
Fledge Documentation

Dianomic Systems

Oct 04, 2021

Contents

1	Quick Start Guide	1
1.1	Introduction to Fledge	1
1.2	Installing Fledge	1
1.3	Starting and stopping Fledge	5
1.4	Troubleshooting Fledge	5
1.5	Running the Fledge GUI	5
1.6	Managing Data Sources	7
1.7	Viewing Data	10
1.8	Sending Data to Other Systems	14
1.9	PI Web API OMF Endpoint	15
1.10	EDS OMF Endpoint	18
1.11	OCS OMF Endpoint	20
1.12	PI Connector Relay	22
1.13	Number Format Hints	29
1.14	Integer Format Hints	29
1.15	Type Name Hints	29
1.16	Type Hint	29
1.17	Tag Name Hint	30
1.18	Datapoint Specific Hint	30
1.19	Asset Framework Location Hint	30
1.20	Adding OMF Hints	30
1.21	Backing up and Restoring Fledge	31
1.22	Troubleshooting and Support Information	32
1.23	Package Uninstallation	33
2	Processing Data	35
2.1	Why Use Filters?	35
2.2	What Can Be Done?	35
2.3	Where Can it Be Done?	36
2.4	Some Useful Filters	46
3	Fledge Architecture	49
3.1	Fledge Core	50
3.2	Storage Layer	50
3.3	Southbound Microservices	50
3.4	Northbound Microservices	50
3.5	Filters	51

3.6	Event Engine	51
3.7	REST API	51
3.8	Graphical User Interface	51
4	Fledge Plugins	53
4.1	South Plugins	53
4.2	North Plugins	54
4.3	Filter Plugins	55
4.4	Notification Rule Plugins	56
4.5	Notification Delivery Plugins	57
5	Securing Fledge	59
5.1	Enabling HTTPS Encryption	59
5.2	Requiring User Login	61
5.3	User Management	67
5.4	Certificate Store	71
6	Buffering & Storage	75
6.1	Configuring The Storage Plugin	76
6.2	Installing A PostgreSQL server	78
6.3	SQLite Plugin Configuration	79
7	Tuning Fledge	81
7.1	South Service Advanced Configuration	81
8	Notifications Service	85
8.1	Notifications	85
8.2	Installing the Notification Service	87
8.3	Starting The Notification Service	88
8.4	Configuring The Notification Service	90
8.5	Using The Notification Service	91
9	Set Point Control	103
9.1	Control Functions	103
9.2	Control Paths	104
10	Troubleshooting the PI-Server integration	113
10.1	Log files	113
10.2	How to check the PI Web API is installed and running	114
10.3	Commands to check the PI WEB API	115
10.4	Error messages and causes	117
10.5	Possible solutions to common problems	118
11	Plugin Developer Guide	121
11.1	Plugins	121
11.2	Writing and Using Plugins	124
11.3	South Plugins	128
11.4	South Plugins in C	137
11.5	C++ Support Classes	148
11.6	Hybrid Plugins	153
11.7	North Plugins	154
11.8	Storage Plugins	164
11.9	Filter Plugins	167
11.10	Notification Delivery Plugins	181
11.11	Testing Your Plugin	185

12 REST API Developers Guide	193
12.1 The Fledge REST API	193
12.2 Administration API Reference	194
12.3 User API Reference	212
13 Building Fledge	217
13.1 Building Developers Guide	217
13.2 Building and using Fledge on Raspbian	262
13.3 Building and using Fledge on RedHat/CentOS	263
13.4 Install Fledge on Red Hat/CentOS using the RPM package	263
13.5 Build of Fledge on Red Hat/CentOS	264
14 Version History	265
14.1 Fledge v1	265
15 Downloads	291
15.1 Packages	291
15.2 Download/Clone from GitHub	291
16 Kerberos authentication	293
16.1 Introduction	293
16.2 PI-Server as the North endpoint	293
16.3 North plugin	293
16.4 Fledge server configuration	294
16.5 Kerberos authentication on RedHat/CentOS	296
17 Plugin Documentation	297
17.1 Fledge South Plugins	297
17.2 Fledge North Plugins	354
17.3 Fledge Filter Plugins	389
17.4 Fledge Notification Rule Plugins	421
17.5 Fledge Notification Delivery Plugins	425

1.1 Introduction to Fledge

Fledge is an open sensor-to-cloud data fabric for the Internet of Things (IoT) that connects people and systems to the information they need to operate their business. It provides a scalable, secure, robust infrastructure for collecting data from sensors, processing data at the edge and transporting data to historian and other management systems. Fledge can operate over the unreliable, intermittent and low bandwidth connections often found in IoT applications.

Fledge 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 Fledge systems
- North services: Transmit data to historians and other systems
- Edge data processing applications

Services can easily be developed and incorporated into the Fledge framework. The Fledge Developer Guides describe how to do this.

1.2 Installing Fledge

Fledge 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 Fledge server:

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

You can obtain Fledge in three ways:

- Dianomic Systems hosts a package repository that allows the Fledge packages to be loaded using the system package manager. This is the recommended method for long term use of Fledge as it gives access to all the Fledge plugins and provides a route for easy upgrade of the Fledge packages. This also has the advantages that once the repository is configured you are able to install new plugins directly from the Fledge user interface without the need to resort to the Linux command line.
- Dianomic Systems offers pre-built, certified binaries of Fledge 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 https://fledge-iot.readthedocs.io/en/latest/92_downloads.html.
- As source code from <https://github.com/fledge-iot/>. Instructions for downloading and building Fledge source code can be found in the Fledge Developer's Guide

In general, Fledge installation will require the following packages:

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

1.2.1 Using the package repository to install Fledge

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.fledge-iot.org/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 to 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.fledge-iot.org/latest/buster/armv7l/ /
```

to the end of the file.

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

```
sudo add-apt-repository "deb http://archives.fledge-iot.org/latest/ubuntu1804/x86_
↳64/ / "
```

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

```
sudo add-apt-repository "deb http://archives.fledge-iot.org/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.fledge-iot.org/latest/ubuntu1804/
↳aarch64/ / "
```

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

```
deb http://archives.fledge-iot.org/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 Fledge 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 fledge package, the fledge user interface and the sinusoid south plugin run the command

```
sudo DEBIAN_FRONTEND=noninteractive apt -y install fledge fledge-gui fledge-south-
↳sinusoid
```

RedHat & CentOS

The RedHat and CentOS flavors of Linux use a different package management system, known as *yum*. Fledge also supports a package management system for the yum package manager.

To add the fledge repository to the yum package manager run the command

```
sudo rpm --import http://archives.fledge-iot.org/RPM-GPG-KEY-fledge
```

CentOS users should then create a file called `fledge.repo` in the directory `/etc/yum.repos.d` and add the following content

```
[fledge]
name=fledge Repository
baseurl=http://archives.fledge-iot.org/latest/centos7/x86_64/
enabled=1
gpgkey=http://archives.fledge-iot.org/RPM-GPG-KEY-fledge
gpgcheck=1
```

Users of RedHat systems should do the same, however the files content is slightly different

```
[fledge]
name=fledge Repository
baseurl=http://archives.fledge-iot.org/latest/rhel7/x86_64/
enabled=1
gpgkey=http://archives.fledge-iot.org/RPM-GPG-KEY-fledge
gpgcheck=1
```

There are a few pre-requisites that need to be installed on these platforms, they differ slightly between the two of them.

On CentOS 7 run the commands

```
sudo yum install -y centos-release-scl-rh
sudo yum install -y epel-release
```

On RedHat 7 run the command

```
sudo yum-config-manager --enable 'Red Hat Software Collections RPMs for Red Hat_
↪Enterprise Linux 7 Server from RHUI'
sudo yum install -y https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.
↪noarch.rpm
```

You can now install and upgrade fledge packages using the yum command. For example to install fledge and the fledge GUI you run the command

```
sudo yum install -y fledge fledge-gui
```

1.2.2 Installing Fledge downloaded packages

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

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

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

```
sudo DEBIAN_FRONTEND=noninteractive apt -y install ./fledge-1.8.0-armv7l.deb
sudo apt -y install ./fledge-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 ./fledge-south-sensehat-1.8.0-armv7l.deb
```

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

1.2.3 Checking package installation

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

```
dpkg -l | grep 'fledge'
```

1.2.4 Run with PostgreSQL

To start Fledge 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/fledge/category/Storage/plugin -d '{"value": "postgres"}'
{
  "description": "The main storage plugin to load",
  "type": "string",
  "order": "1",
  "displayName": "Storage Plugin",
  "default": "sqlite",
  "value": "postgres"
}
```

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

1.3 Starting and stopping Fledge

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

The following command options are available:

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

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

```
/usr/local/fledge/bin/fledge start
```

1.4 Troubleshooting Fledge

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

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

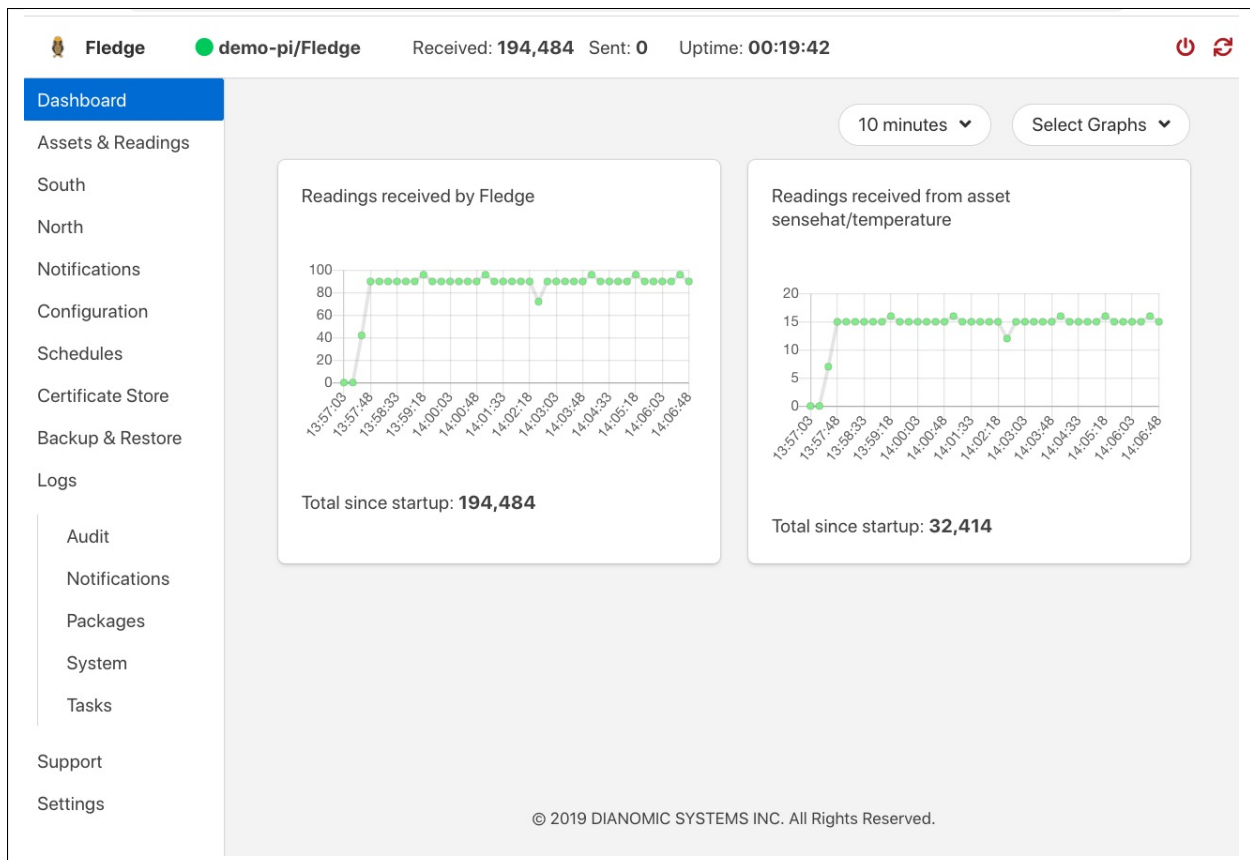
1.5 Running the Fledge GUI

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

You can easily use the Fledge UI to monitor multiple Fledge 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.

1.5.1 Fledge Dashboard



This screen provides an overview of Fledge 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 Fledge system. In addition, you can view graphs of the following system-wide information:

- **Readings:** The total number of data readings collected by Fledge 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

- **Unsparged:** The number of data readings that were purged without being sent to a North service.

1.6 Managing Data Sources

Name	Status	Plugin	Version	Assets	Readings
<u>SenseHAT</u>	enabled	sensehat	1.7.0	sensehat/pressure	33,418
				sensehat/temperature	33,418
				sensehat/humidity	33,418
				sensehat/magnetometer	33,418
				sensehat/gyroscope	33,418
				sensehat/accelerometer	33,418

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 Fledge system. Each data source will display its status, the data assets it is providing, and the number of readings that have been collected.

1.6.1 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 Fledge server. Download the package for the plugin and enter:

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

Once the plugin is installed return to the Fledge GUI and click on “Add+” in the upper right of the South Services screen. Fledge 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 Fledge 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 Fledge 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.

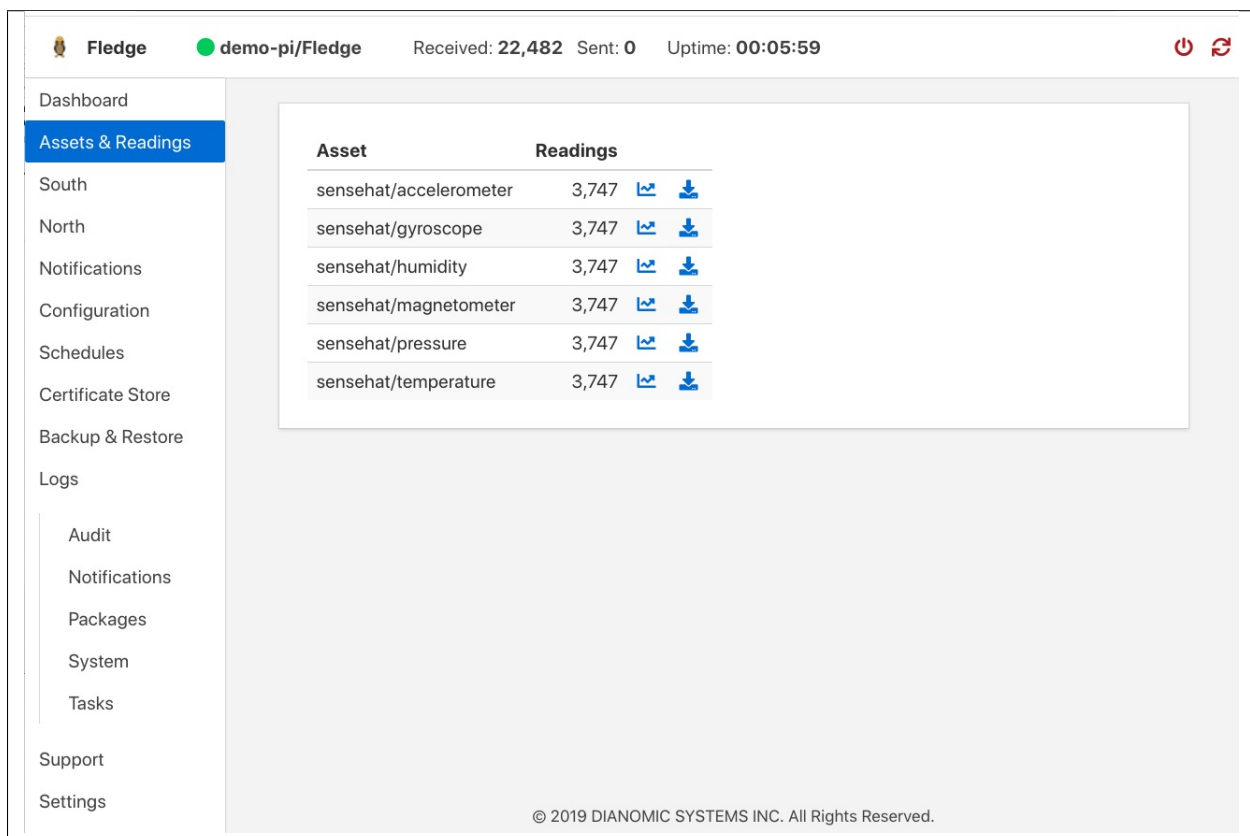
3. If you modify any of the configuration fields, click on the “save” button to save them.
4. The final screen allows you to specify whether the service will be enabled immediately for data collection or await enabling in the future.

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.

1.6.3 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.

1.7 Viewing Data



The screenshot shows the Fledge web interface. The top bar includes the Fledge logo, a green status indicator for 'demo-pi/Fledge', and system statistics: 'Received: 22,482', 'Sent: 0', and 'Uptime: 00:05:59'. A left sidebar contains a menu with items like Dashboard, Assets & Readings (highlighted), South, North, Notifications, Configuration, Schedules, Certificate Store, Backup & Restore, Logs, Audit, Notifications, Packages, System, Tasks, Support, and Settings. The main content area displays a table titled 'Assets & Readings' with the following data:

Asset	Readings	Graph Icon	Download Icon
sensehat/accelerometer	3,747		
sensehat/gyroscope	3,747		
sensehat/humidity	3,747		
sensehat/magnetometer	3,747		
sensehat/pressure	3,747		
sensehat/temperature	3,747		

At the bottom of the page, there is a copyright notice: '© 2019 DIANOMIC SYSTEMS INC. All Rights Reserved.'


You can inspect all the data buffered by the Fledge 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.


1.7.1 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.



1.7.2 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

1.8 Sending Data to Other Systems

The screenshot displays the Fledge web interface. At the top, the status bar shows 'demo-pi/Fledge' with 'Received: 22,938', 'Sent: 0', and 'Uptime: 00:07:14'. The sidebar menu on the left has 'North' selected. The main content area is titled 'North Instances' and contains a table with the following data:

Process	Status	Plugin	Version	Sent
Plant Librarian	enabled			0

At the bottom of the interface, the copyright notice reads: © 2019 DIANOMIC SYSTEMS INC. 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 Fledge system. Each data destination will display its status and the number of readings that have been collected.

1.8.1 Adding Data Destinations

To add a data destination, click on “Create North Instance+” in the upper right of the North Services screen. Fledge 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 Fledge 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.

1.8.2 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.

1.8.3 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.

1.8.4 Using the OMF plugin

OSISoft data historians are one of the most common destinations for Fledge data. Fledge 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.

1.9 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 Fledge 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
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	<input type="checkbox"/>
Asset Framework hierarchies tree	/fledge/data_piwebapi/default
Asset Framework hierarchies rules	<div><div>1</div><div>{ }</div></div>
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

Select PI Web API from the Endpoint options.

- **Basic Information**

- **Endpoint:** Select what you wish to connect to, in this case PI Web API.
- **Send full structure:** Used to control if AF 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 within the PI Server. See [Naming Scheme](#).
- **Server hostname:** The hostname or address of 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. The readings or Fledge'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 Fledge server.

- **Asset Framework**

- **Asset Framework Hierarchies Tree:** The location in the Asset Framework into which the data will be inserted. All data will be inserted at this point in the Asset Framework unless a later rule overrides this.
- **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 equates to no authentication, basic authentication requires a user name and password and Kerberos allows integration with your single sign on environment.
- **PI Web API User Id:** The user name to authenticate with the PI Web API.
- **PI Web API Password:** 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 (Fledge 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 Fledge will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match Fledge 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 Fledge 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.

1.10 EDS OMF Endpoint

To use the OSISoft Edge Data Store first install Edge Data Store on the same machine as your Fledge 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 Fledge 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
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	<input type="checkbox"/>
Asset Framework hierarchies tree	/fledge/data_piwebapi/default
Asset Framework hierarchies rules	<div><div>1</div><div>{ }</div></div>
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

1.10. EDS OMF Endpoint	name_space	19
OCS Namespace		
OCS Tenant ID	ocs_tenant_id	

Select Edge Data Store from the Endpoint options.

- **Basic Information**

- **Endpoint:** Select what you wish to connect to, in this case 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:** The hostname or address of the PI Server. This must be the localhost for EDS.
- **Server port:** The port the Edge Datastore is listening on. Leave as 0 if you are using the default port.
- **Data Source:** Defines which data is sent to the PI Server. The readings or Fledge’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 Fledge 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 (Fledge 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 Fledge will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match Fledge 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 Fledge data types to the data type configured in PI. The defaults 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.

1.11 OCS OMF Endpoint

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

omf

Endpoint

OSIsoft Cloud Services

Server hostname

pi-server

Server port, 0=use the default

0

Producer Token

uid=5ced49c3-3a55-40e7-983f-c6cdcd5c5fd1&crt=20180620084'

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

☐

Asset Framework hierarchies tree

/fledge/data_piwebapi/default

Asset Framework hierarchies rules

1	{}
---	----

PI Web API Authentication Method

anonymous

PI Web API User Id

user_id

PI Web API Password

....

PI Web API Kerberos

piwebapi_kerberos_https.keytab

21

Select OSIsoft Cloud Services from the Endpoint options.

- **Basic Information**

- **Endpoint:** Select what you wish to connect to, in this case 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 the PI Server. The readings or Fledge'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 Fledge server.

- **Authentication**

- **OCS Namespace:** Your namespace within the 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 OCS 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 (Fledge 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 Fledge will time out an HTTP connection attempt.

- **Other (Rarely changed)**


- **Integer Format:** Used to match Fledge 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 Fledge 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.

1.12 PI Connector Relay


The PI Connector Relay was the original mechanism by which OMF data could be ingested into a PI Server, this has since been replaced by the PI Web API OMF endpoint. It is recommended that all new deployments should use the PI Web API endpoint as the Connector Relay has now been discontinued by OSIsoft. To use the Connector Relay, open and sign into the PI Relay Data Connection Manager.



To add a new connector for the Fledge 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 OMF Connector Relay by selecting the new Fledge application, clicking the check box for the OMF Connector Relay and then clicking “Save Configuration”.



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

Now go to the Fledge 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		
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	<input type="checkbox"/>		
Asset Framework hierarchies tree	/fledge/data_piwebapi/default		
Asset Framework hierarchies rules	<table><tr><td>1</td><td>{ }</td></tr></table>	1	{ }
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		

1.12. PI Connector Relay

OCS Namespace	name_space
OCS Tenant ID	ocs_tenant_id

- **Basic Information**

- **Endpoint:** Select what you wish to connect to, in this case the Connector Relay.
- **Server hostname:** The hostname or address of the Connector Relay.
- **Server port:** The port the Connector Relay is listening on. Leave as 0 if you are using the default port.
- **Producer Token:** The Producer Token provided by PI
- **Data Source:** Defines which data is sent to the PI Server. The readings or Fledge'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 Fledge 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 (Fledge 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 Fledge will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match Fledge 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 Fledge 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.

1.12.1 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, in your environment you have objects that are liable to change, i.e. the types of attributes change or the number of attributes change between readings, then you may wish to take a different naming approach to if they do not.

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 for the AF Objects. OMF lets a type be defined, but once defined it can not be changed. A new type must be created rather than changing the existing type. This means a new asset framework object 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 Fledge. 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 data points within each reading within Fledge. 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 the objects in the AF framework and Attributes in the PI server. However if the structure of an asset changes a new AF Object will be created which will have the suffix -type*x* appended to it.

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

Use Attribute Hash 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 Fledge, both type suffices and attribute hashes will be applied to the naming.

1.12.2 Asset Framework Hierarchy Rules

The asset framework rules allow the location of specific assets within the PI Asset Framework to be controlled. There are two basic type 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"
      }
    }
  }
}
```

The name type rules are simply a set of asset name and AF 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 data point having a given value.
- **notequal:** This test looks for a name data point 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 data point called *temperature* it will be stored in the AF hierarchy *temperatures*, if the asset had a data point 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 data point 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 data point 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 data points within the reading. For example if the reading has a data point 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 data point. 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 data point is missing. Therefore in our example above if we used the location `/BuildingA/${room:unknown}` a reading without a `room` data point would be placed in `/BuildingA/unknown`. If no default is given and the data point is missing then the level in the hierarchy is ignored. E.g. if we use our original location `/BuildingA/${room}` and we have the reading

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

this reading would be stored in `/BuildingA`.

1.12.3 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 data point within the asset, *OMFHint*. The hints themselves are encoded as JSON within a string.

1.13 Number Format Hints

A number format hint tells the plugin what number format to insert data into the PI Server as. 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.

1.14 Integer Format Hints

An integer format hint tells the plugin what integer format to insert data into the PI Server as. 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.

1.15 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" }
```

1.16 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" }
```

1.17 Tag Name Hint

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

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

1.18 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 it's value, this 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.

1.19 Asset Framework Location Hint

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

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

1.20 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 .


1.21 Backing up and Restoring Fledge

The screenshot displays the Fledge web interface. At the top, the header shows 'Fledge' with a green status indicator, 'demo-pi/Fledge', and system statistics: 'Received: 23,460 Sent: 0 Uptime: 00:08:44'. On the right of the header are power and refresh icons. The left sidebar contains a menu with items: Dashboard, Assets & Readings, South, North, Notifications, Configuration, Schedules, Certificate Store, Backup & Restore (highlighted in blue), Logs, Audit, Notifications, Packages, System, Tasks, Support, and Settings. The main content area is titled 'Backups' and includes a 'Create New' button with a document icon. Below this is a table with two columns: 'Date & Time' and 'Status'. A single row is visible with the date '2020-03-31 10:10:36' and status 'COMPLETED'. To the right of the status are three links: 'delete', 'restore', and 'download'. At the bottom of the main area, a copyright notice reads '© 2019 DIANOMIC SYSTEMS INC. All Rights Reserved.'

Date & Time	Status	
2020-03-31 10:10:36	COMPLETED	delete restore download

You can make a complete backup of all Fledge 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.

1.22 Troubleshooting and Support Information



Fledge 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 Fledge system.
- **Notifications:** If you are using the Fledge 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 Fledge 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 Fledge 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.

1.23 Package Uninstallation

1.23.1 Debian Platform

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

```
sudo apt -y purge fledge
```

1.23.2 RPM Platform

```
sudo yum -y remove fledge
```

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

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

This is due to the fact that the data directory (`/usr/local/fledge/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 fledge completely from your system, then do `rm -rf /usr/local/fledge` directory.

We have already seen that Fledge 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 Fledge to edit, augment or remove data as it traverses the Fledge system. In the same way Fledge makes extensive use of plugin components to add new sources of data and new destinations for that data, Fledge also uses plugins to add processing filters to the Fledge system.

2.1 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 Fledge by merely configuring filters from the existing pool of available filters into a suitable pipeline without the need to write any new code.

2.2 What Can Be Done?

Data processing is done via plugins that are known as *filters* in Fledge, 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.

2.3 Where Can it Be Done?

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

- In the south service as data arrives in Fledge 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 Fledge system.

More than one filter can be added to a single south or north within a Fledge 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 Fledge or out of Fledge.

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.

2.3.1 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 Fledge 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". At the top right is a close button (X). Below the title bar, there are several sections:

- Asset name:** A text input field containing the value "sinusoid".
- Enabled:** A checkbox that is currently checked.
- Applications:** A section with a "+" icon, indicating where to add new applications.
- Buttons:** "Cancel" and "Save" buttons are located to the right of the "Applications" section. "Export Readings" and "Delete Service" buttons are located at the bottom right.
- Service Info:** A section label at the bottom left.
- Show Advanced Config:** A link located to the right of the "Enabled" checkbox.

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.

The screenshot shows a configuration window titled "Sine South Service" with a close button in the top right corner. A progress bar at the top indicates two steps: "1" (Plugin Name) and "2" (Review Configuration). The "Plugin Name" step is active. Below the progress bar, there is a section for selecting a plugin. On the left, under the label "Plugin", there is a dropdown menu with four options: "asset", "change", "delta", and "expression". The "expression" option is highlighted with a blue background. To the right of the dropdown, the text "Apply an expression to the data stream" is displayed. Below the dropdown, there is a link that says "Install from available plugins". Under the label "Name", there is a text input field containing the placeholder text "name". At the bottom of the window, there are two buttons: "Back" and "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)

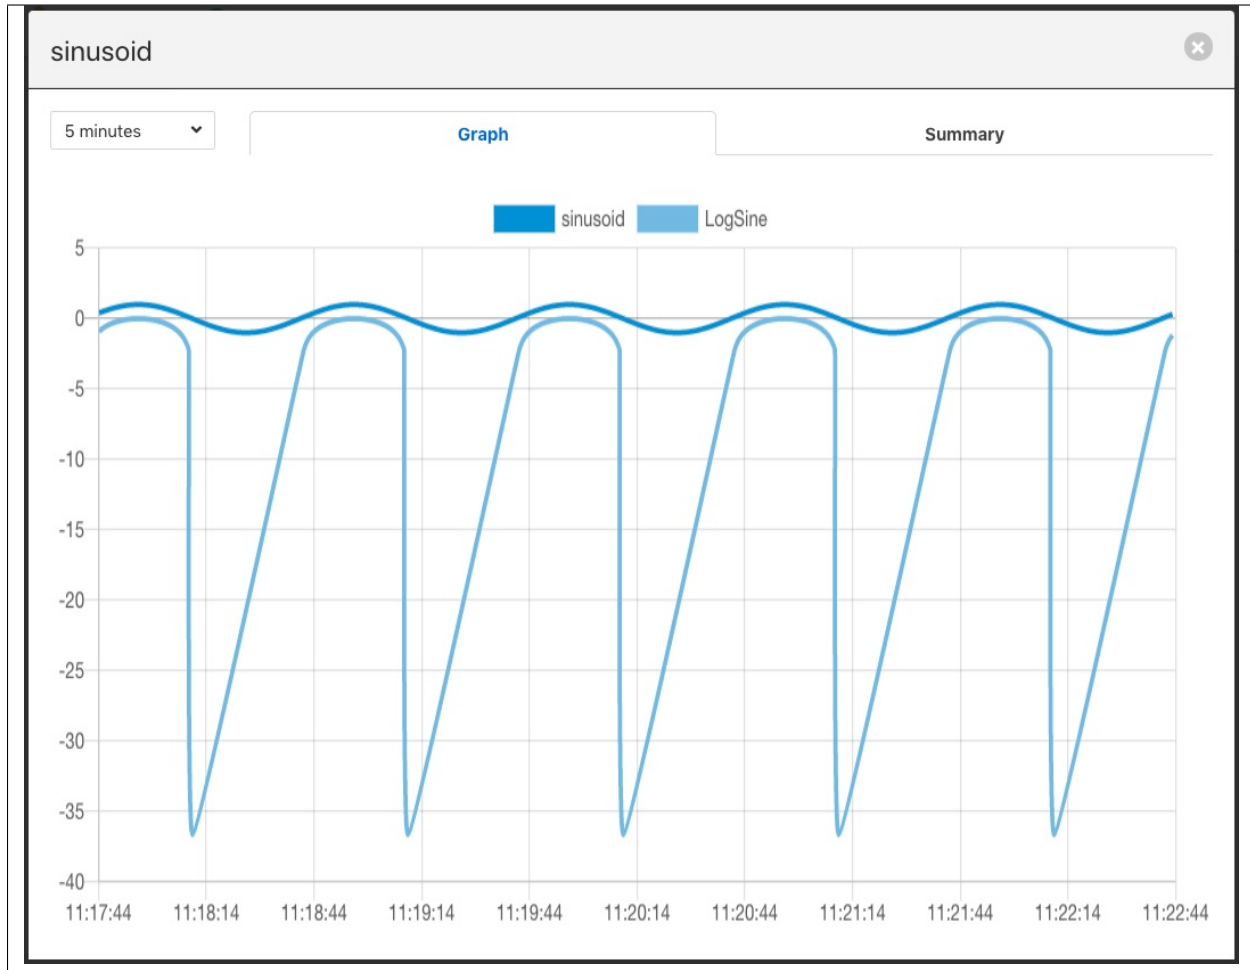
Enabled ☒

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

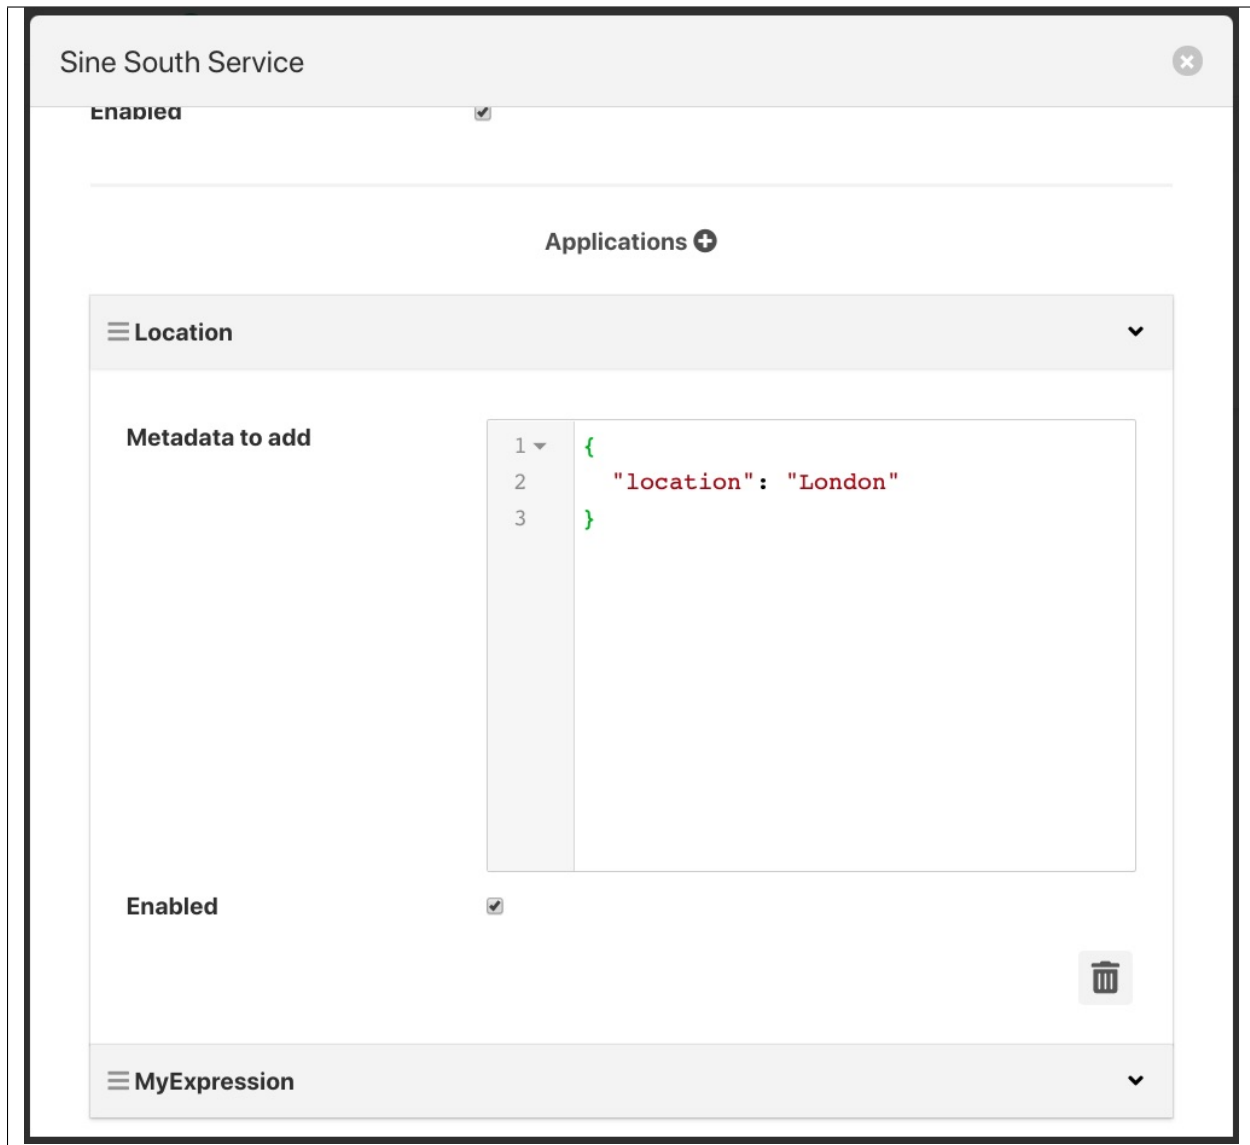
Export Readings

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.

