16.2.9 v1.5.2

Release Date: 2019-04-08

• FogLAMP Core

– New Features:

- * Notification service, notification rule and delivery plugins
- * Addition of a new notification delivery plugin that will create an asset reading when a notification is delivered. This can then be sent to any system north of the FogLAMP instance via the usual mechanisms
- * Bulk insert support for SQLite and Postgres storage plugins

– Enhancements / Bug Fix:

- * Performance improvements for SQLite storage plugin.
- * Improved performance of data browsing where large datasets have been acquired
- * Optimized statistics history collection
- * Optimized purge task
- * The readings count shown on GUI and south page and corresponding API endpoints now shows total readings count and not what is currently buffered by FogLAMP. So these counts don't reduce when purge task runs
- * Static data in the OMF plugin was not being correctly taken from the plugin configuration
- * Reduced the number of informational log messages being sent to the syslog

- **GUI**
 - **New Features:**
 - * Notifications UI
 - **Bug Fix:**
 - * Backup creation time format

16.2.10 v1.5.1

Release Date: 2019-03-12

- **FogLAMP Core**
 - Bug Fix: plugin loading errors
- **GUI**
 - Bug Fix: uptime shows up to 24 hour clock only

16.2.11 v1.5.0

Release Date: 2019-02-21

- **FogLAMP Core**
 - Performance improvements and Bug Fixes
 - Introduction of Safe Mode in case FogLAMP is accidentally configured to generate so much data that it is overwhelmed and can no longer be managed.
- **GUI**
 - re-organization of screens for Health, Assets, South and North
 - bug fixes
- **South**
 - Many Performance improvements, including conversion to C++
 - Modbus plugin
 - many other new south plugins
- **North**
 - Compressed data via OMF
 - Kafka
- **Filters:** Perform data pre-processing, and allow distributed applications to be built on FogLAMP.
 - Delta: only send data upon change
 - Expression: run a complex mathematical expression across one or more data streams
 - Python: run arbitrary python code to modify a data stream
 - Asset: modify Asset metadata
 - RMS: Generate new asset with Root Mean Squared and Peak calculations across data streams
 - FFT (beta): execute a Fast Fourier Transform across a data stream. Valuable for Vibration Analysis

- Many others
- **Event Notification Engine (beta)**
 - Run rules to detect conditions and generate events at the edge
 - Default Delivery Mechanisms: email, external script
 - Fully pluggable, so custom Rules and Delivery Mechanisms can be easily created
- **Debian Packages for All Repo's**

16.2.12 v1.4.1

Release Date: 2018-10-10

16.2.13 v1.4.0

Release Date: 2018-09-25

16.2.14 v1.3.1

Release Date: 2018-07-13

Fixed Issues

- **Open File Descriptors**
 - **open file descriptors:** Storage service did not close open files, leading to multiple open file descriptors

16.2.15 v1.3

Release Date: 2018-07-05

New Features

- **Python version upgrade**
 - **python 3 version:** The minimal supported python version is now python 3.5.3.
- **aiohttp python package version upgrade**
 - **aiohttp package version:** aiohttp (version 3.2.1) and aiohttp_cors (version 0.7.0) is now being used
- **Removal of south plugins**
 - **coap:** coap south plugin was moved into its own repository <https://github.com/foglamp/foglamp-south-coap>
 - **http:** http south plugin was moved into its own repository <https://github.com/foglamp/foglamp-south-http>

Known Issues

- **Issues in Documentation**
 - **plugin documentation:** testing FogLAMP requires user to first install southbound plugins necessary (CoAP, http)

16.2.16 v1.2

Release Date: 2018-04-23

New Features

- **Changes in the REST API**
 - **ping Method:** the ping method now returns uptime, number of records read/sent/purged and if FogLAMP requires REST API authentication.
- **Storage Layer**
 - **Default Storage Engine:** The default storage engine is now SQLite. We provide a script to migrate from PostgreSQL in 1.1.1 version to 1.2. PostgreSQL is still available in the main repository and package, but it will be removed to an operate repository in future versions.
- **Admin and Maintenance Scripts**
 - **foglamp status:** the command now shows what the ping REST method provides.
 - **setenv script:** a new script has been added to simplify the user interaction. The script is in *\$FOGLAMP_ROOT/extras/scripts* and it is called *setenv.sh*.
 - **foglamp service script:** a new service script has been added to setup FogLAMP as a service. The script is in *\$FOGLAMP_ROOT/extras/scripts* and it is called *foglamp.service*.

Known Issues

- **Issues in the REST API**
 - **asset method response:** the asset method returns a JSON object with asset code named `asset_code` instead of `assetCode`
 - **task method response:** the task method returns a JSON object with unexpected element `"exitCode"`

16.2.17 v1.1.1

Release Date: 2018-01-18

New Features

- **Fixed aiohttp incompatibility:** This fix is for the incompatibility of *aiohttp* with *yarl*, discovered in the previous version. The issue has been fixed.
- **Fixed avahi-daemon issue:** Avahi daemon is a pre-requisite of FogLAMP, FogLAMP can now run as a snap or build from source without avahi daemon installed.

Known Issues

- **PostgreSQL with Snap:** the issue described in version 1.0 still persists, see [Known Issues](#) in v1.0.

16.2.18 v1.1

Release Date: 2018-01-09

New Features

- **Startup Script:**
 - `foglamp start` script now checks if the Core microservice has started.
 - `foglamp start` creates a `core.err` file in `$FOGLAMP_DATA` and writes the stderr there.

Known Issues

- **Incompatibility between aiohttp and yarl when FogLAMP is built from source:** in this version we use *aiohttp* 2.3.6 (). This version is incompatible with updated versions of *yarl* (0.18.0+). If you intend to use this version, change the requirements for *aiohttp* for version 2.3.8 or higher.
- **PostgreSQL with Snap:** the issue described in version 1.0 still persists, see [Known Issues](#) in v1.0.

16.2.19 v1.0

Release Date: 2017-12-11

Features

- All the essential microservices are now in place: *Core*, *Storage*, *South*, *North*.
- Storage plugins available in the main repository:
 - **Postgres:** The storage layer relies on PostgreSQL for data and metadata
- South plugins available in the main repository:
 - **CoAP Listener:** A CoAP microservice plugin listening to client applications that send data to FogLAMP
- North plugins available in the main repository:
 - **OMF Translator:** A task plugin sending data to OSIsoft PI Connector Relay 1.0

Known Issues

- **Startup Script:** `foglamp start` does not check if the Core microservice has started correctly, hence it may report that “FogLAMP started.” when the process has died. As a workaround, check with `foglamp status` the presence of the FogLAMP microservices.
- **Snap Execution on Raspbian:** there is an issue on Raspbian when the FogLAMP snap package is used. It is an issue with the snap environment, it looks for a shared object to preload on Raspbian, but the object is not available. As a workaround, a superuser should comment a line in the file `/etc/ld.so.preload`. Add a `#` at the beginning of this line: `/usr/lib/arm-linux-gnueabi/libarmmem.so`. Save the file and you will be able to immediately use the snap.
- **OMF Translator North Plugin for FogLAMP Statistics:** in this version the statistics collected by FogLAMP are not sent automatically to the PI System via the OMF Translator plugin, as it is supposed to be. The issue will be fixed in a future release.
- **Snap installed in an environment with an existing version of PostgreSQL:** the FogLAMP snap does not check if another version of PostgreSQL is available on the machine. The result may be a conflict between the tailored version of PostgreSQL installed with the snap and the version of PostgreSQL generally available on the machine. You can check if PostgreSQL is installed using the command `sudo dpkg -l | grep 'postgres'`. All packages should be removed with `sudo dpkg --purge <package>`.

DOWNLOADS

17.1 Packages

Packages for a number of different Linux platforms are available for both Intel and Arm architectures via the Dianomic web site's download page.

-

OMF KERBEROS AUTHENTICATION

18.1 Introduction

The bundled OMF north plugin in FogLAMP can use a number of different authentication schemes when communicating with the various OSIssoft products. The PI Web API method in the OMF plugin supports the use of a Kerberos scheme.

The FogLAMP *requirements.sh* script installs the Kerberos client to allow the integration with what in the specific terminology is called KDC (the Kerberos server).

18.2 PI Server as the North endpoint

The OSI *Connector Relay* allows token authentication while *PI Web API* supports Basic and Kerberos authentication.

There could be more than one configuration to allow the Kerberos authentication, the easiest one is the Windows server on which the PI Server is executed act as the Kerberos server also.

The Windows Active directory should be installed and properly configured for allowing the Windows server to authenticate Kerberos requests.

18.3 North plugin

The North plugin has a set of configurable options that should be changed, using either the FogLAMP API or the FogLAMP GUI, to select the Kerberos authentication.