2.3.2 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 Fledge 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. Below the title bar, there is a progress indicator with two steps: "1" (highlighted in a green circle) and "2". Step 1 is labeled "Plugin Name" and Step 2 is labeled "Review Configuration".

Under Step 1, there is a "Plugin" dropdown menu with the following options: "fft", "FlirValidity", "metadata" (highlighted), and "python27". To the right of the dropdown, the text "Metadata filter plugin" is displayed. Below the dropdown, there is a link that says "Install from available plugins".

Below the plugin selection, there is a "Name" label and a text input field containing the text "Floor".

At the bottom of the form, there 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*.

PI Server

1 Plugin Name 2 Review Configuration

Metadata to add

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

Enabled ☒

Previous 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.

2.4 Some Useful Filters

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

2.4.1 Scale

The filter *fledge-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.

2.4.2 Metadata

The filter *fledge-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.

2.4.3 Delta

The filter *fledge-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.

2.4.4 Rate

The filter *fledge-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.

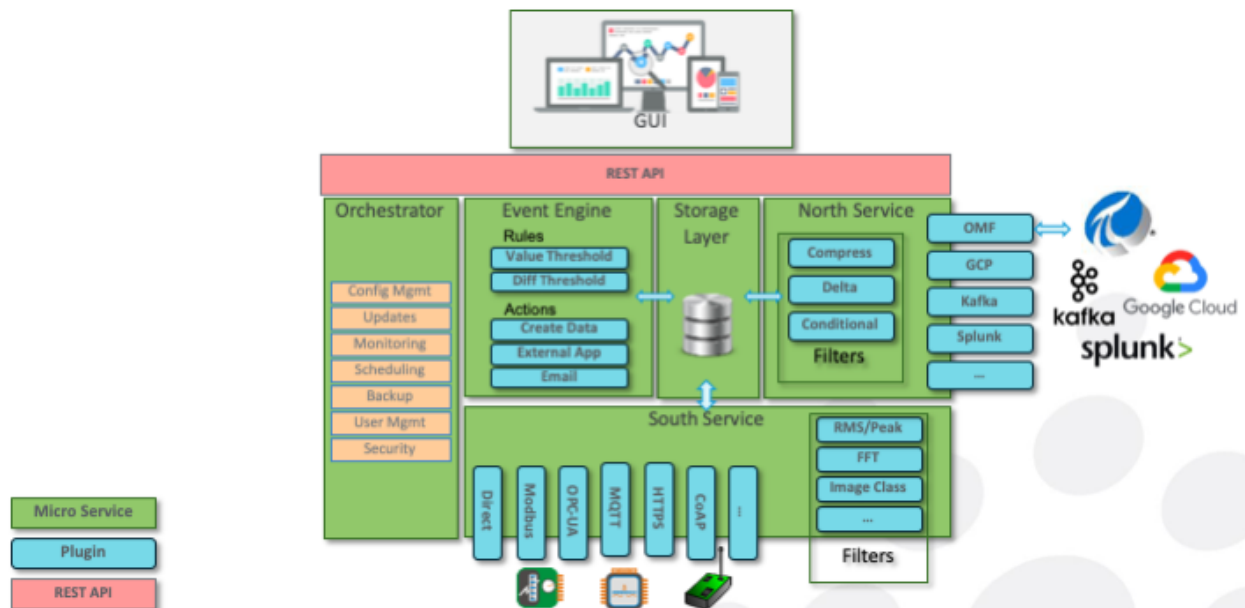
A rate of one reading per minute for example would average all the values for 1 minute and then send that average as the reading at the end of that minute. A condition can be added, when that condition is triggered all data is forwarded at full rate of the incoming data until a further condition is triggered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

CHAPTER 3

Fledge Architecture

The following diagram shows the architecture of Fledge:

- Components in blue are **plugins**. Plugins are light-weight modules that enable Fledge 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 in green are **microservices**. They can co-exist in the same operating environment or they can be distributed across multiple environments.



3.1 Fledge Core

The Core microservice coordinates all of the Fledge 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 Fledge components. Enable software updates across all Fledge components.

Monitoring: monitor all Fledge 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: Fledge 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 Fledge administrators.

Asset Browsing: enable querying of stored asset data.

3.2 Storage Layer

The Storage microservice provides two principal functions: a) maintenance of Fledge 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, in-memory temporary storage may be the best option.

3.3 Southbound Microservices

Southbound microservices offer bi-directional communication of data and metadata between Edge devices, such as sensors, actuators or PLCs and Fledge. Smaller systems may have this service installed onboard Edge devices. Southbound components are typically deployed as always-running services, which continuously wait for new data. Alternatively, they can be deployed as single-shot tasks, which periodically spin up, collect data and spin down.

3.4 Northbound Microservices

Northbound microservices offer bi-directional communication of data and metadata between the Fledge platform and larger systems located locally or in the cloud. Larger systems may be private and public Cloud data services, proprietary solutions or Fledge instances with larger footprints. Northbound 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.

3.5 Filters

Filters are plugins which modify streams of data that flow through Fledge. 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.

3.6 Event Engine

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, or create arbitrary data, or email/slack/hangout/communicate a message.

3.7 REST API

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

3.8 Graphical User Interface

A GUI enables administration of Fledge. All GUI capability is through the REST API, so Fledge 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 Fledge

Assets & Readings: analytics of data in Fledge

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 Fledge

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

Fledge Plugins

The following set of plugins are available for Fledge. 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.

4.1 South Plugins

South plugins add new ways to get data into Fledge, 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: Fledge South Plugins

Name	Description
am2315	Fledge 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
benchmark	A Fledge benchmark plugin to measure the ingestion rates on particular hardware
cc2650	A Fledge south plugin for the Texas Instruments SensorTag CC2650
coap	A south plugin for Fledge that pulls data from a COAP sensor
coral-enviro	A south plugin for the Google Coral Environmental Sensor Board
csv	A Fledge south plugin in C++ for reading CSV files
csv-async	A Fledge 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 Fledge south plugin in C++ that interfaces to a DHT-11 temperature and humidity sensor
dht11	A Fledge south plugin that interfaces a DHT-11 temperature sensor
dnp3	A south plugin for Fledge that implements the DNP3 protocol
envirophat	A Fledge south service for the Raspberry Pi EnviroPhat sensors
expression	A Fledge south plugin that uses a user define expression to generate data
FlirAX8	A Fledge hybrid south plugin that uses fledge-south-modbus-c to get temperature data from a Flir Thermal camera

Continued on next page

Table 1 – continued from previous page

Name	Description
game	The south plugin used for the Fledge lab session game involving remote controlled cars
http	A Python south plugin for Fledge used to connect one Fledge instance to another
iec104	A south plugin to gather data using the IEC 104 protocol.
iec61850	A south plugin for collecting data via the IEC 61850 protocol
ina219	A Fledge 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 Fledge.
lathesim	A simulation plugin used as a demonstration to show how data can be collected within Fledge. This plugin simulates various properties of a lathe.
modbus-c	A Fledge south plugin that implements modbus-tcp and modbus-rtu
modbustcp	A Fledge south plugin that implements modbus-tcp in Python
mqtt	Fledge South MQTT Subscriber Plugin
mqtt-sparkplug	A Fledge south plugin that implements the Sparkplug API over MQTT
opcua	A Fledge south service that pulls data from an OPC-UA server
openweathermap	A Fledge south plugin to pull weather data from OpenWeatherMap
person-detection	Fledge south service plugin that detects person in the live video stream
playback	A Fledge south plugin to replay data stored in a CSV file
pt100	A Fledge south plugin for the PT100 temperature sensor
random	A south plugin for Fledge that generates random numbers
randomwalk	A Fledge south plugin that returns data that with randomly generated steps
roxtec	A Fledge south plugin for the Roxtec cable gland project
s2opcua	An OPCUA south plugin based on the Safe & Secure OPCUA library. This plugin offers similar functionality to the fledge-south-opcua plugin but also offers encryption and authentication.
sensehat	A Fledge south plugin for the Raspberry Pi Sensehat sensors
sensorphone	A Fledge south plugin the task to the iPhone SensorPhone app
sinusoid	A Fledge south plugin that produces a simulated sine wave
systeminfo	A Fledge south plugin that gathers information about the system it is running on.
usb4704	A Fledge south plugin the Advantech USB-4704 data acquisition module
wind-turbine	A Fledge south plugin for a number of sensor connected to a wind turbine demo

4.2 North Plugins

North plugins add new destinations to which data may be sent by Fledge. 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: Fledge 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
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 Fledge instances using HTTP
http-c	A Fledge north plugin that sends data between Fledge instances using HTTP/HTTPS
iec104	A Fledge north plugin for sending data using the IEC-104 protocol.
kafka	A Fledge plugin for sending data north to Apache Kafka
kafka-python	A Python implementation of a north plugin that can send data to Apache Kafka
opcua	A north plugin for Fledge that makes it act as an OPC-UA server for the data it reads from sensors
thingspeak	A Fledge north plugin to send data to Matlab's ThingSpeak cloud
timestream	A timestream north plugin
OMF	Send data to OSIsoft PI Server, Edge Data Store or OSIsoft Cloud Services

4.3 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 Fledge as it is read from devices or written to destination systems.

Table 3: Fledge Filter Plugins

Name	Description
asset	A Fledge processing filter that is used to block or allow certain assets to pass onwards in the data stream
change	A Fledge processing filter plugin that only forwards data that changes by more than a configurable amount
delta	A Fledge 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
expression	A Fledge processing filter plugin that applies a user define formula to the data as it passes through the filter
fft	A Fledge processing filter plugin that calculates a Fast Fourier Transform across sensor data
Flir-Validity	A Fledge processing filter used for processing temperature data from a Flir thermal camera
log	A Fledge filter that converts the readings data to a logarithmic scale. This is the example filter used in the plugin developers guide.
metadata	A Fledge processing filter plugin that adds metadata to the readings in the data stream
omfhint	A filter plugin that allows data to be added to assets that will provide extra information to the OMF north plugin.
python27	A Fledge processing filter that allows Python 2 code to be run on each sensor value.
python35	A Fledge processing filter that allows Python 3 code to be run on each sensor value.
rate	A Fledge processing filter plugin that sends reduced rate data until an expression triggers sending full rate data
rename	A Fledge 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 Fledge processing filter plugin that calculates RMS value for sensor data
scale	A Fledge processing filter plugin that applies an offset and scale factor to the data
scale-set	A Fledge processing filter plugin that applies a set of scale factors to the data
threshold	A Fledge processing filter that only forwards data when a threshold is crossed

4.4 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: Fledge Notification Rule Plugins

Name	Description
average	A Fledge 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.
outofbound	A Fledge notification rule plugin that triggers when sensors values exceed limits set in the configuration of the plugin.
simple-expression	A Fledge notification rule plugin that evaluates an expression based sensor data

4.5 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 deliveries.

Table 5: Fledge Notification Delivery Plugins

Name	Description
alexa-notifyme	A Fledge notification delivery plugin that sends notifications to the Amazon Alexa platform
asset	A Fledge notification delivery plugin that creates an asset in Fledge when a notification occurs
blynk	A Fledge notification delivery plugin that sends notifications to the Blynk service
email	A Fledge notification delivery plugin that sends notifications via email
google-hangouts	A Fledge notification delivery plugin that sends alerts on the Google hangout platform
ifttt	A Fledge notification delivery plugin that triggers an action of IFTTT
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.
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 Fledge notification delivery plugin that runs an arbitrary Python 3 script
setpoint	A fledge notification plugin that invokes a set point operation on a south service.
slack	A Fledge notification delivery plugin that sends notifications via the slack instant messaging platform
telegram	A Fledge notification delivery plugin that sends notifications via the telegram service

The default installation of a Fledge service comes with security features turned off, there are several things that can be done to add security to Fledge. 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 Fledge system. Authentication can be via username and password or by means of an authentication certificate.

5.1 Enabling HTTPS Encryption

Fledge 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,

Fledge Admin and User REST API

Enable HTTP ☒

HTTP Port 8081

HTTPS Port 1995

Certificate Name fledge

Authentication optional

Authentication method any

Auth Certificate ca

Allow Ping ☒

Password Expiry Days 0

Auth Providers

```

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 8081

HTTPS Port 1995

Certificate Name fledge

When this is unticked two options become active on the page, *HTTPS Port* and *Certificate Name*. The HTTPS Port is

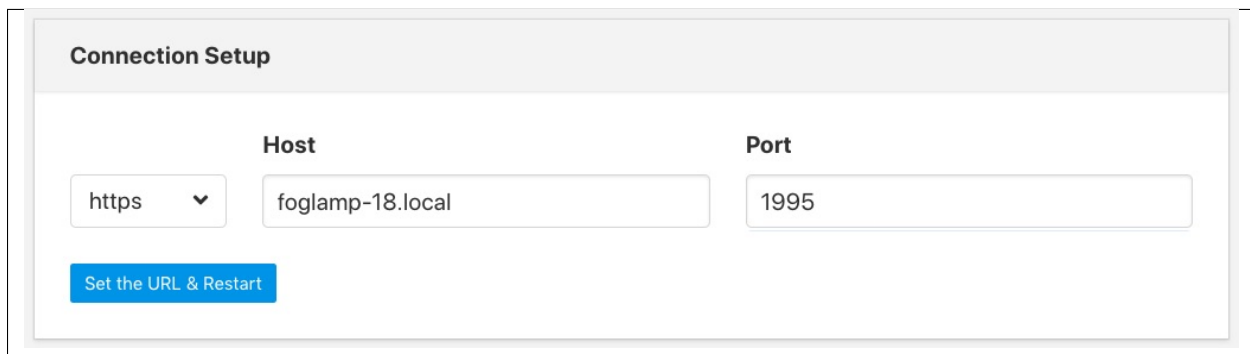
the port that Fledge will listen on for HTTPS requests, the default for this is port 1995.

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 *fledge* that is created as part of the installation process. This certificate is unique per fledge installation but is not signed by a certificate authority. If you require the extra security of using a signed certificate you may use the Fledge *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 Fledge 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_FLEDGE_IP:1995`

Then follow browser instruction in order to allow the connection and close the tab. In the Fledge GUI you should see the green icon (Fledge is running).



The screenshot shows a 'Connection Setup' dialog box. It has a title bar 'Connection Setup'. Below the title bar, there are two columns: 'Host' and 'Port'. Under 'Host', there is a dropdown menu showing 'https' and a text input field containing 'foglamp-18.local'. Under 'Port', there is a text input field containing '1995'. At the bottom left of the dialog, there is a blue button labeled 'Set the URL & Restart'.

5.2 Requiring User Login

In order to set the REST API and GUI to force users to login before accessing Fledge select the *Configuration* option from the left-hand menu and then select *Admin API* from the configuration tree that appears.

Dashboard

Assets & Readings

South

North

Notifications

Configuration

Schedules

Certificate Store

Backup & Restore

Logs

Audit

Notifications

Packages

System

Tasks

Support

Settings

General

- Installation
- Admin API
- Fledge Service

Fledge Admin and User REST API

Enable HTTP ☒

HTTP Port 8081

HTTPS Port 1995

Certificate Name fledge

Authentication optional

Authentication method any

Auth Certificate ca

Allow Ping ☒

Password Expiry Days 0

Auth Providers

```

1 {
2   "providers": [
3     "username",
4     "ldap"
5   ]
6 }

```

Two particular items are of interest in this configuration category that is then displayed; *Authentication* and *Authentication method*

Certificate Name fledge

Authentication

Authentication method

✓ mandatory

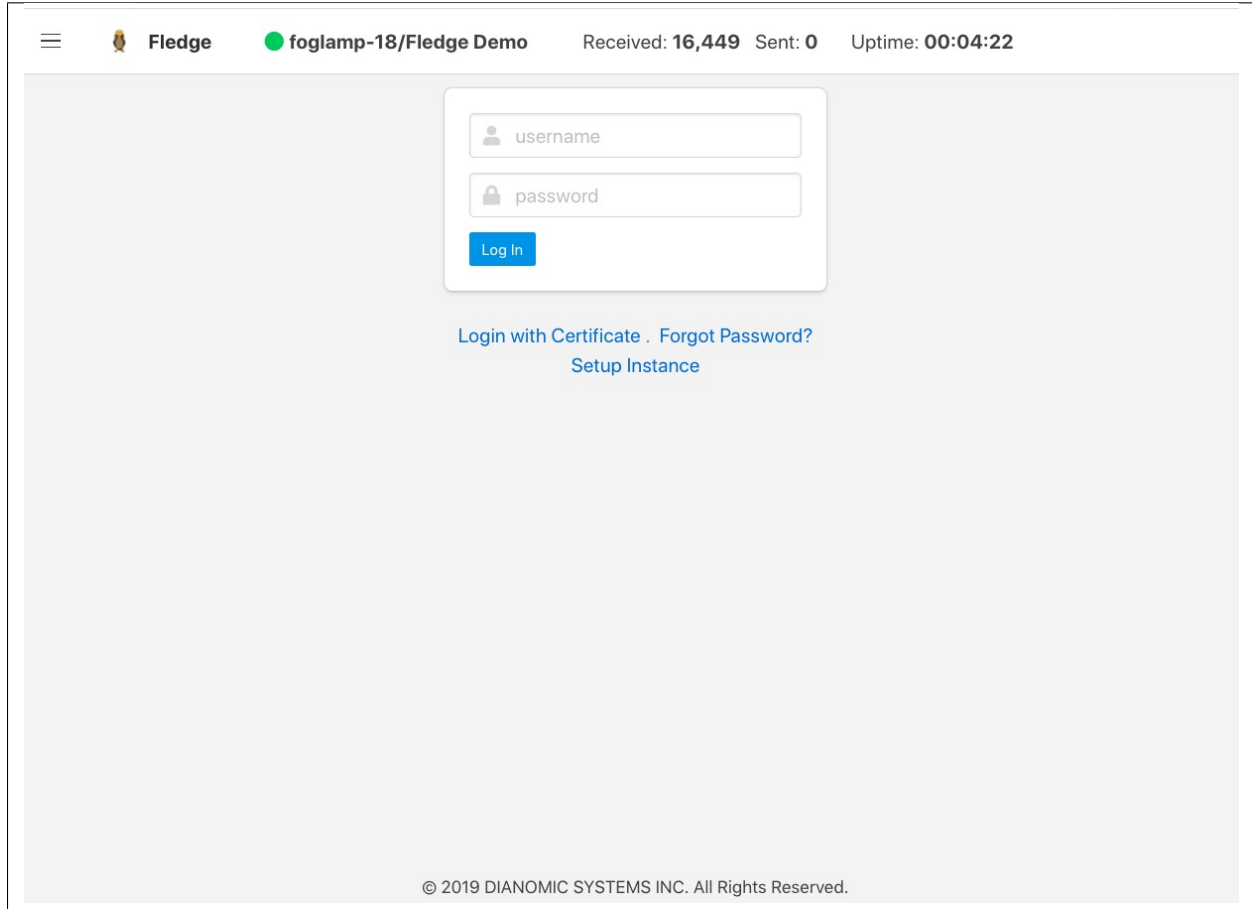
optional

password

Select the *Authentication* field to be mandatory and the *Authentication method* to be password. Click on *Save* at the bottom of the dialog.

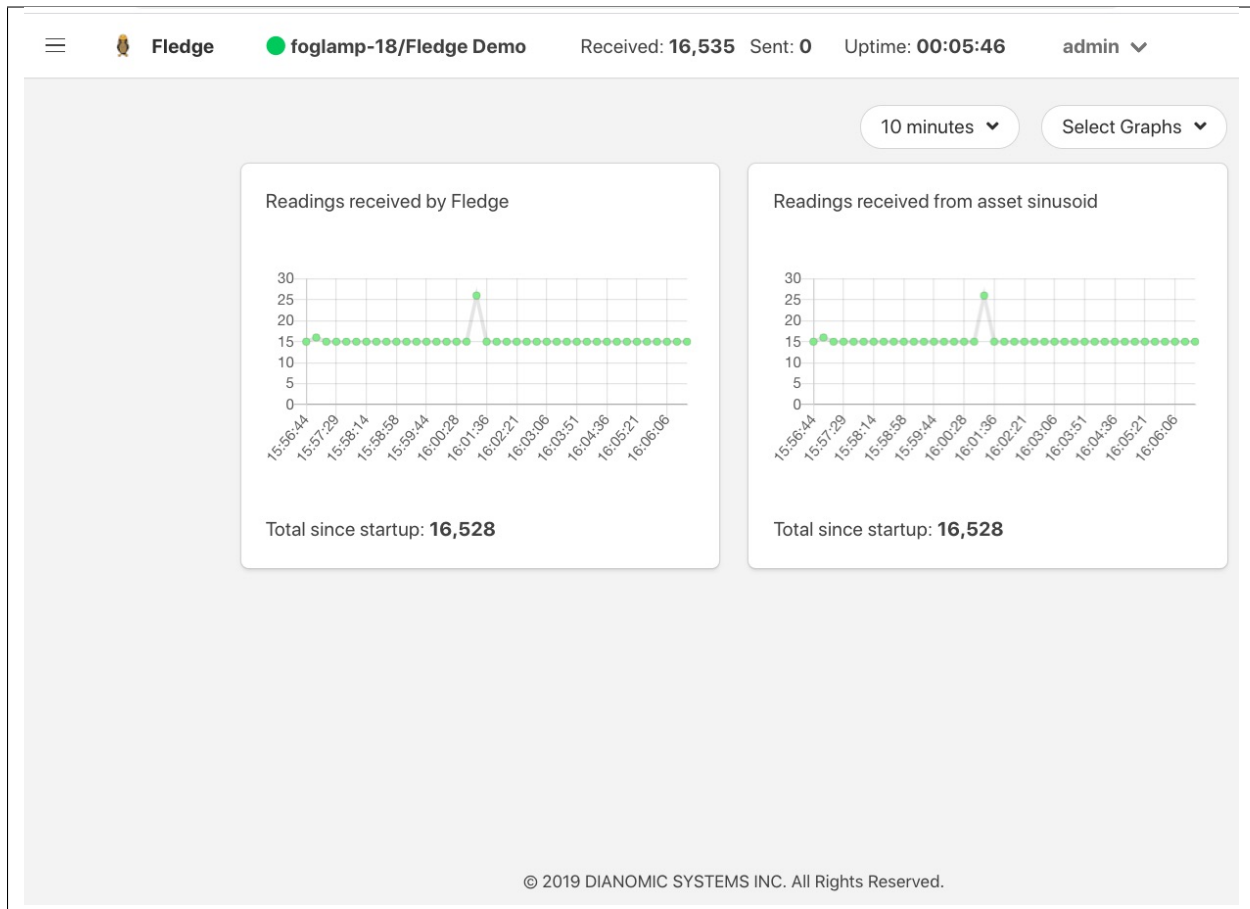
In order for the changes to take effect Fledge must be restarted, this can be done in the GUI by selecting the restart item in the top status bar of Fledge. Confirm the restart of Fledge and wait for it to be restarted.

Once restarted refresh your browser page. You should be presented with a login request.



The screenshot shows the Fledge web interface. At the top, there is a navigation bar with a hamburger menu icon, the Fledge logo, a green status indicator, the text "foglamp-18/Fledge Demo", and system statistics: "Received: 16,449 Sent: 0 Uptime: 00:04:22". The main content area is a light gray background. In the center, there is a white login box containing two input fields: "username" with a user icon and "password" with a lock icon. Below these fields is a blue "Log In" button. Under the login box, there are three links: "Login with Certificate .", "Forgot Password?", and "Setup Instance". At the bottom of the page, there is a copyright notice: "© 2019 DIANOMIC SYSTEMS INC. All Rights Reserved."


The default username is “admin” with a password of “fledge”. Use these to login to Fledge, 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*.

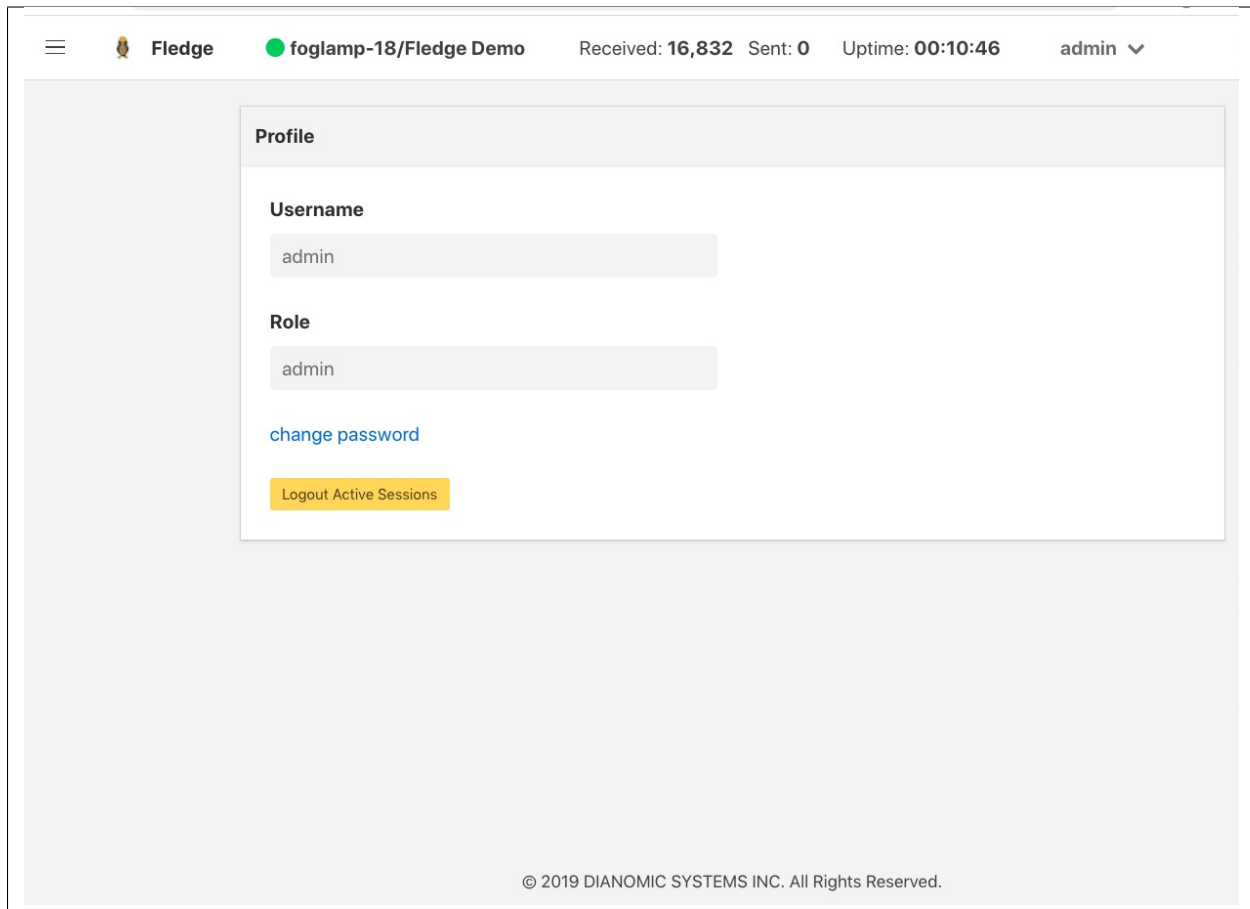
5.2.1 Changing Your Password

The top status bar of the Fledge 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.



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.

5.2.2 Password Rotation Mechanism

Fledge 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 Fledge web interface. The top bar includes the Fledge logo, a status indicator for 'foglamp-18/Fledge Demo', and system metrics: 'Received: 117,981', 'Sent: 0', 'Uptime: 00:12:02', and a 'mark' dropdown. The left sidebar contains navigation links: 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 'Fledge Admin and User REST API' and contains the following configuration options:

- Enable HTTP:** ☒
- HTTP Port:** 8081
- HTTPS Port:** 1995
- Certificate Name:** fledge
- Authentication:** mandatory
- Authentication method:** any
- Auth Certificate:** ca
- Allow Ping:** ☒
- Password Expiry Days:** 30
- Auth Providers:** A code editor showing the following JSON configuration:


```
1 {
2   "providers": [
3     "username",
4     "ldap"
5   ]
6 }
```

Whenever a user logs into Fledge 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.

5.3 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*.

ID	Username	Role	
1	admin	admin	super admin active
2	user	user	change role reset password delete

The user management pages allows

- Adding new users.
- Deleting users.
- Resetting user passwords.
- Changing the role of a user.

Fledge currently supports two roles for users,

- **admin:** a user with admin role is able to fully configure Fledge and also manage Fledge users
- **user:** a user with this role is able to configure Fledge but can not manage users

5.3.1 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 the text 'Create User' and a close button (an 'x' in a circle) on the right. The main area contains a 'Role' label above a dropdown menu showing 'user'. Below the dropdown is a text input field containing 'mark'. Underneath that is a password input field with a masked password '.....'. The bottom of the dialog features a blue 'Save' button. The entire dialog is enclosed in a light gray border with a dark gray shadow.

Create User

Role

user

mark

.....

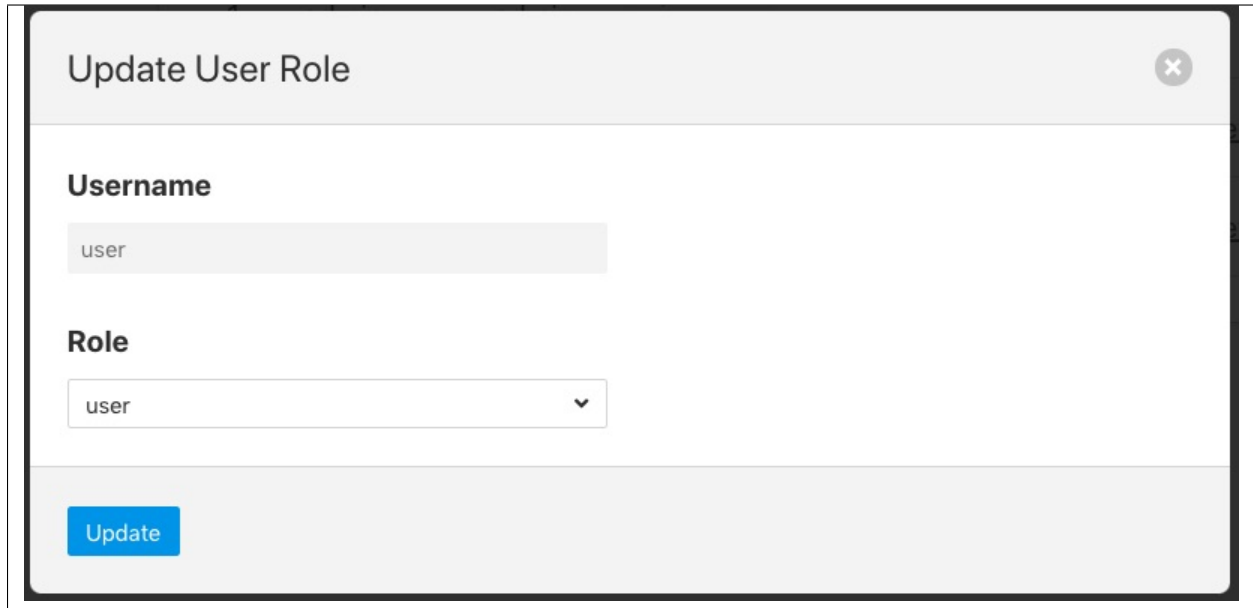
.....

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.

5.3.2 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.

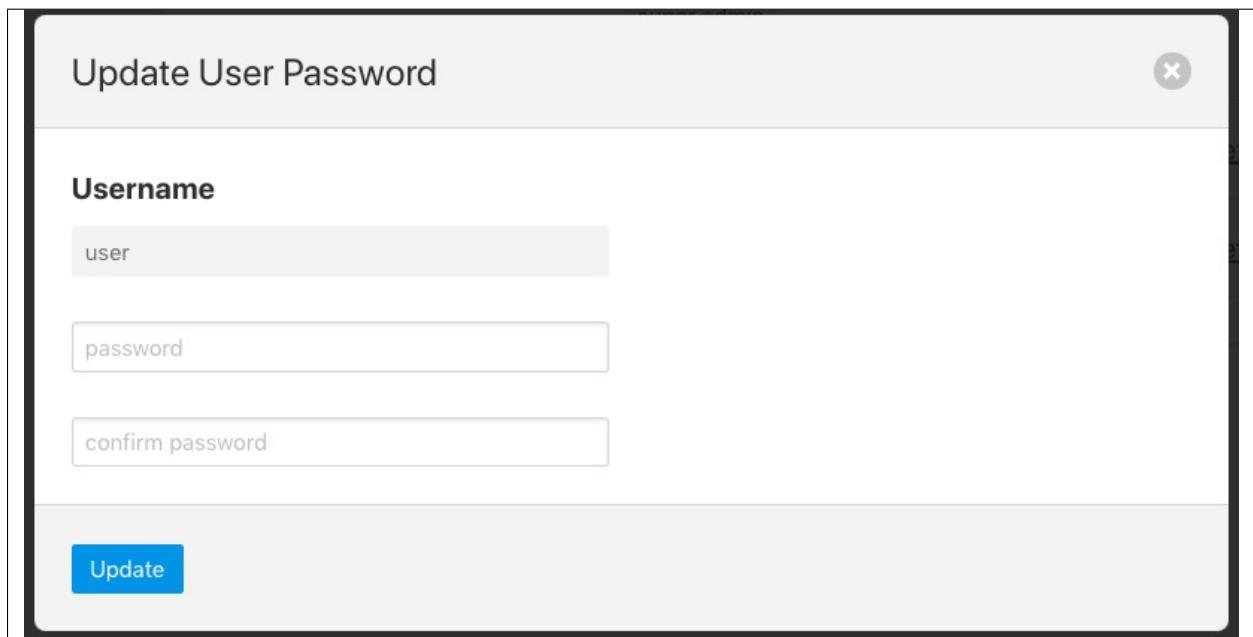


The dialog box titled "Update User Role" has a close button (X) in the top right corner. It contains two input fields: "Username" with the value "user" and "Role" with a dropdown menu showing "user". At the bottom left is a blue "Update" button.

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.

5.3.3 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 dialog box titled "Update User Password" has a close button (X) in the top right corner. It contains three input fields: "Username" with the value "user", "password", and "confirm password". At the bottom left is a blue "Update" button.

Enter the new password and confirm that password by entering it a second time and click on *Update*.

5.3.4 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 Fledge.

5.4 Certificate Store

The Fledge *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.

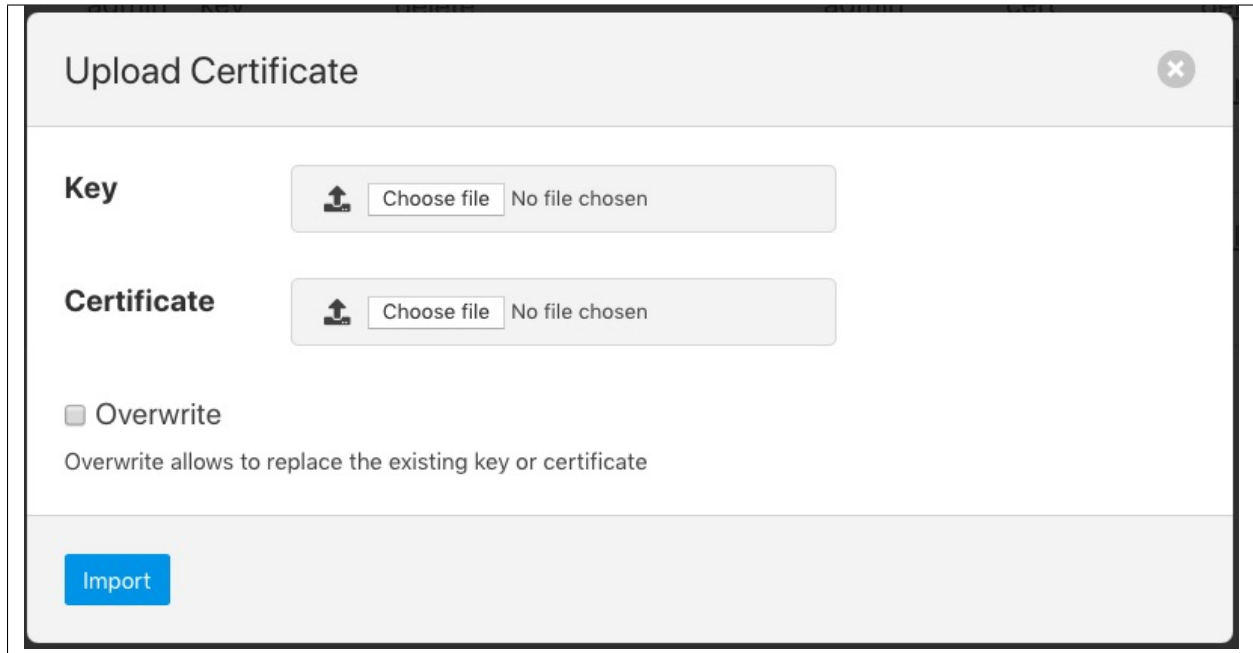
The screenshot shows the Fledge web interface. The top header includes the Fledge logo, a status indicator for 'foglamp-18/Fledge Demo', and system metrics: 'Received: 44,751', 'Sent: 0', 'Uptime: 2 days, 18:37:40', and a user dropdown for 'admin'. The left sidebar contains a list of navigation items, with 'Certificate Store' currently selected and highlighted in blue. The main content area is titled 'Certificate Store' and features an 'Import' button with an upload icon. Below the title, there are two tables. The first table, labeled 'Key', lists four entries: 'admin' (key), 'ca' (key), 'fledge' (key), and 'user' (key). Each entry has a 'delete' link. The second table, labeled 'Certificate', lists four entries: 'admin' (cert), 'ca' (cert), 'fledge' (cert), and 'user' (cert). Each entry also has a 'delete' link. At the bottom of the page, a copyright notice reads: '© 2019 DIANOMIC SYSTEMS INC. All Rights Reserved.'

Key	Extension	
admin	key	delete
ca	key	delete
fledge	key	
user	key	delete

Certificate	Extension	
admin	cert	delete
ca	cert	delete
fledge	cert	
user	cert	delete

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.



The screenshot shows a dialog box titled "Upload Certificate" with a close button (X) in the top right corner. The dialog is divided into two main sections: "Key" and "Certificate". Each section has a file upload icon (a square with an upward arrow) and a text input field. The "Key" section's input field contains the text "Choose file" and "No file chosen". The "Certificate" section's input field also contains the text "Choose file" and "No file chosen". Below these sections is a checkbox labeled "Overwrite" which is currently unchecked. Underneath the checkbox is a line of text: "Overwrite allows to replace the existing key or certificate". At the bottom left of the dialog is a blue button labeled "Import".

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.

Buffering & Storage

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

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

The storage service is configurable, like other services within Fledge and uses plugins to extend the functionality of the storage system. These storage plugins provide the underlying mechanism by which data is stored within Fledge. Fledge 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 Fledge 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 Fledge 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


6.1 Configuring The Storage Plugin

Once installed the storage plugin can be reconfigured in much the same way as any Fledge 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 Fledge 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 Fledge 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 Fledge to start as stop the database server when Fledge is started and stopped. If it is left off then Fledge 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 Fledge installation. These will be documented in the packages that provide the storage plugin.

Storage plugin configurations are not dynamic and Fledge *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 Fledge.

6.1.1 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.

6.2 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.

6.2.1 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 `<fledge_user>`, type:

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

The `-d` argument is important because the user will need to create the Fledge database.

A more generic command is:

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

6.2.2 CentOS/Red Hat Install

On CentOS and Red Hat systems, and other RPM based distributions the command is