The North plugin supports the configurable option *PIServerEndpoint* for allowing to select the target among:

- Connector Relay
- PI Web API
- Edge Data Store
- OSIssoft Cloud Services

The *PIWebAPIAuthenticationMethod* option permits to select the desired authentication among:

- anonymous
- basic
- kerberos

The Kerberos authentication requires a keytab file, the *PIWebAPIKerberosKeytabFileName* option specifies the name of the file expected under the directory:

```
${FOGLAMP_ROOT}/data/etc/kerberos
```

NOTE:

- *A keytab is a file containing pairs of Kerberos principals and encrypted keys (which are derived from the Kerberos password). A keytab file allows to authenticate to various remote systems using Kerberos without entering a password.*

the *AFHierarchyLevel* option allows to specific the first level of the hierarchy that will be created into the Asset Framework and will contain the information for the specific North plugin.

18.4 FogLAMP server configuration

The server on which FogLAMP is going to be executed needs to be properly configured to allow the Kerberos authentication.

The following steps are needed:

- *IP Address resolution for the KDC*
- *Kerberos client configuration*
- *Kerberos keytab file setup*

18.4.1 IP Address resolution of the KDC

The Kerberos server name should be resolved to the corresponding IP Address, editing the */etc/hosts* is one of the possible and the easiest way, sample row to add:

```
192.168.1.51    pi-server.dianomic.com pi-server
```

try the resolution of the name using the usual *ping* command:

```
$ ping -c 1 pi-server.dianomic.com

PING pi-server.dianomic.com (192.168.1.51) 56(84) bytes of data.
64 bytes from pi-server.dianomic.com (192.168.1.51): icmp_seq=1 ttl=128 time=0.317 ms
64 bytes from pi-server.dianomic.com (192.168.1.51): icmp_seq=2 ttl=128 time=0.360 ms
64 bytes from pi-server.dianomic.com (192.168.1.51): icmp_seq=3 ttl=128 time=0.455 ms
```

NOTE:

- *the name of the KDC should be the first in the list of aliases*

18.4.2 Kerberos client configuration

The server on which FogLAMP runs act like a Kerberos client and the related configuration file should be edited for allowing the proper Kerberos server identification. The information should be added into the `/etc/krb5.conf` file in the corresponding section, for example:

```
[libdefaults]
    default_realm = DIANOMIC.COM

[realms]
    DIANOMIC.COM = {
        kdc = pi-server.dianomic.com
        admin_server = pi-server.dianomic.com
    }
```

18.4.3 Kerberos keytab file

The keytab file should be generated on the Kerberos server and copied into the FogLAMP server in the directory:

```
${FOGLAMP_DATA}/etc/kerberos
```

NOTE:

- if **FOGLAMP_DATA** is not set its value should be `$FOGLAMP_ROOT/data`.

The name of the file should match the value of the North plugin option `PIWebAPIKerberosKeytabFileName`, by default `piwebapi_kerberos_https.keytab`

```
$ ls -l ${FOGLAMP_DATA}/etc/kerberos
-rwxrwxrwx 1 foglamp foglamp 91 Jul 17 09:07 piwebapi_kerberos_https.keytab
-rw-rw-r-- 1 foglamp foglamp 199 Aug 13 15:30 README.rst
```

The way the keytab file is generated depends on the type of the Kerberos server, in the case of Windows Active Directory this is an sample command:

```
ktpass -princ HTTPS/pi-server@DIANOMIC.COM -mapuser Administrator@DIANOMIC.COM -pass_
↪Password -crypto AES256-SHA1 -ptype KRB5_NT_PRINCIPAL -out C:\Temp\piwebapi_
↪kerberos_https.keytab
```

18.4.4 Troubleshooting the Kerberos authentication

- 1) check the North plugin configuration, a sample command

```
curl -s -S -X GET http://localhost:8081/foglamp/category/North_Readings_to_PI | jq ". |
↪{URL, \"PIServerEndpoint\", PIWebAPIAuthenticationMethod, PIWebAPIKerberosKeytabFileName,
↪AFHierarchyLevel}\""
```

- 2) check the presence of the keytab file

```
$ ls -l ${FOGLAMP_ROOT}/data/etc/kerberos
-rwxrwxrwx 1 foglamp foglamp 91 Jul 17 09:07 piwebapi_kerberos_https.keytab
-rw-rw-r-- 1 foglamp foglamp 199 Aug 13 15:30 README.rst
```

- 3) verify the reachability of the Kerberos server (usually the PI Server) - Network reachability

```
$ ping pi-server.dianomic.com
PING pi-server.dianomic.com (192.168.1.51) 56(84) bytes of data.
64 bytes from pi-server.dianomic.com (192.168.1.51): icmp_seq=1 ttl=128 time=5.07 ms
64 bytes from pi-server.dianomic.com (192.168.1.51): icmp_seq=2 ttl=128 time=1.92 ms
```

Kerberos reachability and keys retrieval

```
$ kinit -p HTTPS/pi-server@DIANOMIC.COM
Password for HTTPS/pi-server@DIANOMIC.COM:
$ klist
Ticket cache: FILE:/tmp/krb5cc_1001
Default principal: HTTPS/pi-server@DIANOMIC.COM

Valid starting      Expires            Service principal
09/27/2019 11:51:47  09/27/2019 21:51:47  krbtgt/DIANOMIC.COM@DIANOMIC.COM
    renew until 09/28/2019 11:51:46
$
```

FOGLAMP PLUGINS

The following set of plugins are available for FogLAMP. These plugins extend the functionality by adding new sources of data, new destinations, processing filters that can enhance or modify the data, rules for notification delivery and notification delivery mechanisms.

19.1 South Plugins

South plugins add new ways to get data into FogLAMP, a number of south plugins are available ready built or users may add new south plugins of their own by writing them in Python or C/C++.

Table 1: FogLAMP South Plugins

Name	Description
abb	A south plugin to pull data from the ABB cloud
am2315	FogLAMP south plugin for an AM2315 temperature and humidity sensor
b100-modbus-python	A south plugin to read data from a Dynamic Ratings B100 device over Modbus
beckhoff	A Beckhoff ADS data ingress plugin for FogLAMP, this monitors Beckhoff PLCs and returns the state of internal variables within the PLC
benchmark	A FogLAMP benchmark plugin to measure the ingestion rates on particular hardware
cc2650	A FogLAMP south plugin for the Texas Instruments SensorTag CC2650
coap	A south plugin for FogLAMP that pulls data from a COAP sensor
coin-detection	A coin-detection south plugin
coral-enviro	A south plugin for the Google Coral Environmental Sensor Board
csv	A FogLAMP south plugin in C++ for reading CSV files
csv-async	A FogLAMP asynchronous plugin for reading CSV data
csvplayback	Plays a CSV at some configurable speed and each column of the file will become a datapoint of an asset using pandas library.
dht	A FogLAMP south plugin in C++ that interfaces to a DHT-11 temperature and humidity sensor
dht11	A FogLAMP south plugin that interfaces a DHT-11 temperature sensor
digiducer	South plugin for the Digiducer 333D01 vibration sensor
dnp3	A south plugin for FogLAMP that implements the DNP3 protocol
dt9837	A south plugin for the Data Translation DT9837 Series DAQ
edgeml	ML south plugin which forwards the video frames to a model running inside micro k8's; parses the response, generates readings and shows the detection results on browser.
etherip	A south plugin to read tags data from a number of different Allen-Bradley and Rockwell PLCs.