```
sudo yum install -y https://download.postgresql.org/pub/repos/yum/reporpms/EL-7-x86_
↪64/pgdg-redhat-repo-latest.noarch.rpm
sudo yum install -y postgresql96-server
sudo yum install -y postgresql96-devel
sudo yum install -y rh-postgresql96
sudo yum install -y rh-postgresql96-postgresql-devel
sudo /usr/pgsql-9.6/bin/postgresql96-setup initdb
sudo systemctl enable postgresql-9.6
sudo systemctl start postgresql-9.6
```

At this point, Postgres has been configured to start at boot and it should be up and running. You can always check the status of the database server with `systemctl status postgresql-9.6`:

```
sudo systemctl status postgresql-9.6
```

Next, you must create a PostgreSQL user that matches your Linux user.


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

Finally, add `/usr/pgsql-9.6/bin` to your `PATH` environment variable in `$HOME/.bash_profile`. the new `PATH` setting in the file should look something like this:

```
PATH=$PATH:$HOME/.local/bin:$HOME/bin:/usr/pgsql-9.6/bin
```

6.3 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 shows the Fledge configuration interface. On the left, a sidebar contains a menu with 'Advanced' at the top, followed by 'Scheduler', 'Service Monitor', 'Storage' (expanded), and 'sqlite' (selected). The main content area is titled 'Storage Plugin' and features a refresh icon. It contains five configuration fields, each with a label and a text input box: 'Pool Size' (value 5), 'No. Readings per database' (value 15), 'No. databases to allocate in advance' (value 3), 'Database allocation threshold' (value 1), and 'Database allocation size' (value 2). A blue 'Save' button is located at the bottom right of the configuration area.

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 Fledge 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.

Many factors will impact the performance of a Fledge 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 Fledge itself, however it is possible to tune the way Fledge will use certain resources to achieve better performance within the constraints of a deployment environment.

7.1 South Service Advanced Configuration

The south services within Fledge each have a set of advanced configuration options defined for them. These are accessed by editing the configuration of the south service itself. There is a link titled *Show Advanced Config* to the right of the screen between the main configuration parameters and the *Enabled* option. Clicking on this link will show the following panel of advanced configuration options.

[Hide Advanced Config](#)

Maximum Reading Latency (mS)	<input type="text" value="5000"/>
Maximum buffered Readings	<input type="text" value="100"/>
Reading Rate	<input type="text" value="1"/>
Throttle	<input type="checkbox"/>
Reading Rate Per	<input type="text" value="second"/>
Minimum Log Level	<input type="text" value="warning"/>
Enabled	<input checked="" type="checkbox"/>

- *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.

7.1.1 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.

Notifications Service

Fledge supports an optional service, known as the notification service that adds an event engine to the Fledge installation. The notification services observed data as it flows into the Fledge 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.

8.1 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.

8.1.1 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 Fledge. 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.

8.1.2 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.

8.1.3 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.

8.2 Installing the Notification Service

The notification service is not part of the base Fledge installation and is not a plugin, it is a separate microservice dedicated to the detection of events and the sending of notifications. The service is stored in a separate source repository, *fledge-service-notification* and is packaged as a separate binary package for installation.

8.2.1 Building Notification Service

As with *Fledge* itself there is always the option to build the notification service from the source code repository. This is only recommended if you also built your *Fledge* from source code, if you did not then you should first do this before building the notification, otherwise you should install a binary package of the notification service.

The steps involved in building the notification service, assuming you have already built Fledge itself and the environment variable *FLEDGE_ROOT* points to where you built your *Fledge*, are;

```
$ git clone https://github.com/fledge-iot/fledge-service-notification.git
...
$ cd fledge-service-notification
$ ./requirements.sh
...
$ mkdir build
$ cd build
$ cmake ..
...
$ make
...
```

This will result in the creation of a notification service binary, you now need to copy that binary into the *Fledge* installation. There are two options here, one if you used *make install* to create your installation and the other if you are running directly from the build environment.

If you used *make install* to create your *Fledge* installation then simply run *make install* to install your notification service. This should be run from the *build* directory under the *fledge-service-notification* directory.

```
$ make install
```

Note: You may need to run *make install* under a *sudo* command if your user does not have permissions to write to the installation directory. If you use a *DESTDIR=...* option to the *make install* of *Fledge* then you should use the same *DESTDIR=...* option here also.

If you are running your *Fledge* directly from the build environment, then execute the command

```
$ cp ./C/services/notification/fledge.services.notification $FLEDGE_ROOT/services
```

8.2.2 Installing Notification Service Package

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

If you downloaded your 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 ./fledge-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 ./fledge-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 Fledge packages then you can simply run the command

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

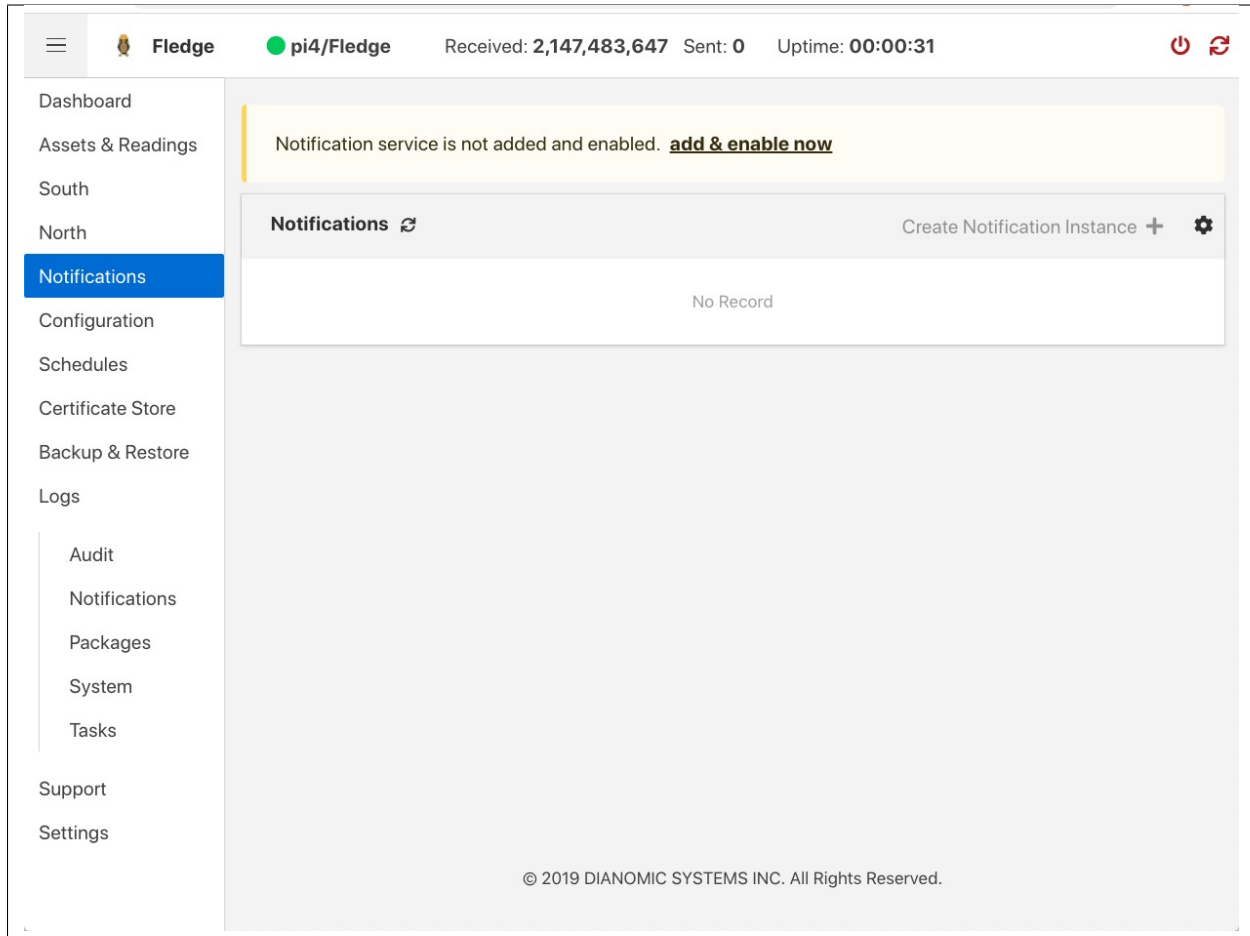
On a Debian/Ubuntu system, or

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

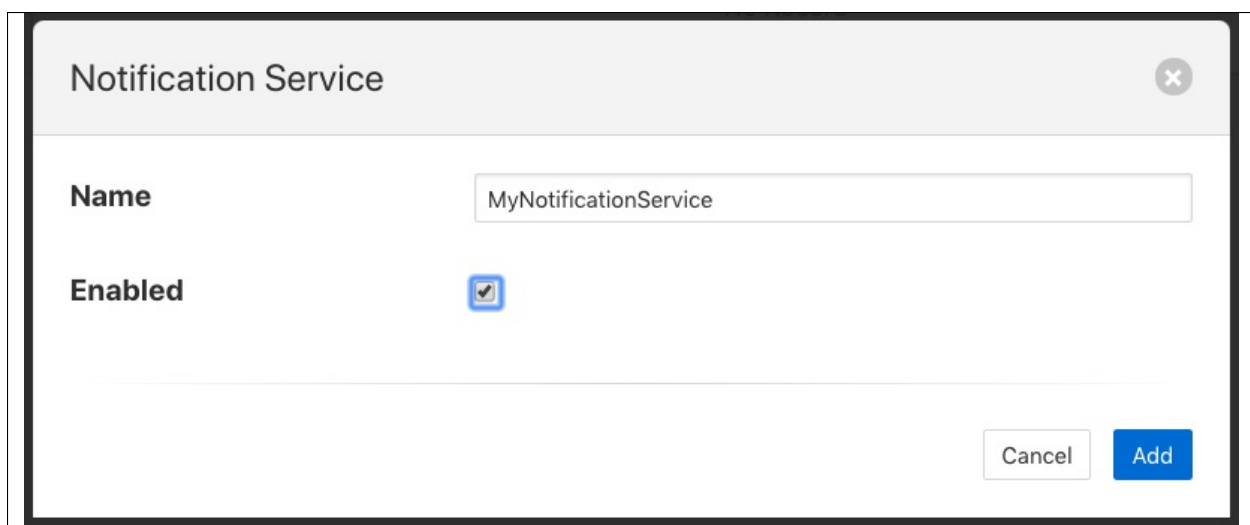
On a RedHat/CentOS system. This will install the latest version of the notification service on your machine.

8.3 Starting The Notification Service

Once installed you must configure Fledge 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 been 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.

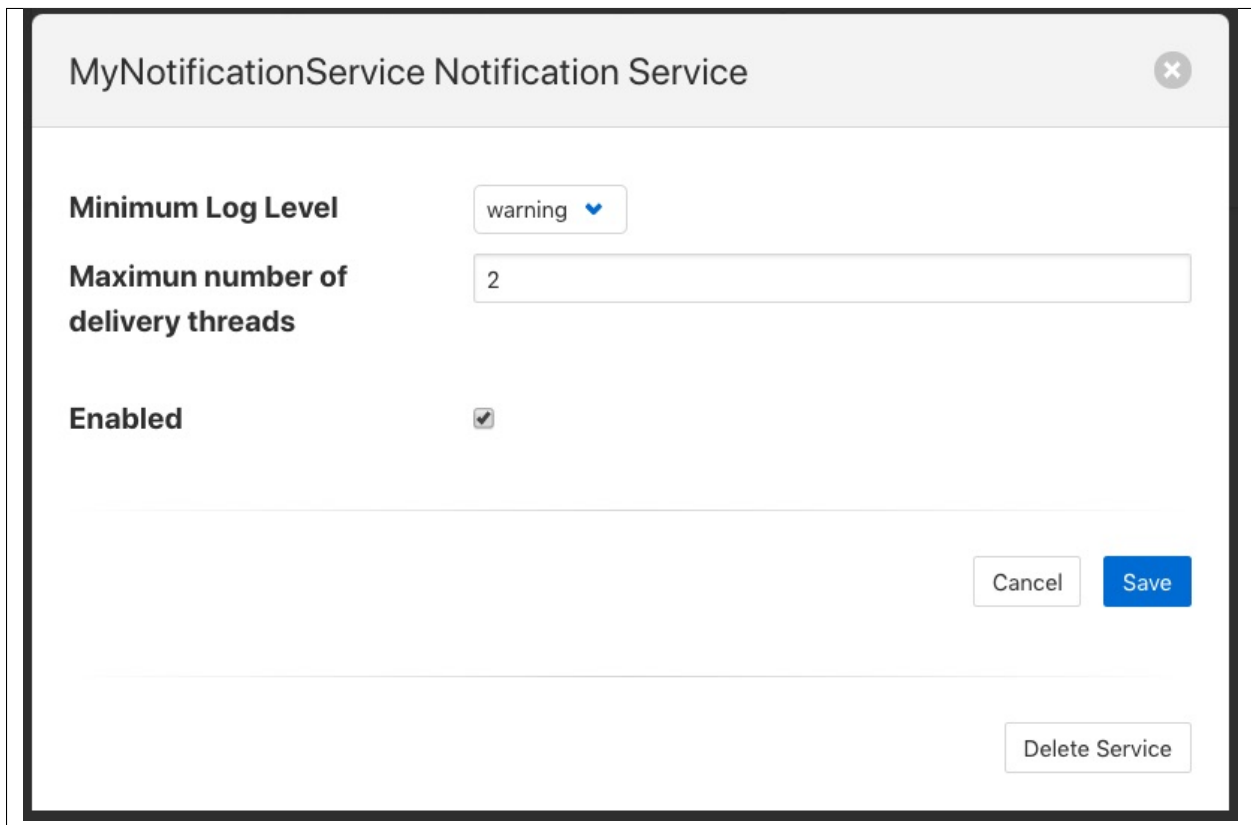


8.4 Configuring The Notification Service

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.

A configuration dialog box titled "MyNotificationService Notification Service" with a close button (X) in the top right corner. The dialog contains three settings: "Minimum Log Level" set to "warning" with a dropdown arrow, "Maximun number of delivery threads" set to "2" in a text input field, and "Enabled" with a checked checkbox. At the bottom right, there are "Cancel" and "Save" buttons. At the bottom center, there is a "Delete Service" button.

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.

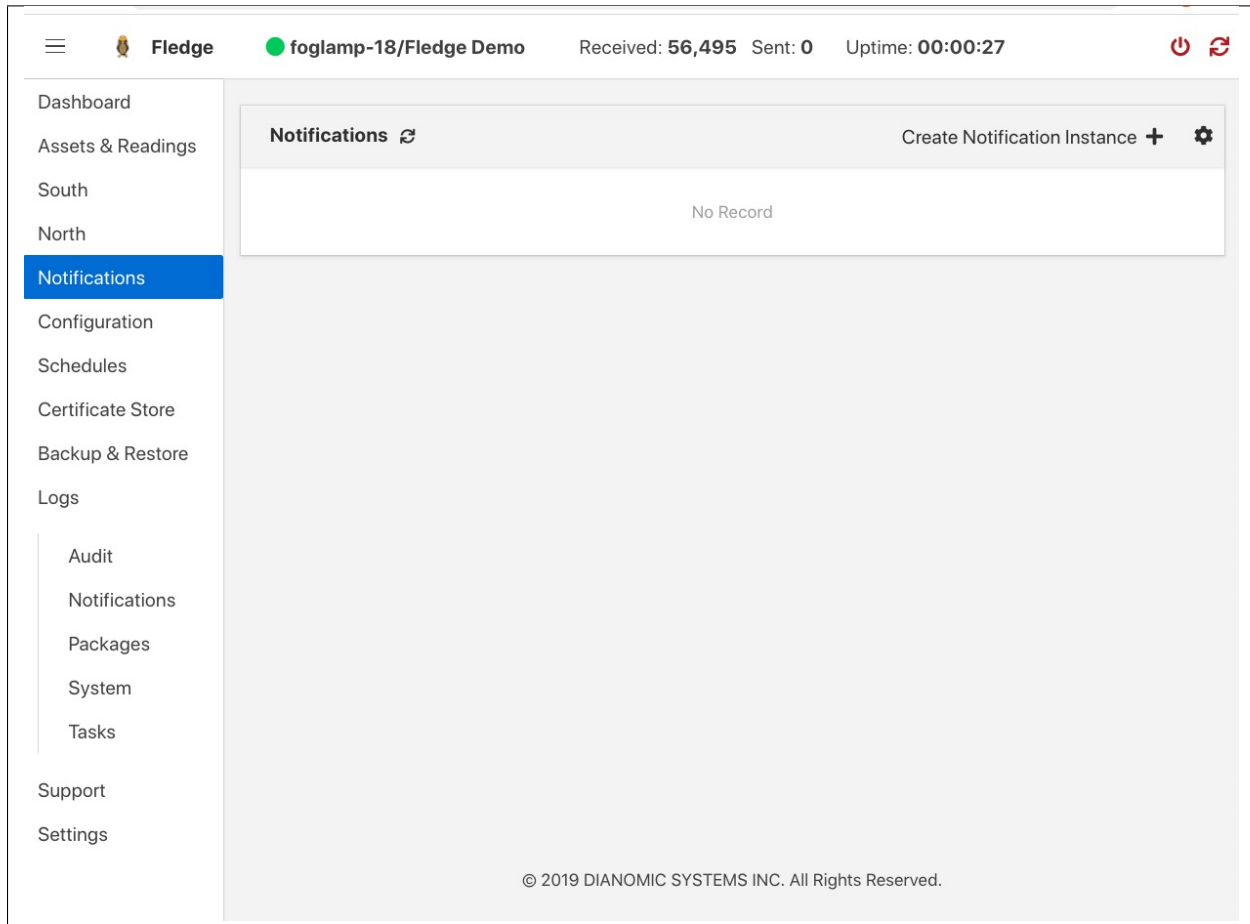
Once you have updated the configuration of the service click on *Save*.

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

8.5 Using The Notification Service

8.5.1 Add A Notification

In order to add a 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 Fledge web interface. The top bar displays the Fledge logo, the instance name 'foglamp-18/Fledge Demo', and system statistics: 'Received: 92,998', 'Sent: 0', and 'Uptime: 18:22:45'. A sidebar on the left contains navigation links: Dashboard, Assets & Readings, South, North, Notifications, Configuration, Schedules, Certificate Store, Backup & Restore, Logs, Audit, Notifications, Packages, System, Tasks, Support, and Settings. The main content area features a progress indicator with four steps: 1. Notification Instance (highlighted), 2. Rule, 3. Delivery Channel, and 4. Done. Below the progress indicator is a form with two input fields: 'Name' (containing 'name') and 'Description' (containing 'description'). At the bottom of the form are 'Back' and 'Next' buttons. The footer text reads '© 2019 DIANOMIC SYSTEMS INC. All Rights Reserved.'

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

This screenshot shows the same 'Notification Instance' setup screen as the previous one, but with user input. The 'Name' field now contains 'Above0.5' and the 'Description' field contains 'Above0.5 notification instance'. The 'Next' button is highlighted in blue, indicating it is the active step in the process.

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 Fledge 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 the Slack Webhook plugin. The form includes three fields: 'Slack Webhook URL' with a text input containing a long URL, 'Message Text' with a text input containing '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.

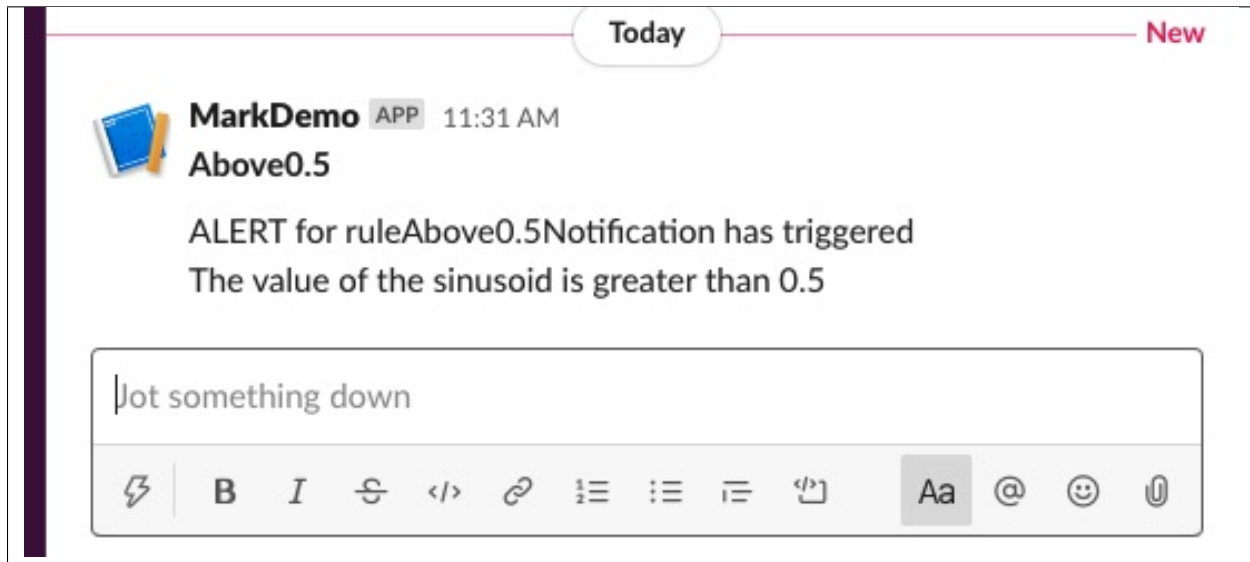
Note: You may disable the delivery of a notification separately to enabling or disabling the notification. This allows you to test the logic of a notification without delivering the notification. Entries will still be made in the notification log when delivery is disabled.

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 has a 'Trigger' dropdown menu set to 'one shot', with a dropdown list showing 'one shot', 'retriggered', and 'toggled'. The 'Enabled' checkbox is checked. A 'Done' button is located 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.

Fledge
● foglamp-18/Fledge Demo
 Received: **98,181** Sent: **0** Uptime: **00:08:34**

Dashboard
Assets & Readings
South
North
Notifications
Configuration
Schedules
Certificate Store
Backup & Restore
Logs

Audit

Notifications

Packages
System
Tasks

Support
Settings

Notification Logs

Count: 13

ALL

INFORMATION

Name

Timestamp	Name	Severity	Source
2020-04-21 10:34:57	Above0.5	INFORMATION	NTFSN
2020-04-21 10:33:57	Above0.5	INFORMATION	NTFSN
2020-04-21 10:32:57	Above0.5	INFORMATION	NTFSN
2020-04-21 10:31:57	Above0.5	INFORMATION	NTFSN
2020-04-21 10:28:49	MyNotificationService	INFORMATION	NTFST
2020-04-21 09:41:48	MyNotificationService	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

You may filter this output using the drop down menus along the top of the page. The list to the left defines the type of event that you filter, clicking on this list will show you the meaning of the different audit types.

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

8.5.2 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.

Fledge

foglamp-18/Fledge Demo

Received: 98,064

Sent: 0

Uptime: 00:06:34

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
<u>Above0.5</u>	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 you 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.

Set Point Control

Fledge 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.

9.1 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.

9.1.1 Set Point

Setting the value within the device is known as a set point action in Fledge. 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.

9.1.2 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.

9.2 Control Paths

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

- As the result of a notification within Fledge itself.
- As a result of a request via the Fledge 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 Fledge.

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

9.2.1 Edge Based Control

Edge based control is the name we use for a class of control applications that take place solely within the Fledge instance at the edge. The data that is required for the control decision to be made is gathered in the Fledge instance, the logic to trigger the control action runs in the Fledge instance and the control action is taken within the Fledge 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 Fledge 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 *fledge-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

Use Control Map ▾

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.

1

2

3

4

Notification Instance

Rule

Delivery Channel

Done

Asset name

MotorTemperature

Datapoint name

temperature

Condition

> ▾

Trigger value

0.0

Evaluation data

Single Item ▾

Window evaluation

Average ▾

Time window

30

Previous

Next

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 within 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" : "$MotorTemperature.desiredSpeed$"
  }
}
```

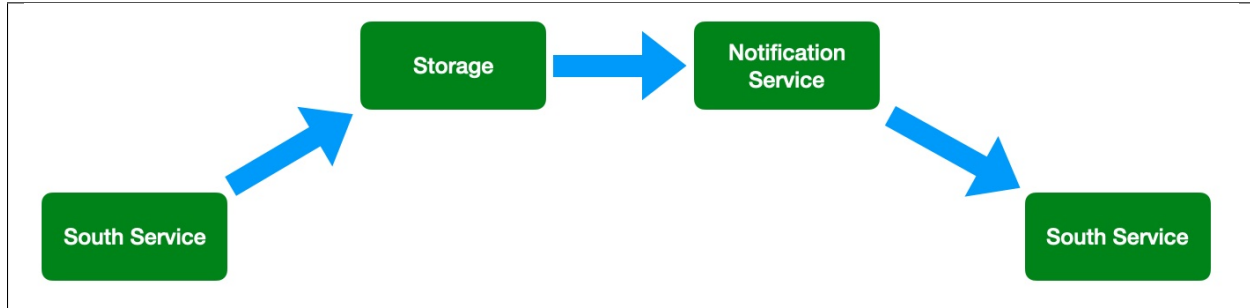
The user then has the choice of adding the desired speed item to the data stored in the north, or adding an asset filter in the north to remove this data point from the data that is sent onward to the north.

Tuning edge control systems

The set point control features of Fledge are not intended to replace real time control applications such as would be seen in PLCs that are typically implemented in ladder logic, however Fledge 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 Fledge. The control decision will be taken in a notification rule and delivered via the *fledge-notify-setpoint* plugin to another south service.

A total of four services within Fledge 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 type="text" value="5000"/>
Maximum buffered Readings	<input type="text" value="100"/>
Reading Rate	<input 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 south 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. Fledge 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 an 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,

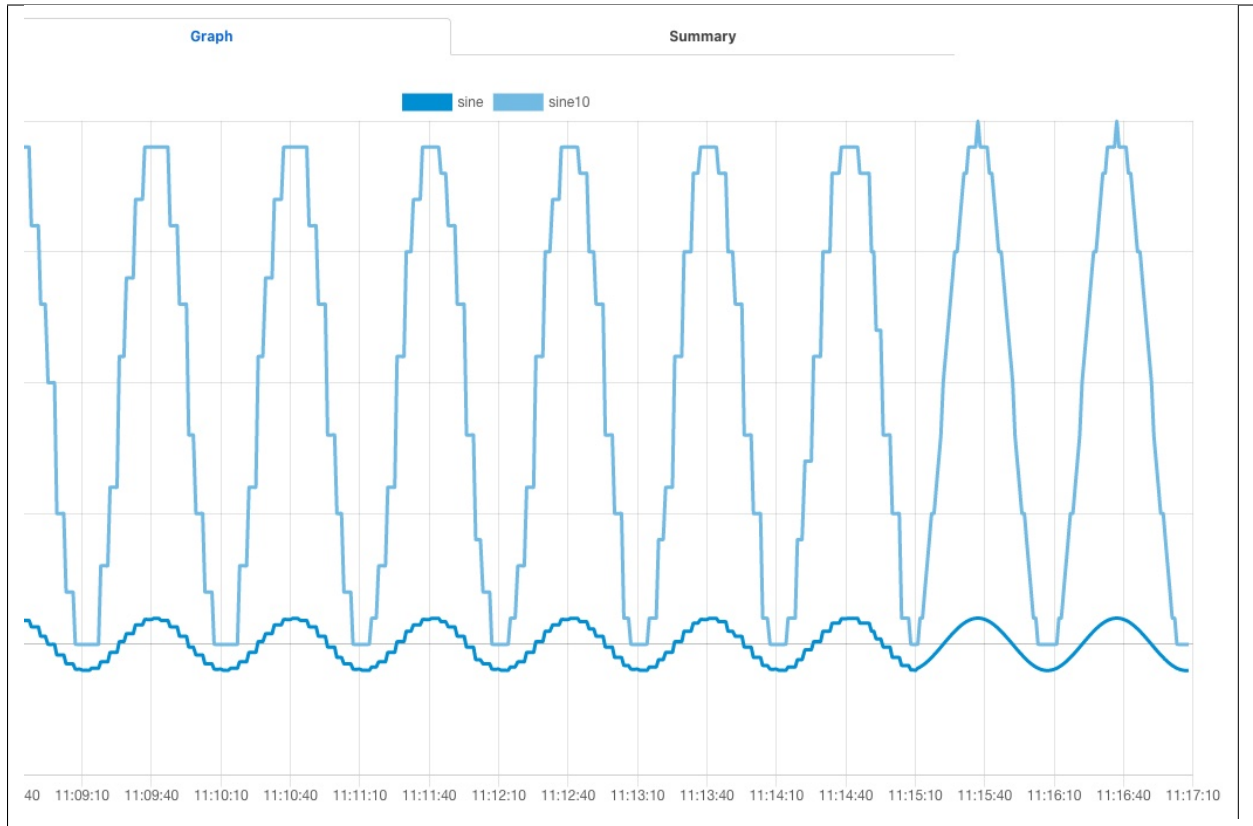


The data observed that is written to the modbus device is not however a clean sine wave as readings have been missed due to the retrigger time eliminating data that arrived in the same buffer.



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.

Troubleshooting the PI-Server integration

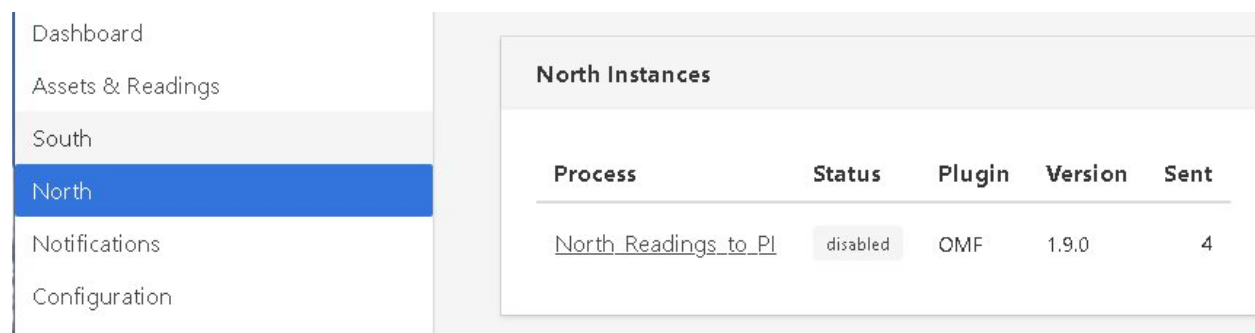
This section describes how to troubleshoot issues with the PI-Server integration using Fledge version $\geq 1.9.1$ and PI Web API 2019 SP1 1.13.0.6518

- *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*

10.1 Log files

Fledge 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 Fledge GUI



```
$ sudo cat /var/log/syslog | grep North_Readings_to_PI
```

Sample message:

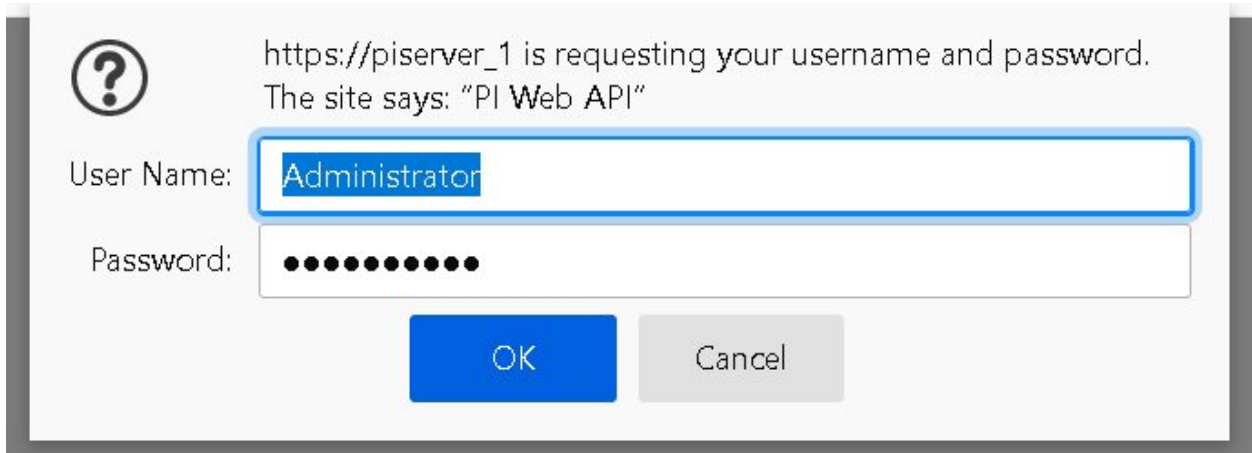
```
user.info, 6,1,Mar 15 08:29:57,localhost,Fledge, North_Readings_to_PI[15506]: INFO: SendingProcess
is starting
```

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
```

10.2 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 Fledge configuration.

The *PI Web API OMF* plugin must be installed to allow the integration with Fledge, in this screenshot the 4th row shows the proper installation of the plugin:

```
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"
  }
}
```

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"
  }
}
```

10.3 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": "F1DSqEe26YHa2kewUjKUX4XLfwV01OLTMyMjbNVTFTMDVQ",
    "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/F1DSqEe26YHa2kewUjKUX4XLfwV01OLTMyMjbNVTFTMDVQ",
      "points": "https://piserver_1/piwabapi/dataservers/F1DSqEe26YHa2kewUjKUX4XLfwV01OLTMyMjbNVTFTMDVQ/points",
      "enumerationsets": "https://piserver_1/piwabapi/dataservers/F1DSqEe26YHa2kewUjKUX4XLfwV01OLTMyMjbNVTFTMDVQ/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": "F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE",
      "Id": 2935,
      "Name": "427300550797094880_measurement_asset_1",
      "Path": "\\WIN-3229M01S95P\\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/F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE",
        "DataServer": "https://piserver_1/piwebapi/dataservers/F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE",
        "Attributes": "https://piserver_1/piwebapi/points/F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE/attributes",
        "InterpolatedData": "https://piserver_1/piwebapi/streams/F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE/interpolated",
        "RecordedData": "https://piserver_1/piwebapi/streams/F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE/recorded",
        "PlotData": "https://piserver_1/piwebapi/streams/F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE/plot",
        "Value": "https://piserver_1/piwebapi/streams/F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE/summary",
        "EndValue": "https://piserver_1/piwebapi/streams/F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE/end"
      }
    },
    {
      "WebId": "F1Dq8e26Yha2kewjKX4KLfwbwAAV010LTHy6tMVTFTMDVQXQyNcHwMDUIMdc5NzcwOTQ4ODBNHUIPQVNVUKVNRUSX0FTU0VUXzE",
      "Id": 2936,
      "Name": "427300550797094880_measurement_asset_2.pl",
      "Path": "\\WIN-3229M01S95P\\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: [FI Web API Help](#) > [AssetServer](#) > [List](#)

```
{
  "Links": {},
  "Items": [
    {
      "WebId": "FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ",
      "Id": "4afc0662-56f7-496e-aae2-1804b5b04a1c",
      "Name": "WIN-3228MU1S05P",
      "Description": "",
      "Path": "\\WIN-3228MU1S05P",
      "IsConnected": true,
      "ServerVersion": "2.10.8.440",
      "ServerTime": "2021-03-15T13:27:31.0903158Z",
      "ExtendedProperties": {},
      "Links": {
        "Self": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ",
        "Databases": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/assetdatabases",
        "NotificationContactTemplates": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/notificationcontacttemplates",
        "NotificationPlugins": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/notificationplugins",
        "SecurityIdentities": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/securityidentities",
        "SecurityMappings": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/securitymappings",
        "UnitClasses": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/unitclasses",
        "AnalysisRulePlugins": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/analysisruleplugins",
        "TimeRulePlugins": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/timeruleplugins",
        "Security": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/security",
        "SecurityEntries": "https://pi-server-1/piwebapi/assetserver/FIRSYgb8SvdWbkmg4hgEtBkFgV010LTHyMjHNVFTMDVQ/securityentries"
      }
    },
    {
      "WebId": "FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ",
      "Id": "6745c0ba-5224-40bc-a166-dd3c876fde90",
      "Name": "pi-server_1",
      "Description": "",
      "Path": "\\pi-server_1",
      "IsConnected": false,
      "ServerVersion": "",
      "ServerTime": null,
      "ExtendedProperties": {},
      "Links": {
        "Self": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ",
        "Databases": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/assetdatabases",
        "NotificationContactTemplates": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/notificationcontacttemplates",
        "NotificationPlugins": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/notificationplugins",
        "SecurityIdentities": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/securityidentities",
        "SecurityMappings": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/securitymappings",
        "UnitClasses": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/unitclasses",
        "AnalysisRulePlugins": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/analysisruleplugins",
        "TimeRulePlugins": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/timeruleplugins",
        "Security": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/security",
        "SecurityEntries": "https://pi-server-1/piwebapi/assetserver/FIRSusxFZyRSvBChZt08h2_ekAUElTRVJWRVJFMQ/securityentries"
      }
    }
  ]
}
```

selecting the database

Help pages: [FI Web API Help](#) > [AssetServer](#) > [GetDatabases](#)

```
{
  "Links": {},
  "Items": [
    {
      "WebId": "FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV",
      "Id": "de1a62d-9eac-4945-82e1-4a1e9ba9d8e",
      "Name": "Configuration",
      "Description": "A store for configuration data.",
      "Path": "\\WIN-3228MU1S05P\\Configuration",
      "ExtendedProperties": {},
      "Links": {
        "Self": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV",
        "Elements": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/elements",
        "EventFrames": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/eventframes",
        "AssetServer": "https://pi-server-1/piwebapi/assetserver/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQ",
        "ElementCategories": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/elementcategories",
        "AttributeCategories": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/attributecategories",
        "TableCategories": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/tablecategories",
        "AnalysisCategories": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/analysiscategories",
        "AnalysisTemplates": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/analysistemplates",
        "EnumerationSets": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/enumerationsets",
        "Tables": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/tables",
        "Security": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/security",
        "SecurityEntries": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/securityentries"
      }
    },
    {
      "WebId": "FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV",
      "Id": "a395ed2d-58c8-49e1-9cd1-6fce75e11e34",
      "Name": "Ecoglamp",
      "Description": "",
      "Path": "\\WIN-3228MU1S05P\\Ecoglamp",
      "ExtendedProperties": {},
      "Links": {
        "Self": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV",
        "Elements": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/elements",
        "EventFrames": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/eventframes",
        "AssetServer": "https://pi-server-1/piwebapi/assetserver/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQ",
        "ElementCategories": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/elementcategories",
        "AttributeCategories": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/attributecategories",
        "TableCategories": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/tablecategories",
        "AnalysisCategories": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/analysiscategories",
        "AnalysisTemplates": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/analysistemplates",
        "EnumerationSets": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/enumerationsets",
        "Tables": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/tables",
        "Security": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/security",
        "SecurityEntries": "https://pi-server-1/piwebapi/assetdatabases/FIRDYgb8SvdWbkmg4hgEtBkFgLeV08hY4UmcW_ODWEHAV010LTHyMjHNVFTMDVQXEPF08BTV/securityentries"
      }
    }
  ]
}
```

Proceed with the drill down operation up to the desired level/asset.

10.4 Error messages and causes

Some error messages and causes:

10.4. 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	Fledge 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	Fledge 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	Fledge is able to interact with PI Web API but there is an attempt to store data in a PI Point that does not exist.

10.5 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 *fledge-filter-asset* on the new north instance
- configure *fledge-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 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 Fledge 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 Fledge will be sent

Fledge 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 Fledge supports.
- Extend the set of services to which Fledge will push accumulated data gathered from those sensors.
- The mechanism by which Fledge buffers data internally.
- Filter plugins may be used to augment, edit or remove data as it flows through Fledge.
- 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 Fledge.

This chapter presents the plugins that are bundled with Fledge, 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 Fledge.

11.1 Plugins

In this version of Fledge you have six types of plugins:

- **South Plugins** - They are responsible for communication between Fledge and the sensors and actuators they support. Each instance of a Fledge 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 Fledge 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 Fledge. 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.

11.1.1 Plugins in this version of Fledge

This version of Fledge 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	OSIsoft 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.

11.1.2 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 **Fledge is running**. This is the required method because the package executed pre and post-installation tasks that require Fledge to run.
- When the plugin is available as **source code**, it should be installed when **Fledge is either running or not**. You will want to manually move the plugin code into the right location where Fledge is installed, add pre-requisites and execute the REST commands necessary to start the plugin **after** you have started Fledge 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 fledge.service
fledge.service - LSB: Fledge
   Loaded: loaded (/etc/init.d/fledge; 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)
```

(continues on next page)

(continued from previous page)

```

CGroup: /system.slice/fledge.service
└─13741 python3 -m fledge.services.core
└─13746 /usr/local/fledge/services/storage --address=0.0.0.0 --port=40138

May 16 01:36:09 ubuntu python3[13741]: Fledge[13741] INFO: scheduler: fledge.services.
↳core.scheduler.scheduler: Process started: Schedule 'stats collection' process
↳'stats coll
                                ['tasks/statistics', '--port=40138', '--
↳address=127.0.0.1', '--name=stats collector']
...
Fledge v1.3.1 running.
Fledge Uptime: 266 seconds.
Fledge records: 0 read, 0 sent, 0 purged.
Fledge does not require authentication.
=== Fledge services:
fledge.services.core
=== Fledge tasks:
$
$ sudo cp fledge-south-openweathermap-1.2-x86_64.deb /var/cache/apt/archives/.
$ sudo apt install /var/cache/apt/archives/fledge-south-openweathermap-1.2-x86_64.deb
Reading package lists... Done
Building dependency tree
Reading state information... Done
Note, selecting 'fledge-south-openweathermap' instead of '/var/cache/apt/archives/
↳fledge-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
  fledge-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 fledge-south-openweathermap.
(Reading database ... 211747 files and directories currently installed.)
Preparing to unpack .../fledge-south-openweathermap-1.2-x86_64.deb ...
Unpacking fledge-south-openweathermap (1.2) ...
Setting up fledge-south-openweathermap (1.2) ...
openweathermap plugin installed.
$
$ fledge status
Fledge v1.3.1 running.
Fledge Uptime: 271 seconds.
Fledge records: 36 read, 0 sent, 0 purged.
Fledge does not require authentication.
=== Fledge services:
fledge.services.core
fledge.services.south --port=42066 --address=127.0.0.1 --name=openweathermap
=== Fledge tasks:
$

```

You may also install new plugins directly from within the Fledge GUI, however you will need to have setup your

Linux machine to include the Fledge package repository in the list of repositories the Linux package manager searches for new packages.

11.2 Writing and Using Plugins

A plugin has a small set of external entry points that must exist in order for Fledge 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 indepth discussion of writing plugins in C/C++ will then follow.

11.2.1 Common Fledge 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 initialised and used. It allows Fledge 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 Fledge interacts with the plugin.
- **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, Storage, Filter, Rule or Delivery.

Note: If you browse the Fledge 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 Fledge 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 Fledge, each plugin will return a set of configuration items that it wishes to use, this will then be used to extend the set of Fledge configuration items. This structure, a JSON document, includes default values but no actual values for each configuration option. The first time Fledge'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'
  }
}
```

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
};

/**
 * Return the information about this plugin
```

(continues on next page)

(continued from previous page)

```
*/
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"
    }
});
```

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['gpioin']['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.

```

(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
    Raises:
        """

    new_handle = new_config['gpioin']['value']
    return new_handle

```

In C/C++ the *plugin_reconfigure* class 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.

11.3 South Plugins

South plugins are used to communicate with sensors and actuators, there are two modes of plugin operation; *asyncio* and *polled*.

11.3.1 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)
        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
```

11.3.2 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 Fledge, 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 Handler

The async handler is defined for incoming message has the responsibility of taking the sensor data and ingesting that into Fledge. Unlike the poll mechanism, this is done from within the handler rather than by passing the data back to the South service itself. A convenient method exists for ingesting readings, *Ingest.add_readings*. This call is passed an asset, timestamp, key and readings document for the asset and will do everything else required to make sure the readings are stored in the Fledge buffer. In the case of our HTTP based example above, the code would create the items needed to generate the arguments to the *Ingest.add_readings* call, by creating data items and retrieving them from the payload sent by the sensor.

```
try:
    if not Ingest.is_available():
        increment_discarded_counter = True
        message = {'busy': True}
    else:
        payload = await request.json()

        asset = 'SensorPhone'
        timestamp = str(datetime.now(tz=timezone.utc))
        messages = payload.get('messages')

        if not isinstance(messages, list):
            raise ValueError('messages must be a list')

        for readings in messages:
            key = str(uuid.uuid4())
            await Ingest.add_readings(asset=asset, timestamp=timestamp, key=key,
            ↳ readings=readings)
except ...
```

It would then respond to the HTTP request and return. Since the handler is embedded in the event loop this will happen in the context of a coroutine and would happen each time a new POST request is received.

```
message['status'] = code
return web.json_response(message)
```

11.3.3 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 -*-

# FLEDGE_BEGIN
# See: http://fledge-iot.readthedocs.io/
# FLEDGE_END

""" Plugin for a DHT11 temperature and humidity sensor attached directly
    to the GPIO pins of a Raspberry Pi

    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 fledge must be able to access /dev/gpiomem,
    the default access for this is owner and group read/write. Either
    Fledge must be added to the group or the permissions altered to
    allow Fledge access to the device.
    """

from datetime import datetime, timezone
import uuid

from fledge.common import logger
from fledge.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'
    }
}
```

(continues on next page)

(continued from previous page)

```

_LOGGER = logger.setup(__name__)
""" Setup the access to the logging system of Fledge """

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
    }

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:

```

(continues on next page)

(continued from previous page)

```

        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

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['gpioin']['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 Fledge and Adding the Plugin

If you have not built Fledge yet, follow the steps described . After the build, you can optionally install Fledge following steps.

- If you have started Fledge from the build directory, copy the structure of the *fledge-south-dht11/python/* directory into the *python* directory:


```
$ cd ~/Fledge
$ cp -R ~/fledge-south-dht11/python/fledge/plugins/south/dht11 python/fledge/plugins/
↳ south/
$
```

- If you have installed Fledge by executing `sudo make install`, copy the structure of the *fledge-south-dht11/python/* directory into the installed *python* directory:

```
$ sudo cp -R ~/fledge-south-dht11/python/fledge/plugins/south/dht11 /usr/local/fledge/
↳ python/fledge/plugins/south/
$
```

Note: If you have installed Fledge 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/fledge/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/fledge/service | jq
```

Let's see what we have collected so far:

```
$ curl -s http://localhost:8081/fledge/asset | jq
[
  {
    "count": 158,
    "asset_code": "dht11"
  }
]
$
```

Finally, let's extract some values:

```
$ curl -s http://localhost:8081/fledge/asset/dht11?limit=5 | jq
[
  {
    "timestamp": "2017-12-30 14:41:39.672",
    "reading": {
      "temperature": 19,
      "humidity": 62
    }
  },
  {

```

(continues on next page)

(continued from previous page)

```
"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
  }
},
{
  "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/fledge/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
    }
  }
],
```

(continues on next page)

(continued from previous page)

```

{
  "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
  }
}
]
$

```

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.



11.4 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 Fledge is not aware in which language the plugin has been written.

11.4.1 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 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;

    if (!handle)
        throw runtime_error("Bad plugin handle");
    return modbus->takeReading();
}
```

11.4.2 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 Fledge. 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 Fledge 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)
{

```

(continues on next page)

(continued from previous page)

```
        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)
    {
        (*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);
    }
```

11.4.3 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 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 to 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;

    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));
        }
    }
}
```

(continues on next page)

(continued from previous page)

```

    }
    Logger::getLogger()->info("Reseeded random number generator");
    return true;
}
Logger::getLogger()->error("Unrecognised operation %s", operation.c_str());
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.

11.4.4 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:

```

/*
 * Fledge 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>

```

(continues on next page)

(continued from previous page)

```

#include <string>
#include <logger.h>
#include <plugin_exception.h>
#include <config_category.h>
#include <rapidjson/document.h>
#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
 */

```

(continues on next page)

(continued from previous page)

```

PLUGIN_INFORMATION *plugin_info()
{
    return &info;
}

/**
 * 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

```

(continues on next page)

(continued from previous page)

```

*/
void plugin_shutdown(PLUGIN_HANDLE *handle)
{
    DHT11 *dht11 = (DHT11*)handle;
    delete dht11;
}
};

```

The full source code, including the *DHT11* class can be found in [GitHub](#)

Building Fledge and Adding the Plugin

If you have not built Fledge yet, follow the steps described . After the build, you can optionally install Fledge following steps.

- Clone the *fledge-south-dht* repository

```

$ git clone https://github.com/fledge-iot/fledge-south-dht.git
...
$

```

- Set the environment variable `FLEDGE_ROOT` to the directory in which you built Fledge

```

$ export FLEDGE_ROOT=~/.fledge
$

```

- Go to the location in which you cloned the *fledge-south-dht* repository and create a build directory and run `cmake` in that directory

```

$ cd ~/.fledge-south-dht
$ mkdir build
$ cd build
$ cmake ..
...
$

```

- Now make the plugin

```

$ make
$

```

- If you have started Fledge from the build directory, copy the plugin into the destination directory

```

$ mkdir -p $FLEDGE_ROOT/plugins/south/dht
$ cp libdht.so $FLEDGE_ROOT/plugins/south/dht
$

```

- If you have installed Fledge by executing `sudo make install`, copy the plugin into the destination directory

```

$ sudo mkdir -p /usr/local/fledge/plugins/south/dht
$ sudo cp libdht.so /usr/local/fledge/plugins/south/dht
$

```

Note: If you have installed Fledge 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/fledge/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.

11.5 C++ Support Classes

A number of support classes exist within the common library that forms part of every Fledge plugin.

11.5.1 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; }
↳ ;
```

(continues on next page)

(continued from previous page)

```

        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);
        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 Fledge storage service and are used by the various services and plugins within Fledge.

- **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 Fledge.
- **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();

```

(continues on next page)

(continued from previous page)

```

    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:
    /**
     * Constructors for the various types
     */
    DatapointValue(const std::string& value;
    DatapointValue(const long value);
    DatapointValue(const double value);
    DatapointValue(const std::vector<double>& values);
    DatapointValue(std::vector<Datapoint*>& values, bool isDict)
    DatapointValue(const DatapointValue& obj)

    DatapointValue& operator=(const DatapointValue& rhs)
    ~DatapointValue();

    void deleteNestedDPV();

    /**
     * Set the value for the various types
     */
    void setValue(long value);
    void setValue(double value);

    /**
     * Return the value as the various types
     */
    std::string toString() const;
    long toInt() const;
    double toDouble() const;

    typedef enum DatapointTag
    {
        T_STRING,
        T_INTEGER,
        T_FLOAT,
        T_FLOAT_ARRAY,
        T_DP_DICT,
        T_DP_LIST
    } dataTagType;
    dataTagType getType() const;
    std::string getTypeStr() const;
    std::vector<Datapoint*>& getDpVec();
}

```


11.5.2 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,
            ↪const std::string def,
            ↪const std::string description,
            ↪std::string& value,
            ↪options);
        void removeItems();
        void removeItemsType(ItemType type);
        void keepItemsType(ItemType type);
        bool extractSubcategory(ConfigCategory &
            ↪subCategories);
        void setDescription(const std::string&
            ↪description);
        std::string getName() const;
        std::string getDescription() const;
        unsigned int getCount() const;
        bool itemExists(const std::string& name)
            ↪const;
        bool setItemDisplayName(const std::string&
            ↪name, const std::string& displayName);
        std::string getValue(const std::string& name)
            ↪const;
        std::string getType(const std::string& name)
            ↪const;
        std::string getDescription(const std::string&
            ↪name) const;
        std::string getDefault(const std::string& name)
            ↪const;
        bool setDefault(const std::string& name,
            ↪const std::string& value);
};
```

(continues on next page)

(continued from previous page)

```

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;
ConfigCategory&
ConfigCategory&
void
void
std::string
↪itemName) const;
enum ItemAttribute
↪MANDATORY_ATTR, FILE_ATTR;
std::string
↪itemName,
ItemAttribute
↪itemAttribute) const;
}

getDisplayNames(const std::string& name)
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& name)
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)
operator=(ConfigCategory const& rhs);
operator+=(ConfigCategory const& rhs);
setItemValueFromDefault();
checkDefaultValuesOnly() const;
itemToJson(const std::string& name,
ItemAttribute attr) const;
{ ORDER_ATTR, READONLY_ATTR,
getItemAttribute(const std::string& name,
ItemAttribute attr) const;

```

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

11.5.3 Logger

The *Logger* class is used to write entries to the syslog system within Fledge. 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.

11.6 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 Fledge 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 Fledge. 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 Fledge 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 Fledge 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 many 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 *fledge-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.

11.7 North Plugins

North plugins are used in North tasks and microservices to extract data buffered in Fledge and send it Northbound, i.e. to a server or a service in the Cloud or in an Enterprise data center. We currently have two North plugins, one to send data to an OSIsoft PI Server and one to the OSIsoft Cloud Service.

11.7.1 The OMF Plugin

The OMF Plugin is used by a North task to send data to an OSIsoft PI server via a PI Connector Relay or PI Web API, it can also send to Edge Data Store or OSIsoft Cloud Services. All these destinations share a single protocol for communication, OMF. OMF stands for OSIsoft Message Format, it is the JSON format defined by OSIsoft to send IoT data to a PI server via a Connector Relay server.

The plugin is designed to send two streams of data:

- The data collected by South microservices and buffered into Fledge
- The statistics generated by Fledge

The streams are managed by two different North tasks using the same plugin, but with a different configuration. The two tasks are registered in the list of scheduled jobs and they can be identified using the `schedule` API call:

```
$ curl -sX GET http://localhost:8081/fledge/schedule
{
  "schedules": [
    {
      "id": "ef8bd42b-da9f-47c4-ade8-751ce9a504be",
```

(continues on next page)

(continued from previous page)

```

    "name": "OMF to PI north",
    "processName": "north_c",
    "type": "INTERVAL",
    "repeat": 30.0,
    "time": 0,
    "day": null,
    "exclusive": true,
    "enabled": false
  },
  {
    "id": "27501b35-e0cd-4340-afc2-a4465fe877d6",
    "name": "Stats OMF to PI north",
    "processName": "north_c",
    "type": "INTERVAL",
    "repeat": 30.0,
    "time": 0,
    "day": null,
    "exclusive": true,
    "enabled": true
  },
  ...
]
}

```

The output of API call above shows three interesting tasks: the two tasks associated to the OMF plugin, the one to send data (*OMF to PI north*) and the one to send statistics (*Stats OMF to PI north*).

The two scheduled tasks are associated to two configuration items that can be retrieved using the category API call. The items are named OMF to PI north and Stats OMF to PI north.

```

$ curl -sX GET http://localhost:8081/fledge/category/OMF%20to%20PI%20north
{
  "enable": {
    "description": "A switch that can be used to enable or disable execution of the
↪ sending process.",
    "type": "boolean",
    "readonly": "true",
    "default": "true",
    "value": "true"
  },
  "streamId": {
    "description": "Identifies the specific stream to handle and the related
↪ information, among them the ID of the last object streamed.",
    "type": "integer",
    "readonly": "true",
    "default": "0",
    "value": "4",
    "order": "16"
  },
  "plugin": {
    "description": "PI Server North C Plugin",
    "type": "string",
    "default": "OMF",
    "readonly": "true",
    "value": "OMF"
  },
  "source": {

```

(continues on next page)

(continued from previous page)

```

        "description": "Defines the source of the data to be sent on the stream, this
↪may be one of either readings, statistics or audit.",
        "type": "enumeration",
        "options": [
            "readings",
            "statistics"
        ],
        "default": "readings",
        "order": "5",
        "displayName": "Data Source",
        "value": "readings"
    },
    ...}
$ curl -sX GET    http://localhost:8081/fledge/category/Stats%20OMF%20to%20PI%20north
{
  "enable": {
    "description": "A switch that can be used to enable or disable execution of the
↪sending process.",
    "type": "boolean",
    "readonly": "true",
    "default": "true",
    "value": "true"
  },
  "streamId": {
    "description": "Identifies the specific stream to handle and the related
↪information, among them the ID of the last object streamed.",
    "type": "integer",
    "readonly": "true",
    "default": "0",
    "value": "5",
    "order": "16"
  },
  "plugin": {
    "description": "PI Server North C Plugin",
    "type": "string",
    "default": "OMF",
    "readonly": "true",
    "value": "OMF"
  },
  "source": {
    "description": "Defines the source of the data to be sent on the stream, this may
↪be one of either readings, statistics or audit.",
    "type": "enumeration",
    "options": [
        "readings",
        "statistics"
    ],
    "default": "readings",
    "order": "5",
    "displayName": "Data Source",
    "value": "statistics"
  },
  ...}
$
```

In order to activate the tasks, you must change their status. First you must collect their id (from the GET method of the schedule API call), then you must use the IDs with the PUT method of the same call:

```
$ curl -sX PUT http://localhost:8081/fledge/schedule/ef8bd42b-da9f-47c4-ade8-751ce9a504be -d '{ "enabled" : true}'
{
  "schedule": {
    "id": "ef8bd42b-da9f-47c4-ade8-751ce9a504be",
    "name": "OMF to PI north",
    "processName": "north_c",
    "type": "INTERVAL",
    "repeat": 30,
    "time": 0,
    "day": null,
    "exclusive": true,
    "enabled": true
  }
}
$ curl -sX PUT http://localhost:8081/fledge/schedule/27501b35-e0cd-4340-afc2-a4465fe877d6 -d '{ "enabled" : true}'
{
  "schedule": {
    "id": "27501b35-e0cd-4340-afc2-a4465fe877d6",
    "name": "Stats OMF to PI north",
    "processName": "north_c",
    "type": "INTERVAL",
    "repeat": 30,
    "time": 0,
    "day": null,
    "exclusive": true,
    "enabled": true
  }
}
$
```

At this point, the configuration has been enriched with default values of the tasks:

```
$ curl -sX GET http://localhost:8081/fledge/category/OMF%20to%20PI%20north
{
  "enable": {
    "description": "A switch that can be used to enable or disable execution of the_
    ↪sending process.",
    "type": "boolean",
    "readonly": "true",
    "default": "true",
    "value": "true"
  },
  "streamId": {
    "description": "Identifies the specific stream to handle and the related_
    ↪information, among them the ID of the last object streamed.",
    "type": "integer",
    "readonly": "true",
    "default": "0",
    "value": "4",
    "order": "16"
  },
  "plugin": {
    "description": "PI Server North C Plugin",
    "type": "string",
    "default": "OMF",

```

(continues on next page)

(continued from previous page)

```

    "readonly": "true",
    "value": "OMF"
  },
  "source": {
    "description": "Defines the source of the data to be sent on the stream, this_
↪may be one of either readings, statistics or audit.",
    "type": "enumeration",
    "options": [
      "readings",
      "statistics"
    ],
    "default": "readings",
    "order": "5",
    "displayName": "Data Source",
    "value": "readings"
  },
  ...}
$ curl -sX GET http://localhost:8081/fledge/category/Stats%20OMF%20to%20PI%20north
{
  "enable": {
    "description": "A switch that can be used to enable or disable execution of the_
↪sending process.",
    "type": "boolean",
    "readonly": "true",
    "default": "true",
    "value": "true"
  },
  "streamId": {
    "description": "Identifies the specific stream to handle and the related_
↪information, among them the ID of the last object streamed.",
    "type": "integer",
    "readonly": "true",
    "default": "0",
    "value": "5",
    "order": "16"
  },
  "plugin": {
    "description": "PI Server North C Plugin",
    "type": "string",
    "default": "OMF",
    "readonly": "true",
    "value": "OMF"
  },
  "source": {
    "description": "Defines the source of the data to be sent on the stream, this may_
↪be one of either readings, statistics or audit.",
    "type": "enumeration",
    "options": [
      "readings",
      "statistics"
    ],
    "default": "readings",
    "order": "5",
    "displayName": "Data Source",
    "value": "statistics"
  },
  ...}

```

(continues on next page)

(continued from previous page)

\$

OMF Plugin Configuration

The following table presents the list of configuration options available for the task that sends data to OMF (category *OMF to PI north*):

Item	Type	Default	Description
AFMap	JSON	{ }	Defines a set of rules to address where assets should be placed in the AF hierarchy.
compression	boolean	true	Compress readings data before sending to PI server
DefaultAFLocation	integer	/fledge/data_piwebapi/default	Defines the hierarchies tree in Asset Framework in which the assets will be created, each level is separated by /, PI Web API only.
enable	boolean	True	A switch that can be used to enable or disable execution of the sending process.
formatInteger	string	int64	OMF format property to apply to the type Integer.
formatNumber	string	float64	OMF format property to apply to the type Number
notBlockingErrors	JSON	“{ “errors400” : [“Redefinition of the type with the same ID is not allowed”, “Invalid value type for the property”, “Property does not exist in the type definition”, “Container is not defined”, “Unable to find the property of the container of type”] }”	These errors are considered not blocking in the communication with the PI Server, the sending operation will proceed with the next block of data if one of these is encountered.
OCSCClientSecret	boolean	ocs_client_secret	Client secret associated to the specific OCS account, it is used to authenticate the source for using the OCS API.
OCSCClientId	string	ocs_client_id	Client id associated to the specific OCS account, it is used to authenticate the source for using the OCS API.
OCSTenantId	string	ocs_tenant_id	Tenant id associated to the specific OCS account
OCSNamespace	string	name_space	Specifies the OCS namespace where the information are stored and it is used for the interaction with the OCS API.
OMFHttpTimeout	integer	10	Timeout in seconds for the HTTP operations with the OMF PI Connector Relay
OMFMaxRetry	integer	1	Seconds between each retry for the communication with the OMF PI Connector Relay, NOTE : the time is doubled at each attempt.
PIWebAPIKerberosKeytabFileName	string	piwebapi_kerberos_https.keytab	Keytab file name used for Kerberos authentication in PI Web API.
PIWebAPIAuthenticationMethod	string	anonymous	Defines the authentication method to be used with the PI Web API.
PIWebAPIPassword	password	password	Password of the user of PI Web API to be used with the basic access authentication.
PIWebAPIUserId	string	user_id	User id of PI Web API to be used with the basic access authentication.
PIServerEndpoint	enumeration	Connector Relay	Select the endpoint among PI Web API, connector Relay, OSIsoft Cloud Services or Edge Data Store
plugin	string	OMF	PI Server North C Plugin

The following table presents the list of configuration options available for the task that sends statistics to OMF (category *Stats OMF to PI north*):

Item	Type	Default	Description
AFMap	JSON	{ }	Defines a set of rules to address where assets should be placed in the AF hierarchy.
compression	boolean	true	Compress readings data before sending to PI server
DefaultAFLocation	integer	/fledge/data_piwebapi/default	Defines the hierarchies tree in Asset Framework in which the assets will be created, each level is separated by /, PI Web API only.
enable	boolean	True	A switch that can be used to enable or disable execution of the sending process.
formatInteger	string	int64	OMF format property to apply to the type Integer.
formatNumber	string	float64	OMF format property to apply to the type Number
notBlockingErrors	JSON	“{ “errors400” : [“Redefinition of the type with the same ID is not allowed”, “Invalid value type for the property”, “Property does not exist in the type definition”, “Container is not defined”, “Unable to find the property of the container of type”] }”	These errors are considered not blocking in the communication with the PI Server, the sending operation will proceed with the next block of data if one of these is encountered.
OCSClientSecret	boolean	ocs_client_secret	Client secret associated to the specific OCS account, it is used to authenticate the source for using the OCS API.
OCSClientId	string	ocs_client_id	Client id associated to the specific OCS account, it is used to authenticate the source for using the OCS API.
OCSTenantId	string	ocs_tenant_id	Tenant id associated to the specific OCS account
OCSNamespace	string	name_space	Specifies the OCS namespace where the information are stored and it is used for the interaction with the OCS API.
OMFHttpTimeout	integer	10	Timeout in seconds for the HTTP operations with the OMF PI Connector Relay
OMFMaxRetry	integer	1	Seconds between each retry for the communication with the OMF PI Connector Relay, NOTE : the time is doubled at each attempt.
PIWebAPIKerberosKeytabFileName	string	piwebapi_kerberos_https.keytab	Keytab file name used for Kerberos authentication in PI Web API.
PIWebAPIAuthenticationMethod	enum	anonymous	Defines the authentication method to be used with the PI Web API.
PIWebAPIPassword	password	password	Password of the user of PI Web API to be used with the basic access authentication.
PIWebAPIUserId	string	user_id	User id of PI Web API to be used with the basic access authentication.
PIServerEndpoint	enumeration	Connector Relay	Select the endpoint among PI Web API, connector Relay, OSIsoft Cloud Services or Edge Data Store
plugin	string	OMF	PI Server North C Plugin

The last parameter to review is the *OMF Type*. The call is the GET method `fledge/category/OMF_TYPES`, which returns an integer value that identifies the measurement type:

```
$ curl -sX GET http://localhost:8081/fledge/category/OMF_TYPES
{
  "type-id": {
    "description": "Identify sensor and measurement types",
    "type": "integer",
    "default": "0001",
    "value": "0001"
  }
}
```

If you change the value, you can easily identify the set of data sent to and then stored into PI.

Changing the OMF Plugin Configuration

Before you send data to the PI server, it is likely that you need to apply more changes to the configuration. The most important items to change are:

- **URL** : the URL to the PI Connector Relay OMF. It is usually composed by the name or address of the Windows server where the Connector Relay service is running, the port associated to the service and the ingress/messages API call. The communication is via HTTPS protocol.
- **producerToken** : the token provided by the Data Collection Manager when the PI administrator sets the use of Fledge.
- **type-id** : the measurement type for the stream of data.
- **source** : this parameter should be set to *readings* (default) when the plugin is used to send data collected by South microservices, and to *statistics* when the plugin is used to send Fledge statistics to the PI system.

An example of the changes to apply to the plugins to send data to the PI system is available [here](#).

Data in the PI System

Once the North plugins have been set properly, you should expect to see data automatically sent and stored in the PI Server. More specifically, the process of the plugin is the following:

- **Assets** buffered in Fledge are stored as *elements* in the PI System. - *PI Asset Framework* is automatically update with the new assets. - JSON objects captured as part of the *reading* in Fledge become *attributes* in the PI Data Archive
- The **Producer Token** is used to authenticate and create the hierarchy of elements in the *PI Asset Framework*
- The configuration object named as **Static Data** is added as a set of *attributes* in the PI Data Archive

System	Object	Value
Fledge	Producer Token	readings_001
	OMF Type	001
	Static Data	{ "Company": "Dianomic", "Location": "Palo Alto" }
	Asset	fogbench/accelerometer
	Reading	[{"reading":{"y":1,"z":1,"x":-1}, "timestamp":"2018-05-14 19:27:06.788"]
PI	Element Template	[OMF.readings_001 Connector.0001_fogbench/accelerometer_typename_sensor]
	Attribute Template	[OMF.readings_001 Connector.0001_fogbench/accelerometer_typename_sensor]
		Company Configuration Item, Excluded, String
		Location Configuration Item, Excluded, String
		x Excluded, Int64
		y Excluded, Int64
		z Excluded, Int64
	Element	fledge > readings_001 > fogbench/accelerometer
	Attributes	Company Dianomic 1970-01-01 00:00:00
		Location Palo Alto 1970-01-01 00:00:00
		x -1 2018-05-14 19:27:06.788
		y -1 2018-05-14 19:27:06.788
		z -1 2018-05-14 19:27:06.788

11.8 Storage Plugins

Storage plugins are used to interact with the Storage Microservice and provide the persistent storage of information for Fledge.

The current version of Fledge 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.

11.8.1 Data and Metadata

Persistency is split in two blocks:

- **Metadata persistency**: it refers to the storage of metadata for Fledge, 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 Fledge, 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 Fledge to change the storage configuration.

The storage configuration file is stored in the Fledge 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 Fledge 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."
  },
  "threads": {
    "value": "1",
    "description": "The number of threads to run"
  },
  "managedStatus": {
    "value": "false",
    "description": "Control if Fledge should manage the storage provider"
  },
  "port": {
    "value": "0",
    "description": "The port to listen on"
  },
  "managementPort": {
```

(continues on next page)

(continued from previous page)

```
"value": "0",  
"description": "The management port to listen on."  
}  
}
```

Fledge 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 Fledge to manage the underlying storage system. If, for example you used a database server and you wished Fledge to start and stop that server as part of the Fledge 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.