continues on next page

Table 1 – continued from previous page

Name	Description
expression	A FogLAMP south plugin that uses a user define expression to generate data
FlirAX8	A FogLAMP hybrid south plugin that uses foglamp-south-modbus-c to get temperature data from a Flir Thermal camera
game	The south plugin used for the FogLAMP lab session game involving remote controlled cars
gw65	FogLAMP plugin for getting vibration data from a set of FLIR GW65 vibration sensors
http	A Python south plugin for FogLAMP used to connect one FogLAMP instance to another
ina219	A FogLAMP south plugin for the INA219 voltage and current sensor
J1708	A plugin that uses the SAE J1708 protocol to load data from the ECU of heavy duty vehicles.
J1939	A CANBUS J1839 plugin to collect data into FogLAMP.
lathesim	A simulation plugin used as a demonstration to show how data can be collected within FogLAMP. This plugin simulates various properties of a lathe.
modbus-c	A FogLAMP south plugin that implements modbus-tcp and modbus-rtu
modbustcp	A FogLAMP south plugin that implements modbus-tcp in Python
mqtt	FogLAMP South MQTT Subscriber Plugin
mqtt-sparkplug	A FogLAMP south plugin that implements the Sparkplug API over MQTT
mqtt-scripted	An MQTT south plugin that allows a Python script to be added to decode the MQTT payload
opcua	A FogLAMP south service that pulls data from an OPC-UA server
openweathermap	A FogLAMP south plugin to pull weather data from OpenWeatherMap
person-detection	FogLAMP south service plugin that detects person in the live video stream
person-detector	Person detection plugin
phidget	FogLAMP south code for different phidgets
piwebapi	A South plugin to ingest data from a PI Server using the PI Web API.
playback	A FogLAMP south plugin to replay data stored in a CSV file
pt100	A FogLAMP south plugin for the PT100 temperature sensor
random	A south plugin for FogLAMP that generates random numbers
randomwalk	A FogLAMP south plugin that returns data that with randomly generated steps
roxtec	A FogLAMP south plugin for the Roxtec cable gland project
rpienviro	A FogLAMP south service for the Raspberry Pi EnviroPhat sensors
s2opcua	An OPCUA south plugin based on the Safe & Secure OPCUA library. This plugin offers similar functionality to the foglamp-south-opcua plugin but also offers encryption and authentication.
s7-python	foglamp-south-s7-python repository
s7	A south plugin that uses the S7 Communications protocol to read data from a Siemens S7 series PLC.
samotics4	A south plugin that will gather data from Samotics 4 API used to monitor AC Motors
sarcos	A south plugin to process the Sarcos XO data files
sensehat	A FogLAMP south plugin for the Raspberry Pi Sensehat sensors
sensorphone	A FogLAMP south plugin the task to the iPhone SensorPhone app
simple-rest	A generic REST south plugin with support for a variety of common rest payloads and Python scripting to manipulate call results.
sinusoid	A FogLAMP south plugin that produces a simulated sine wave
spinnaker	A south plugin that uses the FLIR Spinnaker library to pull radiometric data from FLIR cameras for use within machine learning models
suez	A south plugin to extract data from the Suez Water Insight API

continues on next page

Table 1 – continued from previous page

Name	Description
systeminfo	A FogLAMP south plugin that gathers information about the system it is running on.
usb4704	A FogLAMP south plugin the Advantech USB-4704 data acquisition module
video4linux	A south plugin to ingests images from various devices using the Video4Linux API. Video4Linux supports a wide variety of video capture devices on Linux platforms.
webcam-media	A FogLAMP south plugin that forwards image data, either directly from a webcam or from a directory of images
wind-turbine	A FogLAMP south plugin for a number of sensor connected to a wind turbine demo

19.2 North Plugins

North plugins add new destinations to which data may be sent by FogLAMP. A number of north plugins are available ready built or users may add new north plugins of their own by writing them in Python or C/C++.

Table 2: FogLAMP North Plugins

Name	Description
azure	A north plugin that sends data to Microsoft Azure IoT Core.
gcp	A north plugin to send data to Google Cloud Platform IoT Core
gcp-ps	FogLAMP North Python based gcp plugin for sending an image data to Google Cloud for training machine learning models.
graphite	A north plugin for FogLAMP that sends data to the Graphite Carbon storage system.
harperdb	A north plugin that sends data to the HarperDB SQL/NoSQL data management platform
http	A Python implementation of a north plugin to send data between FogLAMP instances using HTTP
http-c	A FogLAMP north plugin that sends data between FogLAMP instances using HTTP/HTTPS
influxdb	A north plugin for sending data to InfluxDB
influxdbcloud	A north plugin to send data from FogLAMP to the InfluxDBCloud
kafka	A FogLAMP plugin for sending data north to Apache Kafka
kafka-python	A Python implementation of a north plugin that can send data to Apache Kafka
mqtt-scripted	A mqtt-scripted north plugin
png	A plugin to write an image type data points to PNG files in the local filesystem
s7-python	foglamp-north-s7-python repository
splunk	A north plugin for sending data to Splunk
thingspeak	A FogLAMP north plugin to send data to Matlab's ThingSpeak cloud

19.3 Filter Plugins

Filter plugins add new ways in which data may be modified, enhanced or cleaned as part of the ingress via a south service or egress to a destination system. A number of north plugins are available ready built or users may add new north plugins of their own by writing them in Python or C/C++.

It is also possible, using particular filters, to supply expressions or script snippets that can operate on the data as well. This provides a simple way to process the data in FogLAMP as it is read from devices or written to destination systems.

Table 3: FogLAMP Filter Plugins

Name	Description
ADM-LD-prediction	Filter to detect whether a large discharge is required for an ADM centrifuge
asset	A FogLAMP processing filter that is used to block or allow certain assets to pass onwards in the data stream
asset-split	A filter to split an asset with multiple data points into several assets, each with a single data point.
blocktest	A filter designed to aid testing. It combines incoming readings into bigger blocks before sending onwards
change	A FogLAMP processing filter plugin that only forwards data that changes by more than a configurable amount
contrast	A filter that implements automatic and manual contrast adjustment of images.
csv-writer	FogLAMP filter which writes selected readings passing through it out as a rotating sequence of .csv files.
delta	A FogLAMP processing filter plugin that removes duplicates from the stream of data and only forwards new values that differ from previous values by more than a given tolerance
downsample	A data downsampling filter which may be used to reduce the data rate using sampling or averaging techniques.
edgeml	Filter which takes image data, calls out to ML process, and forwards the inference from ML as asset contents.
ema	Generate exponential moving average datapoint: include a rate of current value and a rate of history values
eventrate	A filter designed for use in the north to trigger sending rates based on event notification assets
expression	A FogLAMP processing filter plugin that applies a user define formula to the data as it passes through the filter
fft	A FogLAMP processing filter plugin that calculates a Fast Fourier Transform across sensor data
fft2	Filter for FFT signal processing, finding peak frequencies, etc.
Flir-Validity	A FogLAMP processing filter used for processing temperature data from a Flir thermal camera
greyscale	Convert 24bit RGB images to greyscale images
log	A FogLAMP filter that converts the readings data to a logarithmic scale. This is the example filter used in the plugin developers guide.
metadata	A FogLAMP processing filter plugin that adds metadata to the readings in the data stream
mirror	A filter plugin to mirror image type data points
omfhint	A filter plugin that allows data to be added to assets that will provide extra information to the OMF north plugin.
python27	A FogLAMP processing filter that allows Python 2 code to be run on each sensor value.

continues on next page

Table 3 – continued from previous page

Name	Description
python35	A FogLAMP processing filter that allows Python 3 code to be run on each sensor value.
rate	A FogLAMP processing filter plugin that sends reduced rate data until an expression triggers sending full rate data
rename	A FogLAMP processing filter that is used to modify the name of an asset, datapoint or with both
replace	Filter to replace characters in the names of assets and data points in readings object.
rms	A FogLAMP processing filter plugin that calculates RMS value for sensor data
rms-trigger	An RMS filter that uses a trigger asset rather than a fixed set of readings for each calculation
rotate	Rotate all images found in datapoints within a reading
scale	A FogLAMP processing filter plugin that applies an offset and scale factor to the data
scale-set	A FogLAMP processing filter plugin that applies a set of scale factors to the data
sigfns	Signal processing functions
sigmacleanse	A data cleansing plugin that removes data that differs from the mean value by more than x sigma
simple-python	The simple Python filter plugin is analogous to the expression filter but accept Python code rather than the expression syntax
specgram	FogLAMP filter to generate spectrogram images for vibration data
statistics	Generic statistics filter for FogLAMP data that supports the generation of mean, mode, median, minimum, maximum, standard deviation and variance.
threshold	A FogLAMP processing filter that only forwards data when a threshold is crossed
velocity	Filter to process acceleration data to generate velocity and acceleration envelope
vibration_features	A filter plugin that takes a stream of vibration data and generates a set of features that characterise that data

19.4 Notification Rule Plugins

Notification rule plugins provide the logic that is used by the notification service to determine if a condition has been met that should trigger or clear that condition and hence send a notification. A number of notification plugins are available as standard, however as with any plugin the user is able to write new plugins in Python or C/C++ to extend the set of notification rules.

Table 4: FogLAMP Notification Rule Plugins

Name	Description
average	A FogLAMP notification rule plugin that evaluates an expression based sensor data notification rule plugin that triggers when sensors values depart from the moving average by more than a configured limit.
delta	A delta rule plugin
ML-bad-bearing	Notification rule plugin to detect bad bearing
outofbound	A FogLAMP notification rule plugin that triggers when sensors values exceed limits set in the configuration of the plugin.
periodic	A rule that periodically fires based on a timer when data is observed.
simple-expression	A FogLAMP notification rule plugin that evaluates an expression based sensor data
simple-sigma	A FogLAMP notification rule that will send a notification if the values being monitored differ from the mean for the value by more than a multiple of the current standard deviation.
watchdog	Notification rule designed to be triggered if data for a given asset is not ingested for a period of time.