11.8.2 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 Fledge 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:
 - Fledge 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 Fledge. 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 Fledge. Fledge provides a set of generic API to execute backup and restore operations.

11.9 Filter Plugins

Filter plugins provide a mechanism to alter the data stream as it flows through a fledge 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.

11.9.1 Configuration

Filters use the same configuration mechanism as the rest of Fledge, 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.

11.9.2 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 Fledge 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 Fledge 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)
{
}
```

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 do 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 *FledgeFilter* 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 FledgeFilter {
public:
    FledgeFilter(const std::string& filterName,
                ConfigCategory& filterConfig,
                OUTPUT_HANDLE *outHandle,
                OUTPUT_STREAM output);
    ~FledgeFilter() {};
    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);
public:
    OUTPUT_HANDLE*
        m_data;
    OUTPUT_STREAM
        m_func;
protected:
    std::string
        m_name;
    ConfigCategory
        m_config;
    bool
        m_enabled;
};
```

11.9.3 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.

```
/*
 * Fledge "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",
        "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
};
```

(continues on next page)

(continued from previous page)

```

    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);

    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)

```

(continues on next page)

(continued from previous page)

```

{
    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 `FledgeFilter` class

```

#ifndef _LOG_FILTER_H
#define _LOG_FILTER_H
/*
 * Fledge "Log" filter plugin.
 *
 * Copyright (c) 2020 Dianomic Systems
 *
 * Released under the Apache 2.0 Licence
 *
 * Author: Mark Riddoch
 */
#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 FledgeFilter {
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);

```

(continues on next page)

(continued from previous page)

```

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

```

/*
 * Fledge "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.
 *
 * @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),
                    FledgeFilter(filterName, filterConfig, outHandle, ↪
 * ↪ output)
{
    handleConfig(filterConfig);
}

/**
 * Destructor for this filter class
 */
LogFilter::~LogFilter()
{
    if (m_regex)

```

(continues on next page)

(continued from previous page)

```

        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"));

                // 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_
↳function

                    */
                    if (value.getType() == DatapointValue::T_INTEGER)
                    {

```

(continues on next page)

(continued from previous page)

```

        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
 * 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

```

(continues on next page)

(continued from previous page)

```
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);
    }
}
```

11.9.4 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:
        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
    """
```

11.9.5 Python Filter Example

The following is an example of a Python filter that calculates an exponential moving average.

```
# -*- coding: utf-8 -*-

# Fledge_BEGIN
# See: http://fledge-iot.readthedocs.io/
# Fledge_END
```

(continues on next page)

(continued from previous page)

```

""" 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 fledge.common import logger
import filter_ingest

__author__ = "Massimiliano Pinto"
__copyright__ = "Copyright (c) 2020 Dianomic Systems"
__license__ = "Apache 2.0"
__version__ = "${VERSION}"

_LOGGER = logger.setup(__name__, level = logging.WARN)

# Filter specific objects
the_callback = None
the_ingest_ref = None

# latest ema value
latest = None
# rate value
rate = None
# datapoint name
datapoint = None
# plugin shutdown indicator
shutdown_in_progress = False

_DEFAULT_CONFIG = {
    'plugin': {
        'description': 'Exponential Moving Average filter plugin',
        'type': 'string',
        'default': 'ema',
        'readonly': 'true'
    },
    'enable': {
        'description': 'Enable ema plugin',
        '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"
    },
}

```

(continues on next page)

(continued from previous page)

```

    'datapoint': {
        'description': 'Datapoint name for calculated ema value',
        'type': 'string',
        'default': 'ema',
        'displayName': 'EMA datapoint',
        'order': "1"
    }
}

def compute_ema(reading):
    """ Compute EMA

    Args:
        A reading data
    """
    global rate, latest, datapoint
    for attribute in list(reading):
        if not latest:
            latest = reading[attribute]
            latest = reading[attribute] * rate + latest * (1 - rate)
            reading[datapoint] = latest

def plugin_info():
    """ Returns information about the plugin

    Args:
    Returns:
        dict: plugin information
    Raises:
    """
    return {
        'name': 'ema',
        'version': '1.8.2',
        '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
        ingest_ref:
        callback:

    Returns:
        data: JSON object to be used in future calls to the plugin
    Raises:
    """
    data = copy.deepcopy(config)

    global the_callback, the_ingest_ref, rate, datapoint

    the_callback = callback

```

(continues on next page)

(continued from previous page)

```

the_ingest_ref = ingest_ref
rate = float(config['rate']['value'])
datapoint = config['datapoint']['value']

_LOGGER.debug("plugin_init for filter EMA called")

return data

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
    """
    global rate, datapoint
    rate = float(new_config['rate']['value'])
    datapoint = new_config['datapoint']['value']
    _LOGGER.debug("Old config for ema plugin {} \n new config {}".format(handle, new_
category))
    new_handle = copy.deepcopy(new_config)

    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
    """
    global shutdown_in_progress, the_callback, the_ingest_ref, rate, latest, datapoint
    shutdown_in_progress = True
    time.sleep(1)
    the_callback = None
    the_ingest_ref = None
    rate = None
    latest = None
    datapoint = None

    _LOGGER.info('filter ema plugin shutdown.')

def plugin_ingest(handle, data):
    """ Modify readings data and pass it onward

    Args:
        handle: handle returned by the plugin initialisation call
        data: readings data
    """
    global shutdown_in_progress, the_callback, the_ingest_ref

```

(continues on next page)

(continued from previous page)

```

if shutdown_in_progress:
    return

if handle['enable']['value'] == 'false':
    # Filter not enabled, just pass data onwards
    filter_ingest.filter_ingest_callback(the_callback, the_ingest_ref, data)
    return

# Filter is enabled: compute EMA for each reading
for elem in data:
    compute_ema(elem['readings'])

# Pass data onwards
filter_ingest.filter_ingest_callback(the_callback, the_ingest_ref, data)

_LOGGER.debug("ema filter_ingest done")

```

11.10 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.

11.10.1 Configuration

Notification Delivery plugins use the same configuration mechanism as the rest of Fledge, 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.

11.10.2 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 Fledge 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 Fledge 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;
    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;
    }
}
```

(continues on next page)

(continued from previous page)

```

    // 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(PUGIN_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.

11.11 Testing Your Plugin

The first step in testing your new plugin is to put the plugin in the location in which your Fledge system will be loading it from. The exact location depends on the way you installed your Fledge system and the type of plugin.

If your Fledge system was installed from a package and you used the default installation path, then your plugin must be stored under the directory */usr/local/fledge*. If you installed Fledge in a nonstandard location or you have built it from the source code, then the plugin should be stored under the directory *\$FLEDGE_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/fledge/plugins/south/DHT11*. Within that directory Fledge 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 *\$(FLEDGE_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/fledge/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 call normalise, on a system installed from a package in the default location should be copied into a directory */usr/local/fledge/python/fledge/plugins/filter/normalise*. Within this directory should be a file called *normalise.py* and an empty file called *__init__.py*.

11.11.1 Initial Testing

After you have copied your plugin into the correct location you can test if Fledge is able to see it by running the API call */fledge/plugins/installed*. This will list all the installed plugins and their versions.

```
$ curl http://localhost:8081/fledge/plugins/installed | jq
{
  "plugins": [
```

(continues on next page)

(continued from previous page)

```

{
  "name": "pi_server",
  "type": "north",
  "description": "PI Server North Plugin",
  "version": "1.0.0",
  "installedDirectory": "north/pi_server",
  "packageName": ""
},
{
  "name": "ocs",
  "type": "north",
  "description": "OCS (OSIsoft Cloud Services) North Plugin",
  "version": "1.0.0",
  "installedDirectory": "north/ocs",
  "packageName": ""
},
{
  "name": "http_north",
  "type": "north",
  "description": "HTTP North Plugin",
  "version": "1.8.1",
  "installedDirectory": "north/http_north",
  "packageName": "fledge-north-http-north"
},
{
  "name": "GCP",
  "type": "north",
  "description": "Google Cloud Platform IoT-Core",
  "version": "1.8.1",
  "installedDirectory": "north/GCP",
  "packageName": "fledge-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.

11.11.2 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/utls/get_plugin_info*. If you have installed a package, or run *make install*, you can find it in */usr/local/fledge/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.

11.11.3 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 Fledge
- Use the fledge status script to find the arguments to pass the service

```
$ scripts/fledge status
Fledge v1.8.2 running.
Fledge Uptime: 1451 seconds.
Fledge records: 200889 read, 200740 sent, 120962 purged.
Fledge does not require authentication.
=== Fledge services:
fledge.services.core
fledge.services.storage --address=0.0.0.0 --port=39821
fledge.services.south --port=39821 --address=127.0.0.1 --name=AX8
fledge.services.south --port=39821 --address=127.0.0.1 --name=Sine
=== Fledge tasks:
```

- Note the `--port=` and `--address=` arguments
- Set your `LD_LIBRARY_PATH`. This is normally done in the script that launches Fledge but will need to be run as a manual step when running under the debugger.

```
export LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:/usr/local/fledge/lib
```

If you built from source rather than installing a package you will need to include the libraries you built

```
export LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:${FLEDGE_ROOT}/cmake_build/C/
↪lib
```

- Load the service you wish to use to run your plugin, e.g a south service, under the debugger

```
$ gdb services/fledge.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.

```
(gdb) run --port=39821 --address=127.0.0.1 --name=ServiceName -d
```

Where *ServiceName* is the name you gave your service

- You can now use the debugger in the way you normally would to find any issues.

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 Fledge 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 fledge.services.storage
fledge 23318 0.0 0.3 270848 12388 ? Ssl 10:00 0:01 /usr/
↪local/fledge/services/fledge.services.storage --address=0.0.0.0 --
↪port=33761
fledge 31033 0.0 0.0 13136 1084 pts/1 S+ 10:37 0:00 grep --
↪color=auto fledge.services.storage
```

- Use the process ID of the fledge service as an argument to `gdb`. Note you will need to run `gdb` as root on some systems

```
$ sudo gdb /usr/local/fledge/services/fledge.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>
```

(continues on next page)

(continued from previous page)

```

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/fledge.services.storage...done.
Attaching to program: /usr/local/fledge/services/fledge.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".
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 Fledge in order to disable all processes and services. This will allow you to disable services or to run a particular service manually.

11.11.4 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 Fledge
- Use the fledge status script to find the arguments to pass the service

```

$ scripts/fledge status
Fledge v1.8.2 running.
Fledge Uptime: 1451 seconds.
Fledge records: 200889 read, 200740 sent, 120962 purged.
Fledge does not require authentication.
=== Fledge services:
fledge.services.core
fledge.services.storage --address=0.0.0.0 --port=39821
fledge.services.south --port=39821 --address=127.0.0.1 --name=AX8
fledge.services.south --port=39821 --address=127.0.0.1 --name=Sine
=== Fledge 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/fledge.services.south --port=39821 --address=127.0.0.1
↪--name=ServiceName -d
```

Where *ServiceName* is the name you gave your service

11.11.5 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 Fledge
- Use the fledge status script to find the arguments to pass the service

```
$ scripts/fledge status
Fledge v1.8.2 running.
Fledge Uptime: 1451 seconds.
Fledge records: 200889 read, 200740 sent, 120962 purged.
Fledge does not require authentication.
=== Fledge services:
fledge.services.core
fledge.services.storage --address=0.0.0.0 --port=39821
fledge.services.south --port=39821 --address=127.0.0.1 --name=AX8
fledge.services.south --port=39821 --address=127.0.0.1 --name=Sine
=== Fledge 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

```
$ valgrind --leak-check=full services/fledge.services.south --
↪port=39821 --address=127.0.0.1 --name=ServiceName -d
```

Where *ServiceName* is the name you gave your service

11.11.6 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/fledge` or `${FLEDGE_ROOT}` directory run the command

```
python3 -c 'from fledge.plugins.south.<name>.<name> import plugin_info;
↪print(plugin_info())'
```

Where `<name>` is the name of your plugin.


```
python3 -c 'from fledge.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': 'Asset name',
↳ 'mandatory': 'true'}}
```

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 fledge.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'}}
```


12.1 The Fledge REST API

Users, administrators and applications interact with Fledge via a REST API. This section presents a full reference of the API.

Note: The Fledge REST API should not be confused with the internal REST API used by Fledge tasks and microservices to communicate with each other.

12.1.1 Introducing the Fledge REST API

The REST API is the route into the Fledge appliance, it provides all user and program interaction to configure, monitor and manage the Fledge 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 Fledge appliance
- Access to monitoring statistics for the Fledge appliance
- User and role management for access to the API
- Access to the data buffer contents

Port Usage

In general Fledge 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 *\$FLEDGE_DATA/var/run/*:

```
$ cat data/var/run/fledge.core.pid
{ "adminAPI" : { "protocol" : "HTTP",
                 "port"      : 8081,
                 "addresses" : [ "0.0.0.0" ] },
  "processID" : 3585 }
```

Fledge 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. Fledge should then be restarted to make use of this new port.

Infrastructure

There are two REST API's that allow external access to Fledge, the **Administration API** and the **User API**. The User API is intended to allow access to the data in the Fledge 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 Fledge appliance. This API is used for all configuration operations that occur beyond basic installation.

12.2 Administration API Reference

This section presents the list of administrative API methods in alphabetical order.

12.2.1 Audit Trail

The audit trail API is used to interact with the audit trail log tables in the storage microservice. In Fledge, log information is stored in the system log where the microservice is hosted. All the relevant information used for auditing are instead stored inside Fledge 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 /fledge/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/fledge/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/fledge/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 /fledge/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 Fledge API to utilize the Fledge 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/fledge/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/fledge/audit?severity=FAILURE
{ "totalCount": 1,
  "audit": [ { "timestamp": "2018-04-16 18:32:28.427",
               "source" : "LOGGN",
               "details" : { "message": "Internal System Error" },
               "severity" : "FAILURE" }
            ]
}
$
```

12.2.2 Configuration Management

Configuration management is an important aspect of the REST API, however due to the discoverable form of the configuration of Fledge 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 Fledge is achieved.

category

The *category* interface is part of the Configuration Management for Fledge and it is used to create, retrieve, update and delete configuration categories and items.

GET categor(ies)

GET /fledge/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/fledge/category
{
  "categories":
  [
    {
      "key": "SCHEDULER",
      "description": "Scheduler configuration",
      "displayName": "Scheduler"
    },
    {
      "key": "SMNTR",
      "description": "Service Monitor",
      "displayName": "Service Monitor"
    },
    {
      "key": "rest_api",
      "description": "Fledge Admin and User REST API",
      "displayName": "Admin API"
    },
    {
      "key": "service",
      "description": "Fledge Service",
      "displayName": "Fledge Service"
    },
    {
      "key": "Installation",
      "description": "Installation",
      "displayName": "Installation"
    },
    {
      "key": "General",
      "description": "General",

```

(continues on next page)

(continued from previous page)

```

    "displayName": "General"
  },
  {
    "key": "Advanced",
    "description": "Advanced",
    "displayName": "Advanced"
  },
  {
    "key": "Utilities",
    "description": "Utilities",
    "displayName": "Utilities"
  }
]
}
$

```

GET category

GET /fledge/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 /fledge/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/fledge/category/rest_api
{
  "enableHttp": {
    "description": "Enable HTTP (disable to use HTTPS)",
    "type": "boolean",
    "default": "true",
    "displayName": "Enable HTTP",
    "order": "1",

```

(continues on next page)

(continued from previous page)

```

        "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": "fledge",
        "displayName": "Certificate Name",
        "order": "4",
        "validity": "enableHttp==\"false\"",
        "value": "fledge"
    },
    "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"
        ],
        "default": "any",
        "displayName": "Authentication method",
        "order": "6",
        "value": "any"
    },
    "authCertificateName": {
        "description": "Auth Certificate name",
        "type": "string",
        "default": "ca",

```

(continues on next page)

(continued from previous page)

```

    "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 /fledge/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 /fledge/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 Fledge 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/fledge/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 /fledge/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 /fledge/category call.
- **item** - the 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 Fledge 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/fledge/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 /fledge/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 /fledge/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 Fledge 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/fledge/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 /fledge/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/fledge/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" } }
}
```

12.2.3 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 Fledge tasks.

GET task

GET /fledge/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/fledge/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/fledge/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/fledge/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 /fledge/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/fledge/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/fledge/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 /fledge/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/fledge/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 /fledge/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/fledge/task/ea334d3b-8a33-4a29-845c-8be50efd44a4/
→ cancel
{"id": "ea334d3b-8a33-4a29-845c-8be50efd44a4", "message": "Task cancelled successfully"}
→ $
```

12.2.4 Other Administrative API calls**ping**

The *ping* interface gives a basic confidence check that the Fledge 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 /fledge/ping - return liveness of Fledge

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 Fledge 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	Fledge
hostName	string	Name of host machine	fledge
ipAddresses	list	IPv4 and IPv6 address of host machine	["10.0.0.0","123:234:345:456:567:789:1011:1213"]
health	string	Health of Fledge services	"green"
safeMode	boolean	True if Fledge is started in safe mode (only core and storage services will be started)	2113.076449394226

Example

```
$ curl -s http://localhost:8081/fledge/ping
{
  "uptime": 276818,
  "dataRead": 0,
  "dataSent": 0,
  "dataPurged": 0,
  "authenticationOptional": true,
  "serviceName": "Fledge",
  "hostName": "fledge",
  "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 Fledge device.

GET statistics

GET /fledge/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 Fledge started.

Key	Description
BUFFERED	Readings currently in the Fledge 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 Fledge
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/fledge/statistics
[ {
  "key": "BUFFERED",
  "description": "Readings currently in the Fledge buffer",
  "value": 0
},
...
{
  "key": "UNSNPURGED",
  "description": "Readings that were purged from the buffer before being sent",
  "value": 0
},
... ]
$
```

GET statistics/history

GET /fledge/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 Fledge, 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 Fledge 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 Fledge
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 Fledge since startup with name <i>asset-code</i>

Example

```
$ curl -s http://localhost:8081/fledge/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
    }
  ]
}
```

12.3 User API Reference

The user API provides a mechanism to access the data that is buffered within Fledge. 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 `/fledge/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 `/fledge/authenticate`. This token must be checked for validity and also that the authenticated entity has user or admin permissions.

12.3.1 Browsing Assets

asset

The `asset` method is used to browse all or some assets, based on search and filtering.

GET all assets

`GET /fledge/asset` - Return an array of asset codes buffered in Fledge 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/fledge/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 /fledge/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/fledge/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" },
  { "reading": { "x": 0, "y": 2, "z": 0 }, "timestamp": "2018-04-19 14:20:37.418" },
  { "reading": { "x": -2, "y": -1, "z": 2 }, "timestamp": "2018-04-19 14:20:32.650" },
  { "reading": { "x": 0, "y": 0, "z": 1 }, "timestamp": "2018-04-19 14:06:05.870" },
  { "reading": { "x": 1, "y": 1, "z": 1 }, "timestamp": "2018-04-19 14:06:05.870" },
  { "reading": { "x": 0, "y": 0, "z": -1 }, "timestamp": "2018-04-19 14:06:05.869" },
  ... ]
```

(continues on next page)

(continued from previous page)

```

{ "reading": { "x": 2, "y": -1, "z": 0 }, "timestamp": "2018-04-19 14:06:05.868" },
{ "reading": { "x": -1, "y": -2, "z": 2 }, "timestamp": "2018-04-19 14:06:05.867" },
{ "reading": { "x": 2, "y": 1, "z": 1 }, "timestamp": "2018-04-19 14:06:05.867" },
{ "reading": { "x": 1, "y": -2, "z": 1 }, "timestamp": "2018-04-19 14:06:05.866" },
{ "reading": { "x": 2, "y": -1, "z": 1 }, "timestamp": "2018-04-19 14:06:05.865" },
{ "reading": { "x": 0, "y": -1, "z": 2 }, "timestamp": "2018-04-19 14:06:05.865" },
{ "reading": { "x": 0, "y": -2, "z": 1 }, "timestamp": "2018-04-19 14:06:05.864" },
{ "reading": { "x": -1, "y": -2, "z": 0 }, "timestamp": "2018-04-19 13:45:15.881" } ]
↪ ]
$
$ curl -s http://localhost:8081/fledge/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 /fledge/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 default 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/fledge/asset/fogbench%2Fhumidity/temperature
[ { "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" },
  { "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" },

```

(continues on next page)

(continued from previous page)

```

{ "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/fledge/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 /fledge/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/fledge/asset/fogbench%2Fhumidity/temperature/summary
{ "temperature": { "max": 47, "min": 0, "average": 27 } }
$

```

GET timed average asset reading

GET /fledge/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 default 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/fledge/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/fledge/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 } ]
```

13.1 Building Developers Guide

13.1.1 Introduction

What Is Fledge?

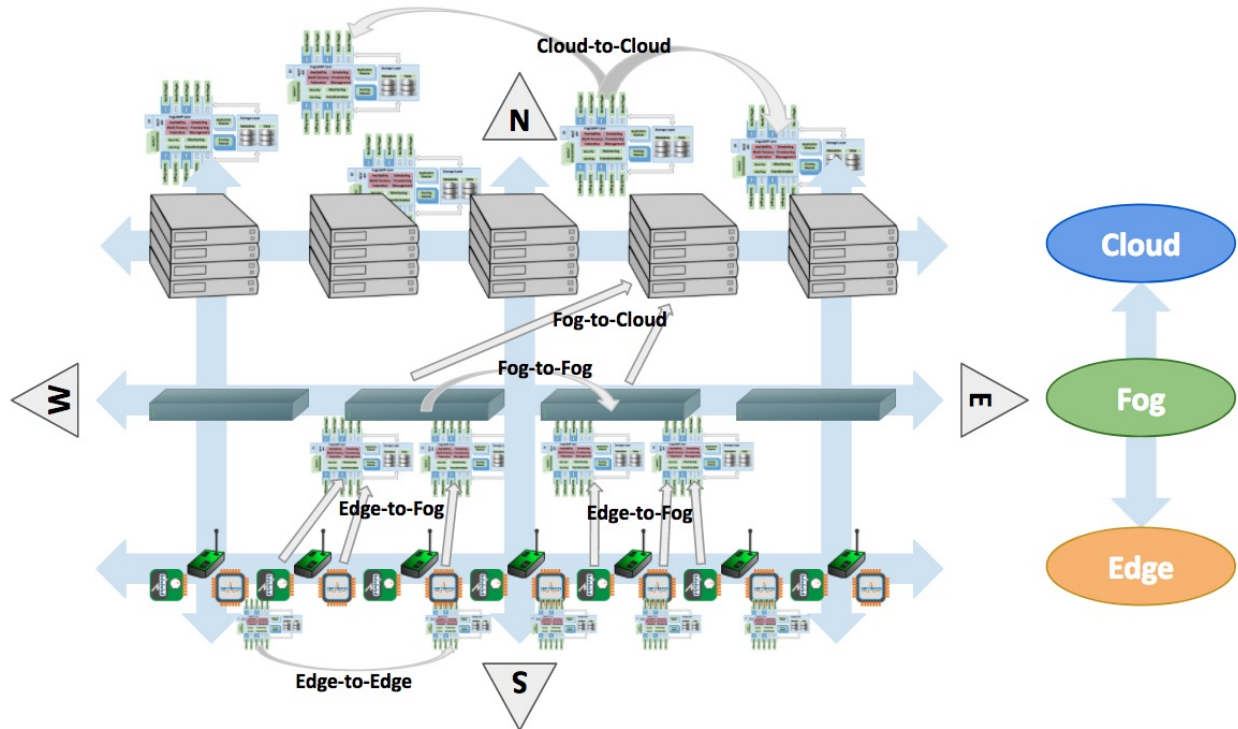
Fledge is an open source platform for the **Internet of Things** and an essential component in **Fog Computing**. It uses a modular **microservices architecture** including sensor data collection, storage, processing and forwarding to historians, Enterprise systems and Cloud-based services. Fledge can 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, Fledge 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.

Fledge 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.

Fledge Positioning in an IoT and IIoT Infrastructure

Fledge 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.
- Fledge provides high availability, scalability and data distribution among **Fog-to-Fog** systems. Due to its portability and modularity, Fledge can be installed on a large number of intermediate servers and gateways, as application instances, appliances, containers or virtualized environments.
- **Cloud to Cloud Fledge 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).

Fledge Features

In a nutshell, these are main features of Fledge:

- 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. Fledge 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.

13.1.2 Building Fledge

Let's get started! In this chapter we will see where to find and how to build, install and run Fledge for the first time.

Fledge Platforms

Due to the use of standard libraries, Fledge can run on a large number of platforms and operating environments, but its primary target is Linux distributions. Our testing environment includes Ubuntu LTS 18.04 and Raspbian, but we have installed and tested Fledge on other Linux distributions. In addition to the native support, Fledge can also run on Virtual Machines, Docker and LXC containers.

General Requirements

This version of Fledge requires the following software to be installed in the same environment:

- **Avahi 0.6.32+**
- **Python 3.6.9+**
- **PostgreSQL 9.5+**
- **SQLite 3.11+**

If you intend to download and build Fledge from source (as explained in this page), you also need *git*. In this version SQLite is default engine, but we have left libraries to easily switch to PostgreSQL, in case you need it. The PostgreSQL plugin will be moved to a different repository in future versions. Other requirements largely depend on the plugins that run in Fledge.

You may also want to install some utilities to make your life easier when you use or test Fledge:

- **curl**: it is used to interact with the REST API
- **jq**: the JSON processor, it helps in formatting the output of the REST API calls

Building Fledge

In this section we will describe how to build Fledge on Ubuntu 18.04 LTS (Server or Desktop). Other Linux distributions, Debian or Red-Hat based, or even other versions of Ubuntu may differ. If you are not familiar with Linux and you do not want to build Fledge from the source code, you can download a ready-made Debian package (the list of packages is [available here](#)).

Build Pre-Requisites

Fledge is currently based on C/C++ and Python code. The packages needed to build and run Fledge are:

- autoconf
- automake
- avahi-daemon
- build-essential
- ca-certificates
- cmake
- cpulimit

- curl
- g++
- git
- krb5-user
- libboost-dev
- libboost-system-dev
- libboost-thread-dev
- libcurl4-openssl-dev
- libssl-dev
- libpq-dev
- libsqlite3-dev
- libtool
- libz-dev
- make
- pkg-config
- postgresql
- python3-dev
- python3-pip
- python3-setuptools
- sqlite3
- uuid-dev

```
$ sudo apt-get update
Get:1 http://security.ubuntu.com/ubuntu xenial-security InRelease [102 kB]
...
All packages are up-to-date.
$
$ sudo apt-get install avahi-daemon ca-certificates curl git cmake g++ make build-
↳essential autoconf automake
Reading package lists... Done
Building dependency tree
...
$
$ sudo apt-get install sqlite3 libsqlite3-dev
Reading package lists... Done
Building dependency tree
...
$
$ sudo apt-get install libtool libboost-dev libboost-system-dev libboost-thread-dev
↳libssl-dev libpq-dev uuid-dev libz-dev
Reading package lists... Done
Building dependency tree
...
$
$ sudo apt-get install python3-dev python3-pip python3-setuptools
```

(continues on next page)

(continued from previous page)

```

Reading package lists... Done
Building dependency tree
...
$
$ sudo apt-get install postgresql
Reading package lists... Done
Building dependency tree
$
$
...
$
$ sudo apt-get install pkg-config cpulimit
Reading package lists... Done
Building dependency tree
$
$
...
$
$ sudo DEBIAN_FRONTEND=noninteractive apt-get install -yq krb5-user
Reading package lists... Done
Building dependency tree
$
$
...
$
$ sudo DEBIAN_FRONTEND=noninteractive apt-get install -yq libcurl4-openssl-dev
Reading package lists... Done
Building dependency tree
$

```

Obtaining the Source Code

Fledge is available on GitHub. The link to the repository is <https://github.com/fledge-iot/Fledge>. In order to clone the code in the repository, type:

```

$ git clone https://github.com/fledge-iot/Fledge.git
Cloning into 'Fledge'...
remote: Counting objects: 15639, done.
remote: Compressing objects: 100% (88/88), done.
remote: Total 15639 (delta 32), reused 58 (delta 14), pack-reused 15531
Receiving objects: 100% (15639/15639), 9.71 MiB | 2.11 MiB/s, done.
Resolving deltas: 100% (10486/10486), done.
Checking connectivity... done.
$

```

The code should be now in your home directory. The name of the repository directory is *Fledge*:

```

$ ls -l Fledge
total 128
drwxr-xr-x  7 ubuntu ubuntu    224 Jan  3 20:08 C
-rw-r--r--  1 ubuntu ubuntu   1480 May  7 00:29 CMakeLists.txt
-rw-r--r--  1 ubuntu ubuntu  11346 Jan  3 20:08 LICENSE
-rw-r--r--  1 ubuntu ubuntu  20660 Mar 13 00:25 Makefile
-rw-r--r--  1 ubuntu ubuntu   9173 May  7 00:29 README.rst
-rwxr-xr-x  1 ubuntu ubuntu    38 May  9 19:50 VERSION
drwxr-xr-x  3 ubuntu ubuntu    96 Jan  3 20:08 contrib
drwxr-xr-x  4 ubuntu ubuntu   128 Jan  3 20:08 data
drwxr-xr-x 15 ubuntu ubuntu   480 Jan  3 20:08 dco-signoffs
drwxr-xr-x 24 ubuntu ubuntu   768 May 11 00:44 docs

```

(continues on next page)

(continued from previous page)

```
drwxr-xr-x  3 ubuntu ubuntu    96 Jan  3 20:08 examples
drwxr-xr-x  4 ubuntu ubuntu   128 Jan  3 20:08 extras
drwxr-xr-x 14 ubuntu ubuntu   448 Jan  3 20:08 python
-rwxr-xr-x  1 ubuntu ubuntu  6804 Mar 13 00:25 requirements.sh
drwxr-xr-x 13 ubuntu ubuntu   416 May  7 00:29 scripts
drwxr-xr-x  7 ubuntu ubuntu   224 Mar 13 00:25 tests
drwxr-xr-x  3 ubuntu ubuntu    96 Jan  3 20:08 tests-manual
$
```

Selecting the Correct Version

The git repository created on your local machine, creates several branches. More specifically:

- The **main** branch is the latest, stable version. You should use this branch if you are interested in using Fledge with the last release features and fixes.
- The **develop** branch is the current working branch used by our developers. The branch contains the latest version and features, but it may be unstable and there may be issues in the code. You may consider to use this branch if you are curious to see one of the latest features we are working on, but you should not use this branch in production.
- The branches with versions **majorID.minorID**, such as *1.0* or *1.4*, contain the code of that specific version. You may use one of these branches if you need to check the code used in those versions.
- The branches with name **FOGL-XXXX**, where 'XXXX' is a sequence number, are working branches used by developers and contributors to add features, fix issues, modify and release code and documentation of Fledge. Those branches are free for you to see and learn from the work of the contributors.

Note that the default branch is *develop*.

Once you have cloned the Fledge project, in order to check the branches available, use the `git branch` command:

```
$ pwd
/home/ubuntu
$ cd Fledge
$ git branch --all
* develop
remotes/origin/1.0
...
remotes/origin/FOGL-822
remotes/origin/FOGL-823
remotes/origin/HEAD -> origin/develop
...
remotes/origin/develop
remotes/origin/main
$
```

Assuming you want to use the latest released, stable version, use the `git checkout` command to select the *master* branch:

```
$ git checkout main
Branch main set up to track remote branch main from origin.
Switched to a new branch 'main'
$
```

You can always use the `git status` command to check the branch you have checked out.

Building Fledge

You are now ready to build your first Fledge project. If you want to install Fledge on CentOS, Fedora or Red Hat, we recommend you to read this section first and then look at [this section](#). Move to the *Fledge* project directory, type the make command and let the magic happen.

```
$ cd Fledge
$ make
mkdir -p cmake_build
cd cmake_build ; cmake /home/ubuntu/Fledge/
-- The C compiler identification is GNU 5.4.0
-- The CXX compiler identification is GNU 5.4.0
...
pip3 install -Ir python/requirements.txt --user --no-cache-dir
...
Installing collected packages: multidict, idna, yarl, async-timeout, chardet, aiohttp,
↳ typing, aiohttp-cors, cchardet, pyjwt, six, pyjq
Successfully installed aiohttp-2.3.8 aiohttp-cors-0.5.3 async-timeout-3.0.0 cchardet-
↳ 2.1.1 chardet-3.0.4 idna-2.6 multidict-4.3.1 pyjq-2.1.0 pyjwt-1.6.0 six-1.11.0
↳ typing-3.6.4 yarl-1.2.6
$
```

Depending on the version of Ubuntu or other Linux distribution you are using, you may have found some issues. For example, there is a bug in the GCC compiler that raises a warning under specific circumstances. The output of the build will be something like:

```
/home/ubuntu/Fledge/C/services/storage/storage.cpp:97:14: warning: ignoring return_
↳ value of 'int dup(int)', declared with attribute warn_unused_result [-Wunused-
↳ result]
    (void)dup(0);      // stdout GCC bug 66425 produces warning
           ^
/home/ubuntu/Fledge/C/services/storage/storage.cpp:98:14: warning: ignoring return_
↳ value of 'int dup(int)', declared with attribute warn_unused_result [-Wunused-
↳ result]
    (void)dup(0);      // stderr GCC bug 66425 produces warning
           ^
```

The bug is documented . For our project, you should ignore it.

The other issue is related to the version of pip (more specifically pip3), the Python package manager. If you see this warning in the middle of the build output:

```
/usr/lib/python3.6/distutils/dist.py:261: UserWarning: Unknown distribution option:
↳ 'python_requires'
    warnings.warn(msg)
```

... and this output at the end of the build process:

```
You are using pip version 8.1.1, however version 9.0.1 is available.
You should consider upgrading via the 'pip install --upgrade pip' command.
```

In this case, what you need to do is to upgrade the pip software for Python3:

```
$ sudo pip3 install --upgrade pip
Collecting pip
  Downloading pip-9.0.1-py2.py3-none-any.whl (1.3MB)
    100% || 1.3MB 1.1MB/s
```

(continues on next page)

(continued from previous page)

```
$
```

Congratulations! You have installed and tested Fledge! If you want to go extra mile (and make the output of the REST API more readable, download the *jq* JSON processor and pipe the output of the *curl* command to it:

```
$ sudo apt install jq
...
$
$ curl -s http://localhost:8081/fledge/statistics | jq
[
  {
    "key": "BUFFERED",
    "description": "Readings currently in the Fledge buffer",
    "value": 0
  },
  {
    "key": "DISCARDED",
    "description": "Readings discarded by the South Service before being placed in_
↪the buffer. This may be due to an error in the readings themselves.",
    "value": 0
  },
  {
    "key": "PURGED",
    "description": "Readings removed from the buffer by the purge process",
    "value": 0
  },
  {
    "key": "READINGS",
    "description": "Readings received by Fledge",
    "value": 0
  },
  {
    "key": "UNSENT",
    "description": "Readings filtered out in the send process",
    "value": 0
  },
  {
    "key": "UNSNPURGED",
    "description": "Readings that were purged from the buffer before being sent",
    "value": 0
  }
]
$
```

Now I Want to Stop Fledge!

Easy, you have learnt `fledge start` and `fledge status`, simply type `fledge stop`:

```
$ scripts/fledge stop
Stopping Fledge.....
Fledge stopped.
$
```

As a next step, let's install Fledge!

Appendix: Setting the PostgreSQL Database

If you intend to use the PostgreSQL database as storage engine, make sure that PostgreSQL is installed and running correctly:

```
$ sudo systemctl status postgresql
postgresql.service - PostgreSQL RDBMS
   Loaded: loaded (/lib/systemd/system/postgresql.service; enabled; vendor preset: _
   → enabled)
   Active: active (exited) since Fri 2017-12-08 15:56:07 GMT; 15min ago
   Main PID: 14572 (code=exited, status=0/SUCCESS)
   CGroup: /system.slice/postgresql.service

Dec 08 15:56:07 ubuntu systemd[1]: Starting PostgreSQL RDBMS...
Dec 08 15:56:07 ubuntu systemd[1]: Started PostgreSQL RDBMS.
Dec 08 15:56:11 ubuntu systemd[1]: Started PostgreSQL RDBMS.
$
$ ps -ef | grep postgres
postgres 14806      1  0 15:56 ?           00:00:00 /usr/lib/postgresql/9.5/bin/postgres -
   → D /var/lib/postgresql/9.5/main -c config_file=/etc/postgresql/9.5/main/postgresql.
   → conf
postgres 14808 14806  0 15:56 ?           00:00:00 postgres: checkpoint process
postgres 14809 14806  0 15:56 ?           00:00:00 postgres: writer process
postgres 14810 14806  0 15:56 ?           00:00:00 postgres: wal writer process
postgres 14811 14806  0 15:56 ?           00:00:00 postgres: autovacuum launcher process
postgres 14812 14806  0 15:56 ?           00:00:00 postgres: stats collector process
ubuntu   15198  1225  0 17:22 pts/0    00:00:00 grep --color=auto postgres
$
```

PostgreSQL 9.5 is the version available for Ubuntu 18.04 when we have published this page. Other versions of PostgreSQL, such as 9.6 or 10.1, work just fine. When you install the Ubuntu package, PostgreSQL is set for a *peer authentication*, i.e. the database user must match with the Linux user. Other packages may differ. You may quickly check the authentication mode set in the `pg_hba.conf` file. The file is in the same directory of the `postgresql.conf` file you may see as output from the `ps` command shown above, in our case `/etc/postgresql/9.5/main`:

```
$ sudo grep '^local' /etc/postgresql/9.5/main/pg_hba.conf
local    all             postgres                  peer
local    all             all                      peer
$
```

The installation procedure also creates a Linux `postgres` user. In order to check if everything is set correctly, execute the `psql` utility as sudo user:

```
$ sudo -u postgres psql -l

                                List of databases
  Name      | Owner   | Encoding | Collate  | Ctype    | Access privileges
-----+-----+-----+-----+-----+-----
 postgres   | postgres | UTF8     | en_GB.UTF-8 | en_GB.UTF-8 |
 template0  | postgres | UTF8     | en_GB.UTF-8 | en_GB.UTF-8 | =c/postgres      +
            |         |         |             |             | postgres=CTc/postgres
 template1  | postgres | UTF8     | en_GB.UTF-8 | en_GB.UTF-8 | =c/postgres      +
            |         |         |             |             | postgres=CTc/postgres
(3 rows)
$
```

Encoding and collations may differ, depending on the choices made when you installed your operating system. Before you proceed, you must create a PostgreSQL user that matches your Linux user. Supposing that your user is `<fledge_user>`, type:

```
$ sudo -u postgres createuser -d <fledge_user>
```

The `-d` argument is important because the user will need to create the Fledge database.

A more generic command is: `$ sudo -u postgres createuser -d $(whoami)`

Finally, you should now be able to see the list of the available databases from your current user:

```
$ psql -l
```

List of databases					
Name	Owner	Encoding	Collate	Ctype	Access privileges
postgres	postgres	UTF8	en_GB.UTF-8	en_GB.UTF-8	
template0	postgres	UTF8	en_GB.UTF-8	en_GB.UTF-8	=c/postgres +
					postgres=Ctc/postgres
template1	postgres	UTF8	en_GB.UTF-8	en_GB.UTF-8	=c/postgres +
					postgres=Ctc/postgres

(3 rows)
\$

Appendix: Building Fledge on CentOS

In this section we present how to prepare a CentOS machine to build and install Fledge. A similar approach can be adopted to build the platform on RedHat and Fedora distributions. Here we refer to CentOS version 17.4.1708, requirements for other versions or distributions might differ.

Pre-Requisites

Pre-requisites on CentOS are similar to the ones on other distributions, but the name of the packages may differ from Debian-based distros. Starting from a minimal installation, this is the list of packages you need to add:

- libtool
- cmake
- boost-devel
- libuuid-devel
- gmp-devel
- mpfr-devel
- libmpc-devel
- sqlite3
- bzip2
- jq

This is the complete list of the commands to execute and the installed packages in CentoOS 17.4.1708.

```
sudo yum install libtool
sudo yum install cmake
sudo yum install boost-devel
sudo yum install libuuid-devel
sudo yum install gmp-devel
sudo yum install mpfr-devel
```

(continues on next page)

(continued from previous page)

```
sudo yum install libmpc-devel
sudo yum install bzip2
sudo yum install jq
sudo yum install libsqlite3x-devel
```

Building and Installing C++ 5.4

Fledge, requires C++ 5.4, CentOS 7 provides version 4.8. These are the commands to build and install the new GCC environment:

```
sudo yum install gcc-c++
curl https://ftp.gnu.org/gnu/gcc/gcc-5.4.0/gcc-5.4.0.tar.bz2 -O
bzip2 -dk gcc-5.4.0.tar.bz
tar xvf gcc-5.4.0.tar
mkdir gcc-5.4.0-build
cd gcc-5.4.0-build
../gcc-5.4.0/configure --enable-languages=c,c++ --disable-multilib
make -j$(nproc)
sudo make install
```

At the end of the procedure, the system will have two versions of GCC installed:

- GCC 4.8, installed in /usr/bin and /usr/lib64
- GCC 5.4, installed in /usr/local/bin and /usr/local/lib64

In order to use the latest version for Fledge, add the following lines at the end of your \$HOME/.bash_profile script:

```
export CC=/usr/local/bin/gcc
export CXX=/usr/local/bin/g++
export LD_LIBRARY_PATH=/usr/local/lib64
```

Installing PostgreSQL 9.6

CentOS provides PostgreSQL 9.2. Fledge has been tested with PostgreSQL 9.5, 9.6 and 10.X. Following <https://www.postgresql.org/download/> instructions, the commands to install the new version of PostgreSQL are:

```
sudo yum install -y https://download.postgresql.org/pub/repos/yum/reposrums/EL-7-x86_
64/pgdg-redhat-repo-latest.noarch.rpm
sudo yum install -y postgresql96-server
sudo yum install -y postgresql96-devel
sudo yum install -y rh-postgresql96
sudo yum install -y rh-postgresql96-postgresql-devel
sudo /usr/pgsql-9.6/bin/postgresql96-setup initdb
sudo systemctl enable postgresql-9.6
sudo systemctl start postgresql-9.6
```

At this point, Postgres has been configured to start at boot and it should be up and running. You can always check the status of the database server with `systemctl status postgresql-9.6`:

```
$ sudo systemctl status postgresql-9.6
[sudo] password for fledge:
```

(continues on next page)

(continued from previous page)

```

postgresql-9.6.service - PostgreSQL 9.6 database server
  Loaded: loaded (/usr/lib/systemd/system/postgresql-9.6.service; enabled; vendor_
  ↳ preset: disabled)
  Active: active (running) since Sat 2018-03-17 06:22:52 GMT; 8min ago
    Docs: https://www.postgresql.org/docs/9.6/static/
  Process: 1036 ExecStartPre=/usr/pgsql-9.6/bin/postgresql96-check-db-dir ${PGDATA}
  ↳ (code=exited, status=0/SUCCESS)
 Main PID: 1049 (postmaster)
   CGroup: /system.slice/postgresql-9.6.service
           └─1049 /usr/pgsql-9.6/bin/postmaster -D /var/lib/pgsql/9.6/data/
             └─1077 postgres: logger process
               └─1087 postgres: checkpoint process
                 └─1088 postgres: writer process
                   └─1089 postgres: wal writer process
                     └─1090 postgres: autovacuum launcher process
                       └─1091 postgres: stats collector process

Mar 17 06:22:52 vbox-centos-test systemd[1]: Starting PostgreSQL 9.6 database server..
  ↳ .
Mar 17 06:22:52 vbox-centos-test postmaster[1049]: < 2018-03-17 06:22:52.910 GMT >
  ↳ LOG:  redirecting log output to logging collector process
Mar 17 06:22:52 vbox-centos-test postmaster[1049]: < 2018-03-17 06:22:52.910 GMT >
  ↳ HINT:  Future log output will appear in directory "pg_log".
Mar 17 06:22:52 vbox-centos-test systemd[1]: Started PostgreSQL 9.6 database server.
$

```

Next, you must create a PostgreSQL user that matches your Linux user.

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

Finally, add `/usr/pgsql-9.6/bin` to your `PATH` environment variable in `$HOME/.bash_profile`. the new `PATH` setting in the file should look something like this:

```
PATH=$PATH:$HOME/.local/bin:$HOME/bin:/usr/pgsql-9.6/bin
```

Installing SQLite3

Fledge requires SQLite version 3.11 or later, CentOS provides an old version of SQLite. We must download SQLite, compile it and install it. The steps are:

- Download the source code of SQLite with `wget`. If you do not have `wget` installed, install it with `sudo yum install wget`: `wget http://www.sqlite.org/2018/sqlite-autoconf-3230100.tar.gz`
- Extract the SQLite tarball: `tar xzvf sqlite-autoconf-3230100.tar.gz`
- Move into the SQLite directory and execute the `configure-make-make install` commands: `cd sqlite-autoconf-3230100 ./configure make sudo make install`

Changing to the PostgreSQL Engine

The CentOS version of Fledge is optimized to work with PostgreSQL as storage engine. In order to achieve that, change the file `configuration.cpp` in the `C/services/storage` directory: line #20, word `sqlite` must be replaced with `postgres`:


```
" { \"plugin\" : { \"value\" : \"postgres\", \"description\" : \"The storage plugin to load\" }, \"
```

Building Fledge

We are finally ready to install Fledge, but we need to apply some little changes to the code and the make files. These changes will be removed in the future, but for the moment they are necessary to complete the procedure.

First, clone the Github repository with the usual command: `git clone https://github.com/fledge-iot/Fledge.git` The project should have been added to your machine under the *Fledge* directory.

We need to apply these changes to *C/plugins/storage/postgres/CMakeLists.txt*:

- Replace `include_directories(../../../../../thirdparty/rapidjson/include /usr/include/postgresql)` with: `include_directories(../../../../../thirdparty/rapidjson/include /usr/pgsql-9.6/include) link_directories(/usr/pgsql-9.6/lib)`

You are now ready to execute the make command, as described [here](#).

Further Notes

Here are some extra notes for the CentOS users.

Commented code The code commented in the previous paragraph is experimental and used for auto-discovery. It has been used for tests with South Microservices running on smart sensors, separated from the Core and Storage Microservices. This means that auto-discovery, i.e. the ability for a South Microservice to automatically identify the other services of Fledge distributed over the network, is currently not available on CentOS.

fledge start When Fledge starts on CentOS, it returns this message:

```
Starting Fledge v1.8.0.Fledge cannot start.
Check /home/fledge/Fledge/data/core.err for more information.
```

Check the *core.err* file, but if it is empty and *fledge status* shows Fledge running, it means that the services are up and running.

```
$ fledge start
Starting Fledge v1.8.0.Fledge cannot start.
Check /home/fledge/Fledge/data/core.err for more information.
$
$ fledge status
Fledge v1.8.0 running.
Fledge uptime: 6 seconds.
Fledge Records: 0 read, 0 sent, 0 purged.
Fledge does not require authentication.
=== Fledge services:
fledge.services.core
=== Fledge tasks:
$
$ cat data/core.err
$
$ ps -ef | grep fledge
...
fledge 6174 1 1 08:03 pts/0 00:00:00 python3 -m fledge.services.core
fledge 6179 1 0 08:03 ? 00:00:00 /home/fledge/Fledge/services/storage --
↪address=0.0.0.0 --port=34037
```

(continues on next page)

(continued from previous page)

```
fledge 6213 6212 0 08:04 pts/0 00:00:00 python3 -m fledge.tasks.statistics --
→port=34037 --address=127.0.0.1 --name=stats collector
...
$
```

fledge stop In CentOS, the command stops all the microservices with the exception of Core (with a `ps -ef` command you can easily check the process still running). You should execute a *stop* and a *kill* command to complete the shutdown on CentOS:

```
$ fledge status
Fledge v1.8.0 running.
Fledge uptime: 6 seconds.
Fledge Records: 0 read, 0 sent, 0 purged.
Fledge does not require authentication.
=== Fledge services:
fledge.services.core
=== Fledge tasks:
$ fledge stop
Stopping Fledge.....
Fledge stopped.
$
$ ps -ef | grep fledge
...
fledge 5782 1 5 07:56 pts/0 00:00:11 python3 -m fledge.services.core
...
$
$ fledge kill
Fledge killed.
$ ps -ef | grep fledge
...
$
```

13.1.3 Fledge Installation

Installing Fledge using defaults is straightforward: depending on the usage, you may install a new version from source or from a pre-built package. In environments where the defaults do not fit, you will need to execute few more steps. This chapter describes the default installation of Fledge and the most common scenarios where administrators need to modify the default behavior.

Installing Fledge from a Build

Once you have built Fledge following the instructions presented, you can execute the default installation with the `make install` command. By default, Fledge is installed from build in the root directory, under `/usr/local/fledge`. Since the root directory `/` is a protected system location, you will need superuser privileges to execute the command. Therefore, if you are not superuser, you should login as superuser or you should use the `sudo` command.

```
$ sudo make install
mkdir -p /usr/local/fledge
Installing Fledge version 1.8.0, DB schema 2
-- Fledge DB schema check OK: Info: /usr/local/fledge is empty right now. Skipping DB_
→schema check.
cp VERSION /usr/local/fledge
cd cmake_build ; cmake /home/fledge/Fledge/
```

(continues on next page)

(continued from previous page)

```
-- Boost version: 1.58.0
-- Found the following Boost libraries:
--   system
--   thread
--   chrono
--   date_time
--   atomic
-- Found SQLite version 3.11.0: /usr/lib/x86_64-linux-gnu/libsqlite3.so
-- Boost version: 1.58.0
-- Found the following Boost libraries:
--   system
--   thread
--   chrono
--   date_time
--   atomic
-- Configuring done
-- Generating done
-- Build files have been written to: /home/fledge/Fledge/cmake_build
cd cmake_build ; make
make[1]: Entering directory '/home/fledge/Fledge/cmake_build'
...
$
```

These are the main steps of the installation:

- Create the `/usr/local/fledge` directory, if it does not exist
- Build the code that has not been compiled and built yet
- Create all the necessary destination directories and copy the executables, scripts and configuration files
- Change the ownership of the `data` directory, if the install user is a superuser (we recommend to run Fledge as regular user, i.e. not as superuser).

Fledge is now present in `/usr/local/fledge` and ready to start. The start script is in the `/usr/local/fledge/bin` directory

```
$ cd /usr/local/fledge/
$ ls -l
total 32
drwxr-xr-x 2 root    root    4096 Apr 24 18:07 bin
drwxr-xr-x 4 fledge fledge 4096 Apr 24 18:07 data
drwxr-xr-x 4 root    root    4096 Apr 24 18:07 extras
drwxr-xr-x 4 root    root    4096 Apr 24 18:07 plugins
drwxr-xr-x 3 root    root    4096 Apr 24 18:07 python
drwxr-xr-x 6 root    root    4096 Apr 24 18:07 scripts
drwxr-xr-x 2 root    root    4096 Apr 24 18:07 services
-rwxr-xr-x 1 root    root      37 Apr 24 18:07 VERSION
$
$ bin/fledge
Usage: fledge {start|start --safe-mode|stop|status|reset|kill|help|version}
$
$ bin/fledge help
Usage: fledge {start|start --safe-mode|stop|status|reset|kill|help|version}
Fledge v1.3.1 admin script
The script is used to start Fledge
Arguments:
  start                - Start Fledge core (core will start other services).
  start --safe-mode    - Start in safe mode (only core and storage services will be
↳ started)
```

(continues on next page)

(continued from previous page)

```
stop          - Stop all Fledge services and processes
kill          - Kill all Fledge services and processes
status        - Show the status for the Fledge services
reset         - Restore Fledge factory settings
               WARNING! This command will destroy all your data!
version       - Print Fledge version
help          - This text
$
$ bin/fledge start
Starting Fledge.....
Fledge started.
$
```

Environment Variables

In order to operate, Fledge requires two environment variables:

- **FLEDGE_ROOT**: the root directory for Fledge. The default is `/usr/local/fledge`
- **FLEDGE_DATA**: the data directory. The default is `$FLEDGE_ROOT/data`, hence whichever value `FLEDGE_ROOT` has plus the `data` sub-directory, or `/usr/local/fledge/data` in case `FLEDGE_ROOT` is set as default value.

The `setenv.sh` Script

In the `extras/scripts` folder of the newly installed Fledge you can find the `setenv.sh` script. This script can be used to set the environment variables used by Fledge and update your `PATH` environment variable. You can call the script from your shell or you can add the same command to your `.profile` script:

```
$ cat /usr/local/fledge/extras/scripts/setenv.sh
#!/bin/sh

##-----
## Copyright (c) 2018 OSIsoft, LLC
##
## Licensed under the Apache License, Version 2.0 (the "License");
## you may not use this file except in compliance with the License.
## You may obtain a copy of the License at
##
##     http://www.apache.org/licenses/LICENSE-2.0
##
## Unless required by applicable law or agreed to in writing, software
## distributed under the License is distributed on an "AS IS" BASIS,
## WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
## See the License for the specific language governing permissions and
## limitations under the License.
##-----

#
# This script sets the user environment to facilitate the administration
# of Fledge
#
# You can execute this script from shell, using for example this command:
#
```

(continues on next page)

(continued from previous page)

```
# source /usr/local/fledge/extras/scripts/setenv.sh
#
# or you can add the same command at the bottom of your profile script
# {HOME}/.profile.
#

export FLEDGE_ROOT="/usr/local/fledge"
export FLEDGE_DATA="${FLEDGE_ROOT}/data"

export PATH="${FLEDGE_ROOT}/bin:${PATH}"
export LD_LIBRARY_PATH="${FLEDGE_ROOT}/lib:${LD_LIBRARY_PATH}"

$ source /usr/local/fledge/extras/scripts/setenv.sh
$
```

The fledge.service Script

Another file available in the *extras/scripts* folder is the *fledge.service* script. This script can be used to set Fledge as a Linux service. If you wish to do so, we recommend to install the Fledge package, but if you have a special build or for other reasons you prefer to work with Fledge built from source, this script will be quite helpful.

You can install Fledge as a service following these simple steps:

- After the `make install` command, copy *fledge.service* with a simple name *fledge* in the */etc/init.d* folder.
- Execute the command `systemctl enable fledge.service` to enable Fledge as a service
- Execute the command `systemctl start fledge.service` if you want to start Fledge

```
$ sudo cp /usr/local/fledge/extras/scripts/fledge.service /etc/init.d/fledge
$ sudo systemctl status fledge.service
fledge.service
  Loaded: not-found (Reason: No such file or directory)
  Active: inactive (dead)
$ sudo systemctl enable fledge.service
fledge.service is not a native service, redirecting to systemd-sysv-install
Executing /lib/systemd/systemd-sysv-install enable fledge
$ sudo systemctl status fledge.service
fledge.service - LSB: Fledge
  Loaded: loaded (/etc/init.d/fledge; bad; vendor preset: enabled)
  Active: inactive (dead)
  Docs: man:systemd-sysv-generator(8)
$ sudo systemctl start fledge.service
$ sudo systemctl status fledge.service
fledge.service - LSB: Fledge
  Loaded: loaded (/etc/init.d/fledge; generated)
  Active: active (running) since Thu 2020-05-28 18:42:07 IST; 9min ago
  Docs: man:systemd-sysv-generator(8)
  Process: 5047 ExecStart=/etc/init.d/fledge start (code=exited, status=0/SUCCESS)
  Tasks: 27 (limit: 4680)
  CGroup: /system.slice/fledge.service
          └─5123 python3 -m fledge.services.core
             └─5331 /usr/local/fledge/services/fledge.services.storage --address=0.0.0.0 -
→-port=34827
             └─8119 /bin/sh tasks/north_c --port=34827 --address=127.0.0.1 --name=OMF to
→PI north
```

(continues on next page)

(continued from previous page)

```
└─8120 ./tasks/sending_process --port=34827 --address=127.0.0.1 --name=OMF
→to PI north

...
$
```

Installing the Debian Package

We have versions of Fledge available as Debian packages for you. Check the to review which versions and platforms are available.

Obtaining and Installing the Debian Package

Check the to find the package to install.

Once you have downloaded the package, install it using the `apt-get` command. You can use `apt-get` to install a local Debian package and automatically retrieve all the necessary packages that are defined as pre-requisites for Fledge. Note that you may need to install the package as superuser (or by using the `sudo` command) and move the package to the `apt` cache directory first (`/var/cache/apt/archives`).

For example, if you are installing Fledge on an Intel x86_64 machine, you can type this command to download the package:

```
$ wget https://fledge-iot.s3.amazonaws.com/1.8.0/ubuntu1804/x86_64/fledge-1.8.0_x86_
→64_ubuntu1804.tgz
--2020-05-28 18:24:12-- https://fledge-iot.s3.amazonaws.com/1.8.0/ubuntu1804/x86_64/
→fledge-1.8.0_x86_64_ubuntu1804.tgz
Resolving fledge-iot.s3.amazonaws.com (fledge-iot.s3.amazonaws.com)... 52.217.40.188
Connecting to fledge-iot.s3.amazonaws.com (fledge-iot.s3.amazonaws.com)|52.217.40.
→188|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 24638625 (23M) [application/x-tar]
Saving to: 'fledge-1.8.0_x86_64_ubuntu1804.tgz'

fledge-1.8.0_x86_64_ubuntu1804.tg 100
→%[=====>] 23.50M 4.30MB/s
→ in 8.3s

2020-05-28 18:24:26 (2.84 MB/s) - 'fledge-1.8.0_x86_64_ubuntu1804.tgz' saved
→[24638625/24638625]
$
```

We recommend to execute an *update-upgrade-update* of the system first, then you may untar the `fledge-1.8.0_x86_64_ubuntu1804.tgz` file and copy the Fledge package in the *apt cache* directory and install it.

```
$ sudo apt update
Hit:1 http://gb.archive.ubuntu.com/ubuntu xenial InRelease
...
$ sudo apt upgrade
...
$ sudo apt update
...
$ sudo cp fledge-1.8.0-x86_64.deb /var/cache/apt/archives/.
...
```

(continues on next page)

(continued from previous page)

```

$ sudo apt install /var/cache/apt/archives/fledge-1.8.0-x86_64.deb
Reading package lists... Done
Building dependency tree
Reading state information... Done
Note, selecting 'fledge' instead of '/var/cache/apt/archives/fledge-1.8.0-x86_64.deb'
The following packages were automatically installed and are no longer required:
...
Unpacking fledge (1.8.0) ...
Setting up fledge (1.8.0) ...
Resolving data directory
Data directory does not exist. Using new data directory
Installing service script
Generating certificate files
Certificate files do not exist. Generating new certificate files.
Creating a self signed SSL certificate ...
Certificates created successfully, and placed in data/etc/certs
Generating auth certificate files
CA Certificate file does not exist. Generating new CA certificate file.
Creating ca SSL certificate ...
ca certificate created successfully, and placed in data/etc/certs
Admin Certificate file does not exist. Generating new admin certificate file.
Creating user SSL certificate ...
user certificate created successfully for admin, and placed in data/etc/certs
User Certificate file does not exist. Generating new user certificate file.
Creating user SSL certificate ...
user certificate created successfully for user, and placed in data/etc/certs
Setting ownership of Fledge files
Calling Fledge package update script
Linking update task
Changing setuid of update_task.apt
Removing task/update
Create link file
Copying sudoers file
Setting setuid bit of cmdutil
Enabling Fledge service
fledge.service is not a native service, redirecting to systemd-sysv-install.
Executing: /lib/systemd/systemd-sysv-install enable fledge
Starting Fledge service
$

```

As you can see from the output, the installation automatically registers Fledge as a service, so it will come up at startup and it is already up and running when you complete the command.

Check the newly installed package:

```

$ sudo dpkg -l | grep fledge
ii fledge      1.8.0      amd64      Fledge, the open source platform for the
↳ Internet of Things
$

```

You can also check the service currently running:

```

$ sudo systemctl status fledge.service
fledge.service - LSB: Fledge
Loaded: loaded (/etc/init.d/fledge; generated)
Active: active (running) since Thu 2020-05-28 18:42:07 IST; 9min ago
Docs: man:systemd-sysv-generator(8)

```

(continues on next page)

(continued from previous page)

```
Process: 5047 ExecStart=/etc/init.d/fledge start (code=exited, status=0/SUCCESS)
Tasks: 27 (limit: 4680)
CGroup: /system.slice/fledge.service
├─5123 python3 -m fledge.services.core
├─5331 /usr/local/fledge/services/fledge.services.storage --address=0.0.0.0 -
↳-port=34827
├─8119 /bin/sh tasks/north_c --port=34827 --address=127.0.0.1 --name=OMF to_
↳PI north
└─8120 ./tasks/sending_process --port=34827 --address=127.0.0.1 --name=OMF_
↳to PI north
...
$
```

Check if Fledge is up and running with the `fledge` command:

```
$ /usr/local/fledge/bin/fledge status
Fledge v1.8.0 running.
Fledge Uptime: 162 seconds.
Fledge records: 0 read, 0 sent, 0 purged.
Fledge does not require authentication.
=== Fledge services:
fledge.services.core
...
=== Fledge tasks:
...
$
```

Don't forget to add the `setenv.sh` available in the `/usr/local/fledge/extras/scripts*` directory to your `.profile` user startup script if you want to have an easy access to the Fledge tools, and...

...Congratulations! This is all you need to do, now Fledge is ready to run.

Upgrading or Downgrading Fledge

Upgrading or downgrading Fledge, starting from version 1.2, is as easy as installing it from scratch: simply follow the instructions in the previous section regarding the installation and the package will take care of the upgrade/downgrade path. The installation will not proceed if there is not a path to upgrade or downgrade from the currently installed version. You should still check the pre-requisites before you apply the upgrade. Clearly the old data will not be lost, there will be a schema upgrade/downgrade, if required.

Uninstalling the Debian Package

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

```
$ sudo apt purge fledge
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following packages were automatically installed and are no longer required:
  libmodbus-dev libmodbus5
Use 'sudo apt autoremove' to remove them.
The following packages will be REMOVED:
  fledge*
```

(continues on next page)

(continued from previous page)

```

0 upgraded, 0 newly installed, 1 to remove and 0 not upgraded.
After this operation, 0 B of additional disk space will be used.
Do you want to continue? [Y/n] y
(Reading database ... 160251 files and directories currently installed.)
Removing fledge (1.8.0) ...
Fledge is currently running.
Stop Fledge service.
Kill Fledge.
Disable Fledge service.
fledge.service is not a native service, redirecting to systemd-sysv-install.
Executing: /lib/systemd/systemd-sysv-install disable fledge
Remove Fledge service script
Reset systemctl
Cleanup of files
Remove fledge sudoers file
(Reading database ... 159822 files and directories currently installed.)
Purging configuration files for fledge (1.8.0) ...
Cleanup of files
Remove fledge sudoers file
dpkg: warning: while removing fledge, directory '/usr/local/fledge' not empty so not
↳ removed
$

```

The command also removes the service installed. You may notice the warning in the last row of the command output: this is due to the fact that the data directory (`/usr/local/fledge/data` by default) has not been removed, in case an administrator might want to analyze or reuse the data.

13.1.4 Testing Fledge

After the installation, you are now ready to test Fledge. An end-to-end test involves three types of tests:

- The **South** side, i.e. testing the collection of information from South microservices and associated plugins
- The **North** side, i.e. testing the tasks that send data North to historians, databases, Enterprise and Cloud systems
- The **East/West** side, i.e. testing the interaction of external applications with Fledge via REST API.

This chapter describes how to tests Fledge in these three directions.

First Checks: Fledge Status

Before we start, let's make sure that Fledge is up and running and that we have the tasks and services in place to execute the tests. First, run the `fledge status` command to check if Fledge has already started. The result of the command can be:

- `Fledge not running.` - it means that we must start Fledge with `fledge start`
- `Fledge starting.` - it means that we have started Fledge but the starting phase has not been completed yet. You should wait for a little while (from few seconds to about a minute) to see Fledge running.
- `Fledge running.` - (plus extra rows giving the uptime and other info. It means that Fledge is up and running, hence it is ready for use.

When you have a running Fledge, check the extra information provided by the `fledge status` command:

```
$ fledge status
Fledge v1.8.0 running.
Fledge Uptime: 9065 seconds.
Fledge records: 86299 read, 86851 sent, 0 purged.
Fledge does not require authentication.
=== Fledge services:
fledge.services.core
fledge.services.storage --address=0.0.0.0 --port=42583
fledge.services.south --port=42583 --address=127.0.0.1 --name=Sine
fledge.services.notification --port=42583 --address=127.0.0.1 --name=Fledge_
↪Notifications
=== Fledge tasks:
fledge.tasks.purge --port=42583 --address=127.0.0.1 --name=purge
tasks/sending_process --port=42583 --address=127.0.0.1 --name=PI Server
$
```

Let's analyze the output of the command:

- **Fledge running.** - The Fledge Core microservice is running on this machine and it is responding to the status command as *running* because other basic microservices are also running.
- **Fledge uptime: 282 seconds.** - This is a simple uptime in second provided by the Core microservice. It is equivalent to the ping method called via the REST API.
- **Fledge records:** - This is a summary of the number of records received from sensors and devices (South), sent to other services (North) and purged from the buffer.
- **Fledge authentication** - This row describes if a user or an application must authenticate to ogLAMP in order to operate with the REST API.

The following lines provide a list of the modules running in this installation of Fledge. They are separated by dots and described in this way:

- The prefix `fledge` is always present and identifies the Fledge modules.
- The following term describes the type of module: *services* for microservices, *tasks* for tasks etc.
- The following term is the name of the module: *core*, *storage*, *north*, *south*, *app*, *alert*
- The last term is the name of the plugin executed as part of the module.
- Extra arguments may be available: they are the arguments passed to the module by the core when it is launched.
- **=== Fledge services:** - This block contains the list of microservices running in the Fledge platform.
 - `fledge.services.core` is the Core microservice itself
 - `fledge.services.south --port=44180 --address=127.0.0.1 --name=COAP` - This South microservice is a listener of data pushed to Fledge via a CoAP protocol
- **=== Fledge tasks:** - This block contains the list of tasks running in the Fledge platform.
 - `fledge.tasks.north.sending_process ... --name=sending process` is a North task that prepares and sends data collected by the South modules to the OSIsoft PI System in OMF (OSIsoft Message Format).
 - `fledge.tasks.north.sending_process ... --name=statistics to pi` is a North task that prepares and sends the internal statistics to the OSIsoft PI System in OMF (OSIsoft Message Format).

Hello, Foggy World!

The output of the `fledge status` command gives you an idea of the modules running in your machine, but let's try to get more information from Fledge.

The Fledge REST API

First of all, we need to familiarize with the Fledge REST API. The API provides a set of methods used to monitor and administer the status of Fledge. Users and developers can also use the API to interact with external applications.

This is a short list of the methods available to the administrators. A more detailed list will be available soon: - **ping** provides the uptime of the Fledge Core microservice - **statistics** provides a set of statistics of the Fledge platform, such as data collected, sent, purged, rejected etc. - **asset** provides a list of asset that have readings buffered in Fledge. - **category** provides a list of the configuration of the modules and components in Fledge.

Useful Tools

Systems Administrators and Developers may already have their favorite tools to interact with a REST API, and they can probably use the same tools with Fledge. If you are not familiar with any tool, we recommend one of these:

- If you are familiar with the Linux shell and command lines, is the simplest and most useful tool available. It comes with every Linux distribution (or you can easily add it if it is not available in the default installation).
- If you prefer to use a browser-like interface, we recommend . Postman is an application available on Linux, MacOS and Windows and allows you to save queries, results, and run a set of queries with a single click.

Hello World!

Let's execute the *ping* method. First, you must identify the IP address where Fledge is running. If you have installed Fledge on your local machine, you can use *localhost*. Alternatively, check the IP address of the machine where Fledge is installed.

Note: This version of Fledge does not have any security setup by default, therefore you may be able to access the entry point for the REST API by any external application, but there may be security setting on your operating environment that prevent access to specific ports from external applications. If you receive an error using the ping method, and the `fledge status` command says that everything is running, it is likely that you are experiencing a security issue.

The default port for the REST API is 8081. Using curl, try this command:

```
$ curl -s http://localhost:8081/fledge/ping ; echo
{"uptime": 10480, "dataRead": 0, "dataSent": 0, "dataPurged": 0,
↪ "authenticationOptional": true, "serviceName": "Fledge", "hostName": "fledge",
↪ "ipAddresses": ["x.x.x.x", "x:x:x:x:x:x:x:x"], "health": "green", "safeMode": false}
$
```

The `echo` at the end of the line is simply used to add an extra new line to the output.

If you are using Postman, select the *GET* method and type `http://localhost:8081/fledge/ping` in the URI address input. If you are accessing a remote machine, replace *localhost* with the correct IP address. The output should be something like:



05_postman_ping.jpg

This is the first message you may receive from Fledge!

Hello from the Southern Hemisphere of the Fledge World

Let's now try something more exciting. The primary job of Fledge is to collect data from the Edge (we call it *South*), buffer it in our storage engine and then we send the data to Cloud historians and Enterprise Servers (we call them *North*). We also offer information to local or networked applications, something we call *East* or *West*.

In order to insert data you may need a sensor or a device that generates data. If you want to try Fledge but you do not have any sensor at hand, do not worry, we have a tool that can generate data as if it is a sensor.

fogbench: a Brief Intro

Fledge comes with a little but pretty handy tool called **fogbench**. The tool is written in Python and it uses the same libraries of other modules of Fledge, therefore no extra libraries are needed. With *fogbench* you can do many things, like inserting data stored in files, running benchmarks to understand how Fledge performs in a given environment, or test an end-to-end installation.

Note: The following instructions assume you have downloaded and installed the CoAP south plugin from <https://github.com/fledge-iot/fledge-south-coap>.