19.5 Notification Delivery Plugins

Notification delivery plugins provide the mechanisms to deliver the notification messages to the systems that will receive them. A number of notification delivery plugins are available as standard, however as with any plugin the user is able to write new plugins in Python or C/C++ to extend the set of notification rules.

Table 5: FogLAMP Notification Delivery Plugins

Name	Description
alexa-notifyme	A FogLAMP notification delivery plugin that sends notifications to the Amazon Alexa platform
asset	A FogLAMP notification delivery plugin that creates an asset in FogLAMP when a notification occurs
blynk	A FogLAMP notification delivery plugin that sends notifications to the Blynk service
config	A notification delivery plugin that allows a configuration item within the local FogLAMP instance to be changed when the notification triggers or is cleared.
control	A control notify plugin
customasset	A FogLAMP notification delivery plugin that creates an event asset in readings.
email	A FogLAMP notification delivery plugin that sends notifications via email
google-hangouts	A FogLAMP notification delivery plugin that sends alerts on the Google hangout platform
ifttt	A FogLAMP notification delivery plugin that triggers an action of IFTTT
jira	A notification plugin that creates tickets in Jira
jsonconfig	A delivery mechanism that updates one element within a JSON configuration type configuration category item.
management	A notification delivery plugin the triggers the FogLAMP management service to check for updates to the configuration of FogLAMP
mqtt	A notification delivery plugin that sends messages via MQTT when a notification is triggered or cleared. This is the example used in the notification delivery plugin writers guide.
north	Deliver notification data via a FogLAMP north task
operation	A notification delivery plugin that will cause an operation to be trigger via the set point control operation API of a south service.
python35	A FogLAMP notification delivery plugin that runs an arbitrary Python 3 script
setpoint	A foglamp notification plugin that invokes a set point operation on a south service.
slack	A FogLAMP notification delivery plugin that sends notifications via the slack instant messaging platform
telegram	A FogLAMP notification delivery plugin that sends notifications via the telegram service
zendesk	A notification delivery plugin that will create tickets within Zendesk ticketing application

GLOSSARY

The following are a set of definitions for terms used within the FogLAMP documentation and code, these are designed to be an aid to understanding some of the principles behind FogLAMP and improve the comprehension of the documentation by ensuring all readers have a common understanding of the terms used.

Asset A representation of a set of device or set of values about a device or entity that is being monitored and possibly controlled by FogLAMP. It may also be used to represent a subset of a device. These values are a collection of *Datapoints* that are the actual values. An asset contains a unique name that is used to reference the data about the asset. An asset is an abstract concept and has no real implementation with the foglamp code, instead a *reading* is used to represent the state of an asset at a point in term. The phase asset is used to represent a time series collection of 0 or more *readings*.

Control Service An optional microservice that is used by the control features of FogLAMP to route control messages from the various sources of control and send them to the *south service* which implements the control path for the *assets* under control.

Core Service The *service* within FogLAMP that is responsible for the oversight of all the other services. It provides configuration management, monitoring, registration and routing services. It is also responsible for the public API into the FogLAMP system and the execution of periodic tasks such as *purge*, statistics and backup.

Datapoint A datapoint is a container for data, each datapoint represents a value that is known about an asset and has a name for that value and the value itself. Values may be one of many types; simpler scalar values, alpha numeric strings, arrays of scalar values, images, arbitrary binary objects or a collection of datapoints.

Filter A combination of a *Filter Plugin* and the configuration that makes that filter perform the processing that is required of it.

Filter Plugin A filter plugin is a *plugin* that implements an operation on one or more *reading* as it passes through the FogLAMP system. This processing may add, remove or augment the data as it passes through FogLAMP. Filters are arrange as linear *pipelines* in either the *south service* as data is ingested into FogLAMP or the *north services* and *tasks* as data is passed upstream to the systems that receive data from FogLAMP.

Microservice A microservice is a small service that implements parts of the FogLAMP functionality. They are also referred to as *services*.

Notification Delivery Plugin A notification delivery plugin is used by the *notification service* to delivery notifications when a *notification rule* triggers. A notification delivery plugin may send notification data to external systems, trigger internal FogLAMP operations or create *reading* data within the FogLAMP *storage service*.

Notification Rule Plugin A notification rule plugin is used by the notification service to determine if a notification should be sent. The rule plugin receives *reading* data from the FogLAMP *storage service*, evaluates a rule against that data and returns a triggered or cleared state to the notification service.

Notification Service An optional *service* within FogLAMP that is responsible for the execution and delivery of notifications when events occurs in the data that is being ingested into FogLAMP.

North An abstract term for any service or system to which FogLAMP sends data that it has ingested. FogLAMP may also receive control messages from the north as well as from other locations.

North Plugin A *plugin* that implements the connection to an upstream system. North plugins are responsible to both implement the communication to the north systems and also the translation from internal data representations to the representation used in the external system.

North Service A *service* responsible for connections upstream from FogLAMP. These are usually systems that will receive data that FogLAMP has ingested and/or processed. There may also be control data flows that operate from the north systems into the FogLAMP system.

North Task A *task* that is run to send data to upstream systems from FogLAMP. It is very similar in operation and concept to a *north service*, but differs from a north service in that it does not always run, it is scheduled using a time based schedule and is designed for situations where connection to the upstream system is not always available or desirable.

Pipeline A linear collection of zero or more *filters* connected between with the *south plugin* that ingests data and the *storage service*, or between the *storage service* and the *north plugin* as data exits FogLAMP to be sent to upstream systems.

Plugin A dynamically loadable code fragment that is used to enhance the capabilities of FogLAMP. These plugins may implement a *south* interface to devices and systems, a *north* interface to systems that receive data from FogLAMP, a *storage plugin* used to buffer *readings*, a *filter plugin* used to process data, a *notification rule* or *notification delivery* plugin. Plugins have well defined interfaces, they can be written by third parties without recourse to modifying the FogLAMP services and are shipped externally to FogLAMP to allow for diverse installations of FogLAMP. Plugins are the major route by which FogLAMP is customized for individual use cases.

Purge The process by which *readings* are removed from the *storage service*.

Reading A reading is the presentation of an *asset* at a point in time. It contains the asset name, two timestamps and the collection of *datapoints* that represent the state of the asset at that point in time. A reading has two timestamps to allow for the time to be recorded when FogLAMP first read the data and also for the device itself to give a time that it sets for when the data was created. Not all devices are capable of reporting timestamps and hence this second timestamp may be the same as the first.

Service FogLAMP is implemented as a set of services, each of which runs constantly and implements a subset of the system functionality. There are a small set of fixed services, such as the *core service* or *storage service*, optional services for enhanced functionality, such as the *notification service* and *control service*. There are also a set of non-fixed services of various types used to interact with downstream or *south* devices and upstream or *north* systems.

South An abstract term for any device or service from which FogLAMP ingests data or over which FogLAMP exerts control.

South Service A *service* responsible for communication with a device or service from which FogLAMP is ingesting data. Each south service connects to a single device and can collect data from that device and optionally send control signals to that device. A south service may represent one or more *assets*.

South Plugin A south plugin is a *plugin* that implements the interface to a device or system from which FogLAMP is collecting data and optionally to which FogLAMP is sending control signals.

Storage Service A *microservice* that implements either permanent or transient storage services used to both buffer *readings* within FogLAMP and also to store FogLAMP's configuration information. The storage services use either one or two *storage plugins* to store the configuration data and the *readings* data.

Storage Plugin A *plugin* that implements the storage requirements of the FogLAMP *storage service*. A plugin may implement the storage of both configuration and *readings* or it may just implement *readings* storage. In this latter case FogLAMP will use two storage plugins, one to store the configuration and the other to store the readings.

Task A task implements functionality that only runs for specific times within FogLAMP. It is used to initiate periodic operations that are not required to be always running. Amongst the tasks that form part of FogLAMP are the *purge task*, *north tasks*, backup and statistics gathering tasks.

INDEX

A

Asset, [603](#)

C

Control Service, [603](#)

Core Service, [603](#)

D

Datapoint, [603](#)

F

Filter, [603](#)

Filter Plugin, [603](#)

M

Microservice, [603](#)

N

North, [604](#)

North Plugin, [604](#)

North Service, [604](#)

North Task, [604](#)

Notification Delivery Plugin, [603](#)

Notification Rule Plugin, [603](#)

Notification Service, [603](#)

P

Pipeline, [604](#)

Plugin, [604](#)

Purge, [604](#)

R

Reading, [604](#)

S

Service, [604](#)

South, [604](#)

South Plugin, [604](#)

South Service, [604](#)

Storage Plugin, [604](#)

Storage Service, [604](#)

T

Task, [605](#)