```
$ git clone https://github.com/fledge-iot/fledge-south-coap
$ cd fledge-south-coap
$ sudo cp -r python/fledge/plugins/south/coap /usr/local/fledge/python/fledge/plugins/
  ↳ south/
$ sudo cp python/requirements-coap.txt /usr/local/fledge/python/
$ sudo pip3 install -r /usr/local/fledge/python/requirements-coap.txt
$ sudo chown -R root:root /usr/local/fledge/python/fledge/plugins/south/coap
$ curl -sX POST http://localhost:8081/fledge/service -d '{"name": "CoAP", "type":
  ↳ "south", "plugin": "coap", "enabled": true}'
```

Depending on your environment, you can call *fogbench* in one of those ways:

- In a development environment, use the script *scripts/extras/fogbench*, inside your project repository (remember to set the *FLEDGE_ROOT* environment variable with the path to your project repository folder).
- In an environment deployed with `sudo make install`, use the script *bin/fogbench*.

You may call the *fogbench* tool like this:

```
$ /usr/local/fledge/bin/fogbench
>>> Make sure south CoAP plugin service is running & listening on specified host and
  ↳ port
usage: fogbench [-h] [-v] [-k {y,yes,n,no}] -t TEMPLATE [-o OUTPUT]
               [-I ITERATIONS] [-O OCCURRENCES] [-H HOST] [-P PORT]
               [-i INTERVAL] [-S {total}]
fogbench: error: the following arguments are required: -t/--template
$
```

...or more specifically, when you call invoke *fogbench* with the *-help* or *-h* argument:

```

$ /usr/local/fledge/bin/fogbench -h
>>> Make sure south CoAP plugin service is running & listening on specified host and
port
usage: fogbench [-h] [-v] [-k {y,yes,n,no}] -t TEMPLATE [-o OUTPUT]
               [-I ITERATIONS] [-O OCCURRENCES] [-H HOST] [-P PORT]
               [-i INTERVAL] [-S {total}]

fogbench -- a Python script used to test Fledge (simulate payloads)

optional arguments:
  -h, --help                show this help message and exit
  -v, --version              show program's version number and exit
  -k {y,yes,n,no}, --keep {y,yes,n,no}
                           Do not delete the running sample (default: no)
  -t TEMPLATE, --template TEMPLATE
                           Set the template file, json extension
  -o OUTPUT, --output OUTPUT
                           Set the statistics output file
  -I ITERATIONS, --iterations ITERATIONS
                           The number of iterations of the test (default: 1)
  -O OCCURRENCES, --occurrences OCCURRENCES
                           The number of occurrences of the template (default: 1)
  -H HOST, --host HOST      CoAP server host address (default: localhost)
  -P PORT, --port PORT      The Fledge port. (default: 5683)
  -i INTERVAL, --interval INTERVAL
                           The interval in seconds for each iteration (default:
                           0)
  -S {total}, --statistics {total}
                           The type of statistics to collect (default: total)

The initial version of fogbench is meant to test the sensor/device interface
of Fledge using CoAP
$

```

In order to use *fogbench* you need a template file. The template is a set of JSON elements that are used to create a random set of values that simulate the data generated by one or more sensors. Fledge comes with a template file named *fogbench_sensor_coap.template.json*. The template is located here:

- In a development environment, look in *data/extras/fogbench* in the project repository folder.
- In an environment deployed using `sudo make install`, look in *\$FLEDGE_DATA/extras/fogbench*.

The template file looks like this:

```

$ cat /usr/local/fledge/data/extras/fogbench/fogbench_sensor_coap.template.json
[
  { "name"           : "fogbench_luxometer",
    "sensor_values" : [ { "name": "lux", "type": "number", "min": 0, "max": 130000,
    ↪"precision":3 } ] },
  { "name"           : "fogbench_pressure",
    "sensor_values" : [ { "name": "pressure", "type": "number", "min": 800.0, "max": ↪
    ↪1100.0, "precision":1 } ] },
  { "name"           : "fogbench_humidity",
    "sensor_values" : [ { "name": "humidity", "type": "number", "min": 0.0, "max": ↪
    ↪100.0 },
                           { "name": "temperature", "type": "number", "min": 0.0, "max": ↪
    ↪50.0 } ] },
  { "name"           : "fogbench_temperature",

```

(continues on next page)

(continued from previous page)

```

    "sensor_values" : [ { "name": "object", "type": "number", "min": 0.0, "max": 50.0 },
    ↪ },
                        { "name": "ambient", "type": "number", "min": 0.0, "max": 50.
    ↪ 0 } ] },
    { "name"           : "fogbench_accelerometer",
      "sensor_values" : [ { "name": "x", "type": "number", "min": -2.0, "max": 2.0 },
                        { "name": "y", "type": "number", "min": -2.0, "max": 2.0 },
                        { "name": "z", "type": "number", "min": -2.0, "max": 2.0 } ] }
    ↪ ,
    { "name"           : "fogbench_gyroscope",
      "sensor_values" : [ { "name": "x", "type": "number", "min": -255.0, "max": 255.0 }
    ↪ ,
                        { "name": "y", "type": "number", "min": -255.0, "max": 255.0 }
    ↪ ,
                        { "name": "z", "type": "number", "min": -255.0, "max": 255.0 }
    ↪ ] },
    { "name"           : "fogbench_magnetometer",
      "sensor_values" : [ { "name": "x", "type": "number", "min": -255.0, "max": 255.0 }
    ↪ ,
                        { "name": "y", "type": "number", "min": -255.0, "max": 255.0 }
    ↪ ,
                        { "name": "z", "type": "number", "min": -255.0, "max": 255.0 }
    ↪ ] },
    { "name"           : "fogbench_mouse",
      "sensor_values" : [ { "name": "button", "type": "enum", "list": [ "up", "down" ] }
    ↪ ] },
    { "name"           : "fogbench_switch",
      "sensor_values" : [ { "name": "button", "type": "enum", "list": [ "up", "down" ] }
    ↪ ] },
    { "name"           : "fogbench_wall clock",
      "sensor_values" : [ { "name": "tick", "type": "enum", "list": [ "tock" ] } ] }
  ]
$

```

In the array, each element simulates a message from a sensor, with a name, a set of data points that have their name, value type and range.

Data Coming from South

Now you should have all the information necessary to test the CoAP South microservice. From the command line, type:

- `$FLEDGE_ROOT/scripts/extras/fogbench -t $FLEDGE_ROOT/data/extras/fogbench/fogbench_sensor_coap.template.json`, if you are in a development environment, with the *FLEDGE_ROOT* environment variable set with the path to your project repository folder
- `$FLEDGE_ROOT/bin/fogbench -t $FLEDGE_DATA/extras/fogbench/fogbench_sensor_coap.template.json`, if you are in a deployed environment, with *FLEDGE_ROOT* and *FLEDGE_DATA* set correctly. - If you have installed Fledge in the default location (i.e. */usr/local/fledge*), type `cd /usr/local/fledge; bin/fogbench -t data/extras/fogbench/fogbench_sensor_coap.template.json`.
- `fledge.fogbench -t /snap/fledge/current/usr/local/fledge/data/extras/fogbench/fogbench_sensor_coap.template.json`, if you have installed a snap version of Fledge.

In development environment the output of your command should be:

```
$ $FLEDGE_ROOT/scripts/extras/fogbench -t data/extras/fogbench/fogbench_sensor_coap.
↳template.json
>>> Make sure south CoAP plugin service is running & listening on specified host and
↳port
Total Statistics:

Start Time: 2017-12-17 07:17:50.615433
End Time:   2017-12-17 07:17:50.650620

Total Messages Transferred: 10
Total Bytes Transferred:    2880

Total Iterations: 1
Total Messages per Iteration: 10.0
Total Bytes per Iteration:    2880.0

Min messages/second: 284.19586779208225
Max messages/second: 284.19586779208225
Avg messages/second: 284.19586779208225

Min Bytes/second: 81848.4099241197
Max Bytes/second: 81848.4099241197
Avg Bytes/second: 81848.4099241197
$
```

Congratulations! You have just inserted data into Fledge from the CoAP South microservice. More specifically, the output informs you that the data inserted has been composed by 10 different messages for a total of 2,880 Bytes, for an average of 284 messages per second and 81,848 Bytes per second.

If you want to stress Fledge a bit, you may insert the same data sample several times, by using the *-I* or *-iterations* argument:

```
$ $FLEDGE_ROOT/scripts/extras/fogbench -t data/extras/fogbench/fogbench_sensor_coap.
↳template.json -I 100
>>> Make sure south CoAP plugin service is running & listening on specified host and
↳port
Total Statistics:

Start Time: 2017-12-17 07:33:40.568130
End Time:   2017-12-17 07:33:43.205626

Total Messages Transferred: 1000
Total Bytes Transferred:    288000

Total Iterations: 100
Total Messages per Iteration: 10.0
Total Bytes per Iteration:    2880.0

Min messages/second: 98.3032852957946
Max messages/second: 625.860558267618
Avg messages/second: 455.15247432732866

Min Bytes/second: 28311.346165188843
Max Bytes/second: 180247.840781074
Avg Bytes/second: 131083.9126062706
$
```

Here we have inserted the same set of data 100 times, therefore the total number of Bytes inserted is 288,000. The performance and insertion rates varies with each iteration and *fogbench* presents the minimum, maximum and average values.

Checking What's Inside Fledge

We can check if Fledge has now stored what we have inserted from the South microservice by using the *asset* API. From curl or Postman, use this URL:

```
$ curl -s http://localhost:8081/fledge/asset ; echo
[{"assetCode": "fogbench_switch", "count": 11}, {"assetCode": "fogbench_temperature",
↪ "count": 11}, {"assetCode": "fogbench_humidity", "count": 11}, {"assetCode":
↪ "fogbench_luxometer", "count": 11}, {"assetCode": "fogbench_accelerometer", "count
↪ ": 11}, {"assetCode": "wall clock", "count": 11}, {"assetCode": "fogbench_
↪ magnetometer", "count": 11}, {"assetCode": "mouse", "count": 11}, {"assetCode":
↪ "fogbench_pressure", "count": 11}, {"assetCode": "fogbench_gyroscope", "count": 11}]
$
```

The output of the asset entry point provides a list of assets buffered in Fledge and the count of elements stored. The output is a JSON array with two elements:

- **assetCode** : the name of the sensor or device that provides the data
- **count** : the number of occurrences of the asset in the buffer

Feeding East/West Applications

Let's suppose that we are interested in the data collected for one of the assets listed in the previous query, for example *fogbench_temperature*. The *asset* entry point can be used to retrieve the data points for individual assets by simply adding the code of the asset to the URI:

```
$ curl -s http://localhost:8081/fledge/asset/fogbench_temperature ; echo
[{"timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41}}, {
↪ "timestamp": "2017-12-18 10:38:12.580", "reading": {"ambient": 33, "object": 7}}]
$
```

Let's see the JSON output on a more readable format:

```
[ { "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41} }
↪,
  { "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41} }
↪,
  { "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41} }
↪,
  { "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41} }
↪,
  { "timestamp": "2017-12-18 10:38:29.652", "reading": {"ambient": 13, "object": 41} }
↪,
```

(continues on next page)

(continued from previous page)

```

{ "timestamp": "2017-12-18 10:38:29.652", "reading": { "ambient": 13, "object": 41} }
↪,
{ "timestamp": "2017-12-18 10:38:29.652", "reading": { "ambient": 13, "object": 41} }
↪,
{ "timestamp": "2017-12-18 10:38:29.652", "reading": { "ambient": 13, "object": 41} }
↪,
{ "timestamp": "2017-12-18 10:38:29.652", "reading": { "ambient": 13, "object": 41} }
↪,
{ "timestamp": "2017-12-18 10:38:29.652", "reading": { "ambient": 13, "object": 41} }
↪,
{ "timestamp": "2017-12-18 10:38:12.580", "reading": { "ambient": 33, "object": 7} } ↪
↪]

```

The JSON structure depends on the sensor and the plugin used to capture the data. In this case, the values shown are:

- **timestamp** : the timestamp generated by the sensors. In this case, since we have inserted 10 times the same value and one time a new value using *fogbench*, the result is 10 timestamps with the same value and one timestamp with a different value.
- **reading** : a JSON structure that is the set of data points provided by the sensor. In this case:
- **ambient** : the ambient temperature in Celsius
- **object** : the object temperature in Celsius. Again, the values are repeated 10 times, due to the iteration executed by *fogbench*, plus an isolated element, so there are 11 readings in total. Also, it is very unlikely that in a real sensor the ambient and the object temperature differ so much, but here we are using a random number generator.

You can dig even more in the data and extract only a subset of the reading. For example, you can select the ambient temperature and limit to the last 5 readings:

```

$ curl -s http://localhost:8081/fledge/asset/fogbench_temperature/ambient?limit=5 ; ↪
↪echo
[ { "ambient": 13, "timestamp": "2017-12-18 10:38:29.652" },
  { "ambient": 13, "timestamp": "2017-12-18 10:38:29.652" },
  { "ambient": 13, "timestamp": "2017-12-18 10:38:29.652" },
  { "ambient": 13, "timestamp": "2017-12-18 10:38:29.652" },
  { "ambient": 13, "timestamp": "2017-12-18 10:38:29.652" } ]
$

```

We have beautified the JSON output for you, so it is more readable.

Note: When you select a specific element in the reading, the timestamp and the element are presented in the opposite order compared to the previous example. This is a known issue that will be fixed in the next version.

Sending Greetings to the Northern Hemisphere

The next and last step is to send data to North, which means that we can take all or some of the data we buffer in Fledge and we can send it to a historian or a database using a North task or microservice.

The OMF Translator

Fledge comes with a North plugin called *OMF Translator*. OMF is the OSIssoft Message Format, which is the message format accepted by the PI Connector Relay OMF. The PI Connector Relay OMF is provided by OSIssoft and it is used to feed the OSIssoft PI System.

- Information regarding OSIssoft are available
- Information regarding OMF are available
- Information regarding the OSIssoft PI System are available

OMF Translator is scheduled as a North task that is executed every 30 seconds (the time may vary, we set it to 30 seconds to facilitate the testing).

Preparing the PI System

In order to test the North task and plugin, first you need to setup the PI system. Here we assume you are already familiar with PI and you have a Windows server with PI installed, up and running. The minimum installation must include the PI System and the PI Connector Relay OMF. Once you have checked that everything is installed and works correctly, you should collect the IP address of the Windows system.

Setting the OMF Translator Plugin

Fledge uses the same *OMF Translator* plugin to send the data coming from the South modules and buffered in Fledge.

Note: In this version, only the South data can be sent to the PI System.

If you are curious to see which categories are available in Fledge, simply type:

```
$ curl -s http://localhost:8081/fledge/category ; echo
{
  "categories":
  [
    {
      "key": "SCHEDULER",
      "description": "Scheduler configuration",
      "displayName": "Scheduler"
    },
    {
      "key": "SMNTR",
      "description": "Service Monitor",
      "displayName": "Service Monitor"
    },
    {
      "key": "rest_api",
      "description": "Fledge Admin and User REST API",
      "displayName": "Admin API"
    },
    {
      "key": "service",
      "description": "Fledge Service",
      "displayName": "Fledge Service"
    },
    {
      "key": "Installation",
      "description": "Installation",
      "displayName": "Installation"
    },
  ],
}
```

(continues on next page)

(continued from previous page)

```

    "key": "General",
    "description": "General",
    "displayName": "General"
  },
  {
    "key": "Advanced",
    "description": "Advanced",
    "displayName": "Advanced"
  },
  {
    "key": "Utilities",
    "description": "Utilities",
    "displayName": "Utilities"
  }
]
}
$

```

For each plugin, you will see corresponding category e.g. For fledge-south-coap the registered category will be { "key": "COAP", "description": "CoAP Listener South Plugin"}. The configuration for the OMF Translator used to stream the South data is initially disabled, all you can see about the settings is:

```

$ curl -sX GET http://localhost:8081/fledge/category/OMF%20to%20PI%20north
{
  "enable": {
    "description": "A switch that can be used to enable or disable execution of the_
↪sending process.",
    "type": "boolean",
    "readonly": "true",
    "default": "true",
    "value": "true"
  },
  "streamId": {
    "description": "Identifies the specific stream to handle and the related_
↪information, among them the ID of the last object streamed.",
    "type": "integer",
    "readonly": "true",
    "default": "0",
    "value": "4",
    "order": "16"
  },
  "plugin": {
    "description": "PI Server North C Plugin",
    "type": "string",
    "default": "OMF",
    "readonly": "true",
    "value": "OMF"
  },
  "source": {
    "description": "Defines the source of the data to be sent on the stream, this_
↪may be one of either readings, statistics or audit.",
    "type": "enumeration",
    "options": [
      "readings",
      "statistics"
    ],
  },

```

(continues on next page)

(continued from previous page)

```

    "default": "readings",
    "order": "5",
    "displayName": "Data Source",
    "value": "readings"
  },
  ...}
$ curl -sX GET http://localhost:8081/fledge/category/Stats%20OMF%20to%20PI%20north
{
  "enable": {
    "description": "A switch that can be used to enable or disable execution of the
→ sending process.",
    "type": "boolean",
    "readonly": "true",
    "default": "true",
    "value": "true"
  },
  "streamId": {
    "description": "Identifies the specific stream to handle and the related
→ information, among them the ID of the last object streamed.",
    "type": "integer",
    "readonly": "true",
    "default": "0",
    "value": "5",
    "order": "16"
  },
  "plugin": {
    "description": "PI Server North C Plugin",
    "type": "string",
    "default": "OMF",
    "readonly": "true",
    "value": "OMF"
  },
  "source": {
    "description": "Defines the source of the data to be sent on the stream, this may
→ be one of either readings, statistics or audit.",
    "type": "enumeration",
    "options": [
      "readings",
      "statistics"
    ],
    "default": "readings",
    "order": "5",
    "displayName": "Data Source",
    "value": "statistics"
  },
  ...}
$
```

At this point it may be a good idea to familiarize with the tool, it will help you a lot in selecting and using data via the REST API. You may remember, we discussed it in the chapter.

First, we can see the list of all the scheduled tasks (the process of sending data to a PI Connector Relay OMF is one of them). The command is:

```

$ curl -s http://localhost:8081/fledge/schedule | jq
{
  "schedules": [
```

(continues on next page)

(continued from previous page)

```

{
  "id": "ef8bd42b-da9f-47c4-ade8-751ce9a504be",
  "name": "OMF to PI north",
  "processName": "north_c",
  "type": "INTERVAL",
  "repeat": 30.0,
  "time": 0,
  "day": null,
  "exclusive": true,
  "enabled": false
},
{
  "id": "27501b35-e0cd-4340-afc2-a4465fe877d6",
  "name": "Stats OMF to PI north",
  "processName": "north_c",
  "type": "INTERVAL",
  "repeat": 30.0,
  "time": 0,
  "day": null,
  "exclusive": true,
  "enabled": true
},
...
]
}
$

```

... which means: “show me all the tasks that can be scheduled”, The output has been made more readable by jq. There are several tasks, we need to identify the one we need and extract its unique id. We can achieve that with the power of jq: first we can select the JSON object that shows the elements of the sending task:

```

$ curl -s http://localhost:8081/fledge/schedule | jq '.schedules[] | select( .name ==
↳ "OMF to PI north") '
{
  "id": "ef8bd42b-da9f-47c4-ade8-751ce9a504be",
  "name": "OMF to PI north",
  "processName": "north_c",
  "type": "INTERVAL",
  "repeat": 30,
  "time": 0,
  "day": null,
  "exclusive": true,
  "enabled": true
}
$

```

Let’s have a look at what we have found:

- **id** is the unique identifier of the schedule.
- **name** is a user-friendly name of the schedule.
- **type** is the type of schedule, in this case a schedule that is triggered at regular intervals.
- **repeat** specifies the interval of 30 seconds.
- **time** specifies when the schedule should run: since the type is INTERVAL, this element is irrelevant.
- **day** indicates the day of the week the schedule should run, in this case it will be constantly every 30 seconds

- **exclusive** indicates that only a single instance of this task should run at any time.
- **processName** is the name of the task to be executed.
- **enabled** indicates whether the schedule is currently enabled or disabled.

Now let's identify the plugin used to send data to the PI Connector Relay OMF.

```
$ curl -s http://localhost:8081/fledge/category | jq '.categories[] | select ( .key_
↪ == "OMF to PI north" )'
{
  "key": "OMF to PI north",
  "description": "Configuration of the Sending Process",
  "displayName": "OMF to PI north"
}
$
```

We can get the specific information adding the name of the task to the URL:

```
$ curl -s http://localhost:8081/fledge/category/OMF%20to%20PI%20north | jq .plugin
{
  "description": "PI Server North C Plugin",
  "type": "string",
  "default": "OMF",
  "readonly": "true",
  "value": "OMF"
}
$
```

Now, the output returned does not say much: this is because the plugin has never been enabled, so the configuration has not been loaded yet. First, let's enable the schedule. From a the previous query of the schedulable tasks, we know the id is *ef8bd42b-da9f-47c4-ade8-751ce9a504be*:

```
$ curl -X PUT http://localhost:8081/fledge/schedule/ef8bd42b-da9f-47c4-ade8-
↪ 751ce9a504be -d '{ "enabled" : true }'
{
  "schedule": {
    "id": "ef8bd42b-da9f-47c4-ade8-751ce9a504be",
    "name": "OMF to PI north",
    "processName": "north_c",
    "type": "INTERVAL",
    "repeat": 30,
    "time": 0,
    "day": null,
    "exclusive": true,
    "enabled": true
  }
}
$
```

Once enabled, the plugin will be executed inside the *OMF to PI north* task within 30 seconds, so you have to wait up to 30 seconds to see the new, full configuration. After 30 seconds or so, you should see something like this:

```
$ curl -s http://localhost:8081/fledge/category/OMF%20to%20PI%20north | jq
{
  "enable": {
    "description": "A switch that can be used to enable or disable execution of the_
↪ sending process.",
    "type": "boolean",
```

(continues on next page)

(continued from previous page)

```

    "readonly": "true",
    "default": "true",
    "value": "true"
  },
  "streamId": {
    "description": "Identifies the specific stream to handle and the related_
↪information, among them the ID of the last object streamed.",
    "type": "integer",
    "readonly": "true",
    "default": "0",
    "value": "4",
    "order": "16"
  },
  "plugin": {
    "description": "PI Server North C Plugin",
    "type": "string",
    "default": "OMF",
    "readonly": "true",
    "value": "OMF"
  },
  "PIServerEndpoint": {
    "description": "Select the endpoint among PI Web API, Connector Relay, OSIssoft_
↪Cloud Services or Edge Data Store",
    "type": "enumeration",
    "options": [
      "PI Web API",
      "Connector Relay",
      "OSIssoft Cloud Services",
      "Edge Data Store"
    ],
    "default": "Connector Relay",
    "order": "1",
    "displayName": "Endpoint",
    "value": "Connector Relay"
  },
  "ServerHostname": {
    "description": "Hostname of the server running the endpoint either PI Web API or_
↪Connector Relay",
    "type": "string",
    "default": "localhost",
    "order": "2",
    "displayName": "Server hostname",
    "validity": "PIServerEndpoint != \"Edge Data Store\" && PIServerEndpoint != \"
↪OSIssoft Cloud Services\"",
    "value": "localhost"
  },
  "ServerPort": {
    "description": "Port on which the endpoint either PI Web API or Connector Relay_
↪or Edge Data Store is listening, 0 will use the default one",
    "type": "integer",
    "default": "0",
    "order": "3",
    "displayName": "Server port, 0=use the default",
    "validity": "PIServerEndpoint != \"OSIssoft Cloud Services\"",
    "value": "0"
  },
  "producerToken": {

```

(continues on next page)

(continued from previous page)

```

    "description": "The producer token that represents this Fledge stream",
    "type": "string",
    "default": "omf_north_0001",
    "order": "4",
    "displayName": "Producer Token",
    "validity": "PIServerEndpoint == \"Connector Relay\"",
    "value": "omf_north_0001"
  },
  "source": {
    "description": "Defines the source of the data to be sent on the stream, this may
↪be one of either readings, statistics or audit.",
    "type": "enumeration",
    "options": [
      "readings",
      "statistics"
    ],
    "default": "readings",
    "order": "5",
    "displayName": "Data Source",
    "value": "readings"
  },
  "StaticData": {
    "description": "Static data to include in each sensor reading sent to the PI
↪Server.",
    "type": "string",
    "default": "Location: Palo Alto, Company: Dianomic",
    "order": "6",
    "displayName": "Static Data",
    "value": "Location: Palo Alto, Company: Dianomic"
  },
  "OMFRetrySleepTime": {
    "description": "Seconds between each retry for the communication with the OMF PI
↪Connector Relay, NOTE : the time is doubled at each attempt.",
    "type": "integer",
    "default": "1",
    "order": "7",
    "displayName": "Sleep Time Retry",
    "value": "1"
  },
  "OMFMaxRetry": {
    "description": "Max number of retries for the communication with the
↪OMF PI Connector Relay",
    "type": "integer",
    "default": "3",
    "order": "8",
    "displayName": "Maximum Retry",
    "value": "3"
  },
  "OMFHttpTimeout": {
    "description": "Timeout in seconds for the HTTP operations with the OMF PI
↪Connector Relay",
    "type": "integer",
    "default": "10",
    "order": "9",
    "displayName": "HTTP Timeout",
    "value": "10"
  },

```

(continues on next page)

(continued from previous page)

```

"formatInteger": {
  "description": "OMF format property to apply to the type Integer",
  "type": "string",
  "default": "int64",
  "order": "10",
  "displayName": "Integer Format",
  "value": "int64"
},
"formatNumber": {
  "description": "OMF format property to apply to the type Number",
  "type": "string",
  "default": "float64",
  "order": "11",
  "displayName": "Number Format",
  "value": "float64"
},
"compression": {
  "description": "Compress readings data before sending to PI server",
  "type": "boolean",
  "default": "true",
  "order": "12",
  "displayName": "Compression",
  "value": "false"
},
"DefaultAFLocation": {
  "description": "Defines the hierarchies tree in Asset Framework in which the
↳ assets will be created, each level is separated by /, PI Web API only.",
  "type": "string",
  "default": "/fledge/data_piwebapi/default",
  "order": "13",
  "displayName": "Asset Framework hierarchies tree",
  "validity": "PIServerEndpoint == \"PI Web API\"",
  "value": "/fledge/data_piwebapi/default"
},
"AFMap": {
  "description": "Defines a set of rules to address where assets should be placed
↳ in the AF hierarchy.",
  "type": "JSON",
  "default": "{ }",
  "order": "14",
  "displayName": "Asset Framework hierarchies rules",
  "validity": "PIServerEndpoint == \"PI Web API\"",
  "value": "{ }"
},
"notBlockingErrors": {
  "description": "These errors are considered not blocking in the communication
↳ with the PI Server, the sending operation will proceed with the next block of data
↳ if one of these is encountered",
  "type": "JSON",
  "default": "{ \"errors400\" : [ \"Redefinition of the type with the same ID is
↳ not allowed\", \"Invalid value type for the property\", \"Property does not exist
↳ in the type definition\", \"Container is not defined\", \"Unable to find the
↳ property of the container of type\" ] }",
  "order": "15",
  "readonly": "true",
  "value": "{ \"errors400\" : [ \"Redefinition of the type with the same ID is not
↳ allowed\", \"Invalid value type for the property\", \"Property does not exist in
↳ the type definition\", \"Container is not defined\", \"Unable to find the property
↳ of the container of type\" ] }"

```

(continues on next page)

(continued from previous page)

```

    },
    "PIWebAPIAuthenticationMethod": {
      "description": "Defines the authentication method to be used with the PI Web API.
↪",
      "type": "enumeration",
      "options": [
        "anonymous",
        "basic",
        "kerberos"
      ],
      "default": "anonymous",
      "order": "17",
      "displayName": "PI Web API Authentication Method",
      "validity": "PIServerEndpoint == \"PI Web API\"",
      "value": "anonymous"
    },
    "PIWebAPIUserId": {
      "description": "User id of PI Web API to be used with the basic access_
↪authentication.",
      "type": "string",
      "default": "user_id",
      "order": "18",
      "displayName": "PI Web API User Id",
      "validity": "PIServerEndpoint == \"PI Web API\" && PIWebAPIAuthenticationMethod_
↪== \"basic\"",
      "value": "user_id"
    },
    "PIWebAPIPassword": {
      "description": "Password of the user of PI Web API to be used with the basic_
↪access authentication.",
      "type": "password",
      "default": "password",
      "order": "19",
      "displayName": "PI Web API Password",
      "validity": "PIServerEndpoint == \"PI Web API\" && PIWebAPIAuthenticationMethod_
↪== \"basic\"",
      "value": "*****"
    },
    "PIWebAPIKerberosKeytabFileName": {
      "description": "Keytab file name used for Kerberos authentication in PI Web API.",
      "type": "string",
      "default": "piwebapi_kerberos_https.keytab",
      "order": "20",
      "displayName": "PI Web API Kerberos keytab file",
      "validity": "PIServerEndpoint == \"PI Web API\" && PIWebAPIAuthenticationMethod_
↪== \"kerberos\"",
      "value": "piwebapi_kerberos_https.keytab"
    },
    "OCSNamespace": {
      "description": "Specifies the OCS namespace where the information are stored and_
↪it is used for the interaction with the OCS API",
      "type": "string",
      "default": "name_space",
      "order": "21",
      "displayName": "OCS Namespace",
      "validity": "PIServerEndpoint == \"OSIsoft Cloud Services\"",
      "value": "name_space"
    }
  }

```

(continues on next page)

(continued from previous page)

```

},
"OCSTenantId": {
  "description": "Tenant id associated to the specific OCS account",
  "type": "string",
  "default": "ocs_tenant_id",
  "order": "22",
  "displayName": "OCS Tenant ID",
  "validity": "PIServerEndpoint == \"OSIsoft Cloud Services\"",
  "value": "ocs_tenant_id"
},
"OCSClientId": {
  "description": "Client id associated to the specific OCS account, it is used to_
↪authenticate the source for using the OCS API",
  "type": "string",
  "default": "ocs_client_id",
  "order": "23",
  "displayName": "OCS Client ID",
  "validity": "PIServerEndpoint == \"OSIsoft Cloud Services\"",
  "value": "ocs_client_id"
},
"OCSClientSecret": {
  "description": "Client secret associated to the specific OCS account, it is used_
↪to authenticate the source for using the OCS API",
  "type": "password",
  "default": "ocs_client_secret",
  "order": "24",
  "displayName": "OCS Client Secret",
  "validity": "PIServerEndpoint == \"OSIsoft Cloud Services\"",
  "value": "*****"
}
}
$

```

You can look at the descriptions to have a taste of what you can control with this plugin. The default configuration should be fine, with the exception of the *ServerHostname*, which of course should refer to the IP address of the machine and the port used by the PI Connector Relay OMF. The PI Connector Relay OMF 1.0 used the HTTP protocol with port 8118 and version 1.2, or higher, uses the HTTPS and port 5460. Assuming that the port is 5460 and the IP address is 192.168.56.101, you can set the new *ServerHostname* with this PUT method:

```


$ curl -sH'Content-Type: application/json' -X PUT -d '{ "ServerHostname": "192.168.56.
↪101" }' http://localhost:8081/fledge/category/OMF%20to%20PI%20north | jq
"ServerHostname": {
  "description": "Hostname of the server running the endpoint either PI Web API or_
↪Connector Relay",
  "type": "string",
  "default": "localhost",
  "order": "2",
  "displayName": "Server hostname",
  "validity": "PIServerEndpoint != \"Edge Data Store\" && PIServerEndpoint != \"
↪OSIsoft Cloud Services\"",
  "value": "192.168.56.101"
}
$

```


You can note that the *value* element is the only one that can be changed in *URL* (the other elements are factory settings). Now we are ready to send data North, to the PI System.

Sending Data to the PI System

The last bit to accomplish is to start the PI Connector Relay OMF on the Windows Server. The output may look like this screenshot, where you can see the Connector Relay debug window on the left and the PI Data Explorer on the right.



Wait a few seconds ... et voilà! Readings and statistics are in the PI System:



Congratulations! You have experienced an end-to-end test of Fledge, from South with sensor data through Fledge and East/West applications and finally to North towards Historians.

13.1.5 Fledge Utilities and Scripts

The Fledge platform comes with a set of utilities and scripts to help users, developers and administrators with their day-by-day operations. These tools are under heavy development and you may expect incompatibilities in future versions, therefore it is highly recommended to check the revision history to verify the changes in new versions.

fledge

`fledge` is the first utility available with the platform, it is the control center for all the admin operations on Fledge.

In the current implementation, *fledge* provides these features:

- *start* Fledge
- *stop* Fledge
- *kill* Fledge processes
- Check the *status* of Fledge, i.e. whether it is running, starting or not running
- *reset* Fledge to its factory settings

Starting Fledge

`fledge start` is the command to start Fledge. Since only one core microservice of Fledge can be executed in the same environment, the command checks if Fledge is already running, and if it does, it ends. The command also checks the presence of the `FLEDGE_ROOT` and `FLEDGE_DATA` environment variables. If the variables have not been set, it verifies if Fledge has been installed in the default position, which is `/usr/local/fledge` or a position defined by the installed package, and it will set the missing variables accordingly. It will also take care of the `PYTHONPATH` variable.

In more specific terms, the command executes these steps:

- Check if Fledge is already running
- Check if the storage layer is *managed* or *unmanaged*. “managed” means that the storage layer relies on a storage system (i.e. a database, a set of files or in-memory structures) that are under exclusive control of Fledge. “unmanaged” means that the storage system is generic and potentially shared with other applications.
- Check if the storage plugin and the related storage system (for example a PostgreSQL database) is available.
- Check if the metadata structure that is necessary to execute Fledge is already available in the storage layer. If the metadata is not available, it creates the data model and sets the factory settings that are necessary to start and use Fledge.
- Start the core microservice.
- Wait until the core microservice starts the Storage microservice and the initial required process that are necessary to handle other tasks and microservices.

Safe Mode

It is possible to start Fledge in safe mode by passing the flag `--safe-mode` to the start command. In safe mode Fledge will not start any of the south services or schedule any tasks, such as purge or north bound tasks. Safe mode allows Fledge to be started and configured in those situations where a previous misconfiguration has rendered it impossible to start and interact with Fledge.

Once started in safe mode any configuration changes should be made and then Fledge should be restarted in normal mode to test those configuration changes.

Stopping Fledge

`fledge stop` is the command used to stop Fledge. The command waits until all the tasks and services have been completed, then it stops the core service.

If Fledge Does Not Stop

If Fledge does not stop, i.e. if by using the process status command `ps` you see Fledge processes still running, you can use `fledge kill` to kill them.

Note: The command issues a `kill -9` against the processes associated to Fledge. This is not recommended, unless Fledge cannot be stopped. The `stop` command. In other words, *kill* is your last resort before a reboot. If you must use the kill command, it means that there is a problem: please report this to the Fledge project slack channel.

Checking the Status of Fledge

`fledge status` is used to provide the current status of tasks and microservices on the machine. The output is something like:

```
$ fledge status
Fledge running.
Fledge uptime: 2034 seconds.
=== Fledge services:
```

(continues on next page)

(continued from previous page)

```
fledge.services.core
fledge.services.south --port=33074 --address=127.0.0.1 --name=HTTP_SOUTH
fledge.services.south --port=33074 --address=127.0.0.1 --name=COAP
=== Fledge tasks:
$ fledge_use_from_here stop
Fledge stopped.
$ fledge_use_from_here status
Fledge not running.
$
```

- The first row always indicates if Fledge is running or not
- The second row provides the uptime in seconds
- The next set of rows provides information regarding the microservices running on the machine
- The last set of rows provides information regarding the tasks running on the machine

Resetting Fledge

It may occur that you want to restore Fledge to its factory settings, and this is what `fledge reset` does. The command also destroys all the data and all the configuration currently stored in Fledge, so you must use it at your own risk!

Fledge can be restored to its factory settings only when it is not running, hence you should stop it first.

The command forces you to insert the word *YES*, all in uppercase, to continue:

```
$ fledge reset
This script will remove all data stored in the server.
Enter YES if you want to continue: YES
$
```

13.1.6 Fledge Tasks

Tasks are part of the Fledge IoT platform. They are like services, but with a clear distinction:

- *services* are started at a certain point (usually at startup) and they are likely to continue to work until Fledge stops.
- *tasks* are started when required, they execute a job and then they terminate.

In simple terms, a service is meant to always listen and react to requests, while a task is triggered by an event and then when job is terminated, the tasks ends.

That said, tasks and services shared these same features:

- They are both started by the Fledge scheduler. It is likely that services are started at startup, while tasks can start at a given time or interval.
- They both use the internal API to communicate with other services.
- They both use the same pluggable architecture to separate a common logic, usually associated to the internal features of Fledge, from a more generic logic, usually closer to the type of operations that must be performed.

In this chapter we present a set of tasks that are commonly available in Fledge.

Purge

The *Purge* task is triggered by the scheduler to purge old data that is still stored (buffered) in Fledge. The logic applied to the task is relatively simple:

- The task is called exclusively (i.e. there cannot be more than one *Purge* task running at any given time) by the Fledge scheduler every hour (by default).
- Data that is older than a certain date/time is removed.
- Optionally, data is removed if the total size of the stored objects is bigger than 1GByte (default)
- Optionally, data is not removed if it has not been extracted and used by any North task or service yet.
- All purge operations are stored in the audit log.

Purge Schedule

Purge is one of the tasks launched by the Fledge scheduler. You can retrieve information about the scheduling by calling the *GET* method of the *schedule* call. The name and the process name of the task are both *purge*:

```
$ curl -sX GET http://localhost:8081/fledge/schedule
...
{ "id"          : "cea17db8-6ccc-11e7-907b-a6006ad3dba0",
  "name"        : "purge",
  "time"        : 0,
  "enabled"     : true,
  "repeat"      : 3600,
  "type"        : "INTERVAL",
  "exclusive"   : true,
  "processName" : "purge",
  "day"         : null },
...
$
```

As you can see from the JSON output, the task is scheduled to be executed every hour (3,600 seconds). In order to change the interval between *Purge* tasks, you can call the *PUT* method of the *schedule* call by passing the associated *id*. For example, in order to change the task to be executed any 5 minutes (i.e. 300 seconds) you should call:

```
$ curl -sX PUT http://localhost:8081/fledge/schedule/cea17db8-6ccc-11e7-907b-a6006ad3dba0 -d '{"repeat": 300}'
{"schedule": { "id": "cea17db8-6ccc-11e7-907b-a6006ad3dba0",
               "name"      : "purge",
               "time"      : 0,
               "enabled"   : true,
               "repeat"    : 300,
               "type"      : "INTERVAL",
               "exclusive" : true,
               "processName": "purge",
               "day"       : null }
}
$
```

Purge Configuration

The configuration of the *Purge* task is stored in the metadata structures of Fledge and it can be retrieve using the *GET* method of the *category/PURGE_READ* call. This is the command used to retrieve the configuration in JSON format:

```
$ curl -sX GET http://localhost:8081/fledge/category/PURGE_READ
{ "retainUnsent" : { "type": "boolean",
                    "default": "False",
                    "description": "Retain data that has not been sent to any
↳historian yet.",
                    "value": "False" },
  "age"           : { "type": "integer",
                    "default": "72",
                    "description": "Age of data to be retained, all data that is
↳older than this value will be removed, unless retained. (in Hours)",
                    "value": "72" },
  "size"          : { "type": "integer",
                    "default": "1000000",
                    "description": "Maximum size of data to be retained, the oldest
↳data will be removed to keep below this size, unless retained. (in Kbytes)",
                    "value": "1000000" } }
```

Changes can be applied using the *PUT* method for each parameter call. For example, in order to change the retention policy for data that has not been sent to historians yet, you can use this call:

```
$ curl -sX PUT http://localhost:8081/fledge/category/PURGE_READ/retainUnsent -d '{
↳"value": "True"}'
{ "type": "boolean",
  "default": "False",
  "description": "Retain data that has not been sent to any historian yet.",
  "value": "True" }
```

The following table shows the list of parameters that can be changed in the *Purge* task:

Item	Type	Default	Description
retainUnsent	boolean	False	Retain data that has not been sent to “North” yet When <i>True</i> , data that has not yet been retrieved by any North service or task, will not be purged. When <i>False</i> , data is purged without checking whether it has been sent to a North destination yet or not.
age	integer	72	Age in hours of the data to be retained. Data that is older than this value, will be purged.
size	integer	1000000	Size in KBytes of data that will be retained in Fledge. Older data will be removed to keep the data stored in Fledge below this size.

13.2 Building and using Fledge on Raspbian

Fledge requires the use of Python 3.5.3+ in order to support the asynchronous IO mechanisms used by Fledge. Earlier Raspberry Pi Raspbian distributions support Python 3.4 as the latest version of Python. In order to build and run Fledge on Raspbian the version of Python must be updated manually if your distribution has an older version.

NOTE: These steps must be executed *in addition* to what is described in the README file when you install Fledge on Raspbian.

Check your Python version by running the command

```
$ python3 --version
$
```


If your version is less than 3.5.3 then follow the instructions below to update your Python version.

Install and update the build tools required for Python to be built

```
$ sudo apt-get update
$ sudo apt-get install build-essential tk-dev
$ sudo apt-get install libncurses5-dev libncursesw5-dev libreadline6-dev
$ sudo apt-get install libdb5.3-dev libgdbm-dev libsqlite3-dev libssl-dev
$ sudo apt-get install libbz2-dev libexpat1-dev liblzma-dev zlib1g-dev
$
```

Now build and install the new version of Python

```
$ wget https://www.python.org/ftp/python/3.5.3/Python-3.5.3.tgz
$ tar zxvf Python-3.5.3.tgz
$ cd Python-3.5.3
$ ./configure
$ make
$ sudo make install
```

Confirm the Python version

```
$ python3 --version
$ pip3 --version
```

These should both return a version number as 3.5.3+, if not then check which python3 and pip3 you are running and replace these with the newly built versions. This may be caused by the newly built version being installed in /usr/local/bin and the existing python3 and pip3 being in /usr/bin. If this is the case then remove the /usr/bin versions

```
$ sudo rm /usr/bin/python3 /usr/bin/pip3
```

You may also link to the new version if you wish

```
$ sudo ln -s /usr/bin/python3 /usr/local/bin/python3
$ sudo ln -s /usr/bin/pip3 /usr/local/bin/pip3
```

Once python3.5 has been installed you may follow the instructions in the README file to build, install and run Fledge on Raspberry Pi using the Raspbian distribution.

13.3 Building and using Fledge on RedHat/CentOS

Fledge can be built or installed on Red Hat or CentOS, it is currently tested against:

- Red Hat 7
- CentOS 7

You may follow the instructions in the README file to build, install and run Fledge on Red Hat or CentOS.

13.4 Install Fledge on Red Hat/CentOS using the RPM package

The Fledge RPM is available in the download page of the documentation available at [download page](#).

The RPM can also be created from the Fledge sources through the repository [fledge-pkg](#) using the make_rpm script and following the instruction in the README.rst.

13.4.1 Installation on Red Hat

It is necessary to install a Red Hat package before Fledge can be installed successfully. The installation sequence is as follows:

```
$ sudo yum-config-manager --enable 'Red Hat Software Collections RPMs for Red Hat_
↳Enterprise Linux 7 Server from RHUI'
$ sudo yum -y localinstall ~/fledge-1.8.0-1.00.x86_64.rpm
```

13.4.2 Installation on CentOS

It is necessary to install a CentOS package before Fledge can be installed successfully. The installation sequence is as follows:

```
$ sudo yum install -y centos-release-scl-rh
$ sudo yum -y localinstall ~/fledge-1.8.0-1.00.x86_64.rpm
```

Note: By default, `/var/log/messages` are created with read-write permissions for ‘root’ user only. Make sure to set the correct READ permissions.

`sudo chmod 644 /var/log/messages`

13.5 Build of Fledge on Red Hat/CentOS

A gcc version newer than 4.9.0 is needed to properly use <regex> and build Fledge.

The *requirements.sh* script, executed as follows:

```
$ sudo ./requirements.sh
```

installs *devtoolset-7* that provides the newer compiler.

It must be enabled before building Fledge using:

```
$ source scl_source enable devtoolset-7
```

It is possible to use the following command to verify which version is currently active:

```
$ gcc --version
```

The previously installed gcc will be by default enabled again after a logoff/login.

13.5.1 Build and use Fledge with PostgreSQL for Red Hat/CentOS

The *rh-postgresql96* environment should be enabled using:

```
$ source scl_source enable rh-postgresql96
```

before building Fledge if the intention is to use the Postgres plugin.

14.1 Fledge v1

14.1.1 v1.9.2

Release Date: 2021-09-29

- **Fledge 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 OPC/UA south plugin has been created based on the Safe and Secure OPC/UA library. This plugin supports authentication and encryption mechanisms.
 - * 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 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 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.
- 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.
 - * The CSV playback south plugin installation on CentOS 7 platforms has now been corrected.
 - * 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 example path shown in the DHT11 plugin in the developers guide was incorrect, this has now been fixed.

14.1.2 v1.9.1

Release Date: 2021-05-27

- **Fledge 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 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.
- * 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.
- * Fixed a bug that treated arrays of bytes as if they were strings in the OPC/UA south plugin.
- * 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.
- * An issue with the SQLite In Memory storage plugin that caused database locks under high load conditions has been resolved.

14.1.3 v1.9.0

Release Date: 2021-02-19

- **Fledge 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 Fledge 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.

14.1.4 v1.8.2

Release Date: 2020-11-03

- **Fledge 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.

14.1.5 v1.8.1

Release Date: 2020-07-08

- **Fledge 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 Fledge 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 Fledge. 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 been 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 dealing 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 Fledge-iot 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 fledge-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.

14.1.6 v1.8.0

Release Date: 2020-05-08

- **Fledge 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 Fledge instances in different time zones has been resolved.
 - * A new API entry point has been added to the Fledge 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 Fledge. 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 Fledge.
 - * Data on multiple assets may now be requested in a single call to the asset growing API within Fledge.
 - * 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 Fledge.
 - * 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 fledge 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 Fledge. This slowed down the user interface and put load n the server. This has now been removed and has improved the performance of examining the buffered data within the Fledge instance.
 - * The FogBench utility used to send data to Fledge 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 cause 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 Fledge 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 Fledge buffer.

• GUI

– New Features:

- * The user interface for viewing logs has been improve 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 Fledge 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 rely, 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 fledge-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 Twin-CAT 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 Fledge's into the PI Server using PI Web API, this has now been resolved.
 - * The filter Fledge 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 Fledge 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.

14.1.7 v1.7.0

Release Date: 2019-08-15

- **Fledge 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.
 - * Fledge 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 Fledge 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 Fledge installation.
 - * An API has been added that allows the caller to retrieve the list of plugins that are available in the Fledge package repository.
 - * The /fledge/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, `FLEDGE_PLUGIN_PATH` has been added to allow plugins to be stored in multiple locations or locations outside of the usual Fledge installation directory.
 - * Added support for `FLEDGE_PLUGIN_PATH` environment variable, that would be used for searching additional directory paths for plugins/filters to use with Fledge.
 - * Fledge 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 *fledge status* command was corrupt, showing random text after the number of seconds for which fledge 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 Fledge 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 Fledge 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 Fledge instance must have connectivity to the package repository to allow this feature to succeed.
- * The GUI now supports using certificates to authenticate with the Fledge 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 Fledge 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 Fledge 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 Fledge 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 Fledge 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 Fledge 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 Fledge, 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 fledge-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 Fledge. 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 from the system to alleviate this problem.

The plugins involved are,

- fledge-north-http Vs fledge-north-http-north
- fledge-south-http Vs fledge-south-http-south
- fledge-south-Csv Vs fledge-south-csv
- fledge-south-Expression Vs fledge-south-expression
- fledge-south-dht Vs fledge-south-dht11V2
- fledge-south-modbus Vs fledge-south-modbus

14.1.8 v1.6.0

Release Date: 2019-05-22

• Fledge Core

– New Features:

- * The scope of the Fledge 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 Fledge has been added to the packaging options for Fledge in this version of Fledge.
- * Red Hat Enterprise Linux packages have been made available from this release of Fledge onwards. These packages include all the applicable plugins and notification service for Fledge.
- * The Fledge API now supports the creation of configuration snapshots which can be used to create configuration checkpoints and rollback configuration changes.
- * The Fledge administration API has been extended to allow the installation of new plugins via API.

– Improvements/Bug Fix:

- * A bug that prevents multiple Fledge's on the same network being discoverable via multicast DNS lookup has been fixed.
- * Set, unset optional configuration attributes

- **GUI**

- **New Features:**

- * The Fledge 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 Fledge 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 Fledge 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 Fledge 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 Fledge 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 chatroom.
 - * 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 hard 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.

14.1.9 v1.5.2

Release Date: 2019-04-08

- **Fledge 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 Fledge 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 Fledge. 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

14.1.10 v1.5.1

Release Date: 2019-03-12

- **Fledge Core**

- Bug Fix: plugin loading errors

- **GUI**

- Bug Fix: uptime shows up to 24 hour clock only

14.1.11 v1.5.0

Release Date: 2019-02-21

- **Fledge Core**
 - Performance improvements and Bug Fixes
 - Introduction of Safe Mode in case Fledge 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 Fledge.
 - 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**

14.1.12 v1.4.1

Release Date: 2018-10-10

14.1.13 v1.4.0

Release Date: 2018-09-25

14.1.14 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

14.1.15 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/fledge-iot/fledge-south-coap>
 - **http:** http south plugin was moved into its own repository <https://github.com/fledge-iot/fledge-south-http>

Known Issues

- **Issues in Documentation**
 - **plugin documentation:** testing Fledge requires user to first install southbound plugins necessary (CoAP, http)

14.1.16 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 Fledge 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**

- **fledge 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 `$FLEDGE_ROOT/extras/scripts` and it is called `setenv.sh`.
- **fledge service script**: a new service script has been added to setup Fledge as a service. The script is in `$FLEDGE_ROOT/extras/scripts` and it is called `fledge.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"`

14.1.17 v1.1.1

Release Date: 2018-01-18

New Features

- **Fixed aiohttp incompatibility**: This fix is for the incompatibility of *aiohttp* with *yaml*, discovered in the previous version. The issue has been fixed.
- **Fixed avahi-daemon issue**: Avahi daemon is a pre-requisite of Fledge, Fledge 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.

14.1.18 v1.1

Release Date: 2018-01-09

New Features

- **Startup Script**:
 - `fledge start` script now checks if the Core microservice has started.
 - `fledge start` creates a `core.err` file in `$FLEDGE_DATA` and writes the stderr there.

Known Issues

- **Incompatibility between aiohttp and yarl when Fledge 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.

14.1.19 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 Fledge
- 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:** `fledge start` does not check if the Core microservice has started correctly, hence it may report that “Fledge started.” when the process has died. As a workaround, check with `fledge status` the presence of the Fledge microservices.
- **Snap Execution on Raspbian:** there is an issue on Raspbian when the Fledge 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/libarmmmem.so`. Save the file and you will be able to immediately use the snap.
- **OMF Translator North Plugin for Fledge Statistics:** in this version the statistics collected by Fledge 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 Fledge 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>`.

15.1 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.

-

15.2 Download/Clone from GitHub

Fledge and the Fledge tools are on GitHub. You can view and download them here:

- **Fledge:** This is the main project for the Fledge platform.
- **Fledge GUI:** This is an experimental GUI that connects to the Fledge REST API to configure and administer the platform and to retrieve the data buffered in it.

There are many south, north, and filter plugins available on github:

16.1 Introduction

The bundled OMF north plugin in Fledge 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 Fledge *requirements.sh* script installs the Kerberos client to allow the integration with what in the specific terminology is called KDC (the Kerberos server).

16.2 PI-Server as the North endpoint

The OSI *Connector Relay* allows token authentication while *PI Web API* supports Basic and Kerberos.

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.

16.3 North plugin

The North plugin has a set of configurable options that should be changed, using either the Fledge API or the Fledge 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:

```
${FLEDGE_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.

16.4 Fledge server configuration

The server on which Fledge 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*

16.4.1 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*

16.4.2 Kerberos client configuration

The server on which Fledge 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
    }
```

16.4.3 Kerberos keytab file

The keytab file should be generated on the Kerberos server and copied into the Fledge server in the directory:

```
${FLEDGE_DATA}/etc/kerberos
```

NOTE:

- if **FLEDGE_DATA** is not set its value should be `$FLEDGE_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 ${FLEDGE_DATA}/etc/kerberos
-rwxrwxrwx 1 fledge fledge  91 Jul 17 09:07 piwebapi_kerberos_https.keytab
-rw-rw-r-- 1 fledge fledge 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
```

16.4.4 Troubleshooting the Kerberos authentication

- 1) check the North plugin configuration, a sample command

```
curl -s -S -X GET http://localhost:8081/fledge/category/North_Readings_to_PI | jq ".|
↪{URL, \"PIServerEndpoint\", PIWebAPIAuthenticationMethod, PIWebAPIKerberosKeytabFileName,
↪AFHierarchyLevel}\"
```

- 2) check the presence of the keytab file

```
$ ls -l ${FLEDGE_ROOT}/data/etc/kerberos
-rwxrwxrwx 1 fledge fledge  91 Jul 17 09:07 piwebapi_kerberos_https.keytab
-rw-rw-r-- 1 fledge fledge 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
$
```

16.5 Kerberos authentication on RedHat/CentOS

RedHat and CentOS version 7 provide by default an old version of curl and the related libcurl and it does not support Kerberos, output of the curl provided by CentOS:

```
$ curl -V
curl 7.29.0 (x86_64-redhat-linux-gnu) libcurl/7.29.0 NSS/3.36 zlib/1.2.7 libidn/1.28_
↳libssh2/1.4.3
Protocols: dict file ftp ftps gopher http https imap imaps ldap ldaps pop3 pop3s rtsp_
↳scp sftp smtp smtps telnet tftp
Features: AsynchDNS GSS-Negotiate IDN IPv6 Largefile NTLM NTLM_WB SSL libz unix-
↳sockets
```

The *requirements.sh* evaluates if the default version 7.29.0 is installed and in this case it will download the sources, build and install the version 7.65.3 to provide Kerberos authentication, output of the curl after the upgrade:

```
$ curl -V
curl 7.65.3 (x86_64-unknown-linux-gnu) libcurl/7.65.3 OpenSSL/1.0.2k-fips zlib/1.2.7
Release-Date: 2019-07-19
Protocols: dict file ftp ftps gopher http https imap imaps pop3 pop3s rtsp smb smbs_
↳smtp smtps telnet tftp
Features: AsynchDNS GSS-API HTTPS-proxy IPv6 Kerberos Largefile libz NTLM NTLM_WB_
↳SPNEGO SSL UnixSockets
```

The sources are downloaded from the curl repository [curl sources](#), the curl homepage is available at [curl homepage](#).

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

17.1 Fledge South Plugins

17.1.1 AM2315 Temperature & Humidity Sensor



The *fledge-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°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*

The screenshot shows a configuration window with a progress bar at the top. The progress bar has three steps: 1. Plugin & Service Name, 2. Review Configuration (highlighted with a green circle), and 3. Done. Below the progress bar, there is a form with two input fields. The first field is labeled 'Asset Name' and contains the text 'am2315/%M/'. The second field is labeled 'I2C Address' and contains the text '0x5C'. 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. 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

17.1.2 CC2650 SensorTag



The *fledge-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 Fledge. 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*

The screenshot shows the 'Review Configuration' step of the Fledge plugin configuration. The progress bar at the top indicates three steps: 1. Plugin & Service Name, 2. Review Configuration (current), and 3. Done. The configuration form includes the following fields:

- Bluetooth Address:** B0:91:22:EA:79:04
- Asset Name Prefix:** CC2650/%M/
- Shutdown Threshold:** 10
- Connection Timeout:** 3
- Temperature Sensor:** ☒
- Temperature Sensor Name:** temperature
- Luminance Sensor:** ☐
- Luminance Sensor Name:** luminance
- Humidity Sensor:** ☐
- Humidity Sensor Name:** humidity
- Pressure Sensor:** ☐
- Pressure Sensor Name:** pressure
- Movement Sensor:** ☐
- Gyroscope Sensor Name:** gyroscope
- Accelerometer Sensor Name:** accelerometer
- Magnetometer Sensor Name:** magnetometer
- Battery Data:** ☐
- Battery Sensor Name:** battery

At the bottom of the form are 'Previous' and 'Next' buttons.

- 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*

17.1.3 CoAP

The *fledge-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 shows the Fledge configuration interface for the CoAP plugin. At the top, there is a progress bar with three steps: 1. Plugin & Service Name, 2. Review Configuration, and 3. Done. The 'Review Configuration' step is currently active. Below the progress bar, there is a form with two input fields: 'Port' and 'URI'. The 'Port' field contains the value '5683' and the 'URI' field contains the value 'sensor-values'. At the bottom of the form, there are two buttons: 'Previous' and 'Next'.

- 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*

17.1.4 Simple CSV Plugin

The *fledge-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*

The screenshot shows the Fledge configuration interface for the CSV plugin. At the top, there is a progress bar with 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 four fields: 'Asset Name' with the value 'Vibration', 'Datapoint' with the value 'ch', 'Multi-Column' which is an unchecked checkbox, and 'Path Of File' which is an empty text box. 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
 - **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 Fledge process should have permission to read this file.
- Click *Next*
- Enable the service and click on *Done*

17.1.5 CSV Playback

The plugin plays a csv file inside some given directory in file system (The default being FLEDGE_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.

The screenshot displays the 'Review Configuration' step of the Fledge CSV Playback plugin setup. The configuration parameters are as follows:

- Asset name:** vibration
- CSV directory name:** FLEDGE_DATA
- CSV file pattern:** (empty)
- Header processing method:** do_not_skip
- Data point for header rows:** metadata
- Number of rows to skip or pass in datapoint:** 1
- Dynamic columns:** ☐
- Column processing method:** pick_from_file
- Auto generate prefix:** column
- Column names and data types for explicit:** (empty)

- **‘assetName’:** type: string default: **‘vibration’**: The output asset that contains the readings.
- **‘csvDirName’:** type: string default: **‘FLEDGE_DATA’**: The directory where CSV file exists. Default is FLEDGE_DATA or FLEDGE_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"/>
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 FLEDGE_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/fledge/service -d @- << EOF | jq
↪ '. '
{
  "name": "csv_player",
  "type": "south",
  "plugin": "csvplayback",
  "enabled": false,
  "config": {
    "assetName": {"value": "My_csv_asset"},
    "csvDirName": {"value": "FLEDGE_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=$FLEDGE_ROOT/python/fledge/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=$FLEDGE_ROOT/python/fledge/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/fledge/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)

17.1.6 DHT11 (C version)



The *fledge-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 *fledge-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*

1 Plugin & Service Name 2 Review Configuration 3 Done


Asset Name: dht11

Rpi Pin: 7

Previous 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*

17.1.7 DHT11 (Python version)



The *fledge-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 *fledge-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*

The screenshot shows a three-step configuration process. Step 1, 'Plugin & Service Name', is highlighted with a green circle and line. Step 2, 'Review Configuration', is also highlighted with a green circle and line. Step 3, 'Done', is marked with a grey circle. The 'Review Configuration' step displays two input fields: 'Asset Name' with the value 'dht11' and 'GPIO Pin' with the value '4'. At the bottom, there are 'Previous' and 'Next' buttons.

- 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*

17.1.8 DNP3 Master Plugin

The *fledge-south-dnp3* allows Fledge 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 Fledge. The DNP3 plugin also handles unsolicited messages transmitted by the outstation.

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name prefix

Master link Id

Outstation address

Outstation port

Outstation link Id

Data scan ☐

Scan interval

Network timeout

[Previous](#) [Next](#)

- **Asset Name prefix:** An asset name prefix that is prepended to the DNP3 objects retrieved from the DNP3 outstations to create the Fledge asset name.
- **Master link id:** The master link id Fledge 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 Fledge plugin with these parameters should allow Fledge 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.

17.1.9 Enviro pHAT Plugin



The *fledge-south-enviophat* 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 Fledge 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 *envirophat* from the plugin list
- Name your service and click *Next*

The screenshot displays the 'Review Configuration' step of the Fledge plugin setup. A progress bar at the top indicates the current step is 2 of 3. The configuration form includes the following settings:

Field	Value
Asset Name Prefix	e_
RGB Sensor	<input checked="" type="checkbox"/>
RGB Sensor Name	rgb
Magnetometer Sensor	<input checked="" type="checkbox"/>
Magnetometer Sensor Name	magnetometer
Accelerometer Sensor	<input checked="" type="checkbox"/>
Accelerometer Sensor Name	accelerometer
Weather Sensor	<input checked="" type="checkbox"/>
Weather Sensor Name	weather

Navigation buttons at the bottom are 'Previous' and 'Next'.

- 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*

17.1.10 Expression South Plugin

The *fledge-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 'Review Configuration' step of the plugin setup. The configuration fields are as follows:

Field	Value
Asset Name	Expression
Expression	<code>clamp(-1.0,sin(2 * pi * x) + cos(x / 2 * pi),+1.0)</code>
Minimum Value	-5
Maximum Value	5
Step Value	0.001

The parameters that can be configured are;

- **Asset Name:** The name of the asset to be created inside Fledge.
- **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 *fledge-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)

17.1.11 Flir AX8 Thermal Imaging Camera



The *fledge-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.

The screenshot shows the configuration screen for the FlirAX8 plugin. At the top, there's a progress bar with three steps: 1. Plugin & Service Name, 2. Review Configuration, and 3. Done. The configuration fields are: Asset Name (AX8), Server Address (192.168.0.48), Port (502), and Slave ID (1). Navigation buttons 'Previous' and 'Next' are at the bottom.

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 Fledge 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 *fledge-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.

17.1.12 South HTTP

The *fledge-south-http* plugin allows data to be received from another Fledge instance or external system using a REST interface. The Fledge 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 Fledge 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 Fledge format.

The primary purpose of this plugin is for Fledge to Fledge communication however, there is no reason to prevent other applications that wish to send data into a Fledge 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 Fledge 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: ☒

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 Fledge 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 Fledge .

- 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 Fledge.

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
    }
  }
]
```

17.1.13 INA219 Voltage & Current Sensor



The *fledge-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*

The screenshot shows a configuration window for the 'ina219' plugin. At the top, there is a progress bar with three steps: 1. Plugin & Service Name, 2. Review Configuration (currently active), and 3. Done. Below the progress bar, there is a form with three input fields: 'Asset Name' with the value 'electrical', 'I2C Address' with the value '64', and 'Voltage Range' with the value '32V2A'. At the bottom of the form, there are two buttons: 'Previous' and '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

17.1.14 Lathe Simulation

The *fledge-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.

- **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.

17.1.15 Modbus South Plugin

The *fledge-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 Fledge. 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 Fledge,

- 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 Fledge 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 Fledge. 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*.

Set Point Control

The *fledge-south-modbus-c* plugin supports the Fledge 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 agisnt 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 rad 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 floatign point values.

17.1.16 South MQTT

The *fledge-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

The screenshot shows the Fledge configuration interface for the MQTT plugin. The interface is divided into three steps: 1. Plugin & Service Name, 2. Review Configuration, and 3. Done. The 'Review Configuration' step is active. The configuration fields are: MQTT Broker host (localhost), MQTT Broker Port (1883), Keep Alive Interval (60), Topic To Subscribe (Room1/conditions), QoS Level (0), and Asset Name (mqtt-). There are 'Previous' and 'Next' buttons at the bottom.

- **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

17.1.17 MQTT Sparkplug B

The *fledge-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 Fledge 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*

The screenshot shows a configuration wizard for the MQTT Sparkplug B plugin. At the top, there are three numbered steps: 1. Plugin & Service Name, 2. Review Configuration (which is the current step), and 3. Done. Below the steps, there is a form with the following fields:

- Asset Name:** mqtt
- MQTT Host:** chariot.groov.com
- MQTT Port:** 1883
- Username:** opto
- Password:** masked with dots
- Topic:** spBv1.0/Opto22/DDATA/groovEPIC_workshop/Strategy

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

17.1.18 OPC/UA South Plugin

The *fledge-south-opcua* plugin allows Fledge 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 Fledge.

- 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 ☒

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 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 Fledge. 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.
- **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.

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 Fledge. Although 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 Fledge 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 Fledge.

Object names, variable names and NamespaceIndexes can be easily retrieved browsing the given OPC/UA server using OPC UA clients, such as .

17.1.19 Person Detection Plugin

The *fledge-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.

The screenshot shows the configuration interface for the 'rtrr South Service'. The fields are as follows:

- TFlite Model File:** detect.tflite
- Labels File:** coco_labels.txt
- Asset Name:** person_detection_4
- Enable Edge TPU:** ☐
- 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:** ☐

- **TFlite Model File:** This is the name of the tflite model file that should be placed in `python/fledge/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
```

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_fledge_is_running:the_port_for_web_streaming*

17.1.20 Playback Plugin

The *fledge-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*

1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name: sample

CSV file name with extension: some.csv

Header Row: ☒

Header columns: None

Cherry pick column with same/new name: 1, {}

Historic timestamps: ☐

Pick timestamp delta from file: ☐

Timestamp column name: ts

Timestamp format: %Y-%m-%d %H:%M:%S.%f

Ingest mode: batch

Sample Rate: 100

Burst Interval (ms): 1000

Burst size: 1

Read file in a loop: ☐

Previous Next

- 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 fledge 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 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*.

17.1.21 PT100 Temperature Sensor

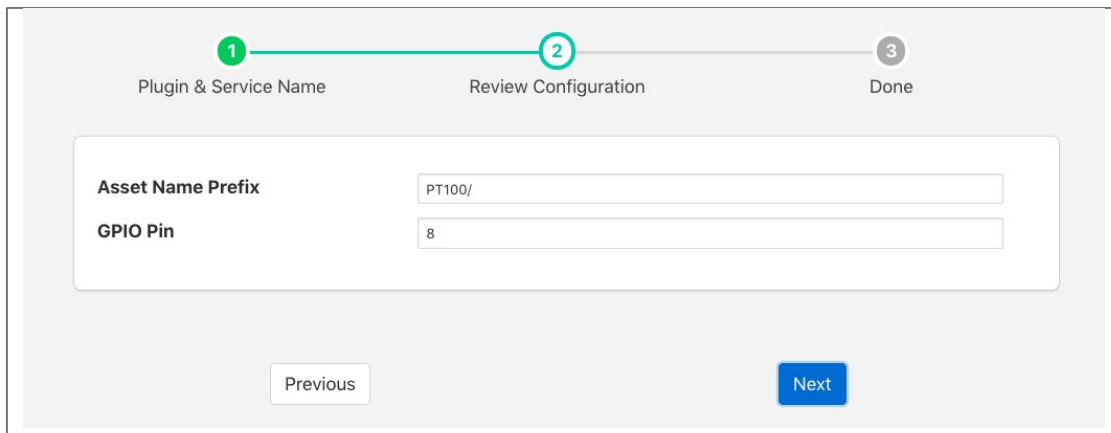


The *fledge-south-pt100* is a south plugin for the PT-100 temperature sensor. The PT100 is a resistance temperature detectors (RTDs) consist of a fine wire (typically platinum) wrapped around a ceramic core, exhibiting a linear increase in resistance as temperature rises. The sensor connects via a MAX31865 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*



1 Plugin & Service Name 2 Review Configuration 3 Done

Asset Name Prefix PT100/

GPIO Pin 8

Previous 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.

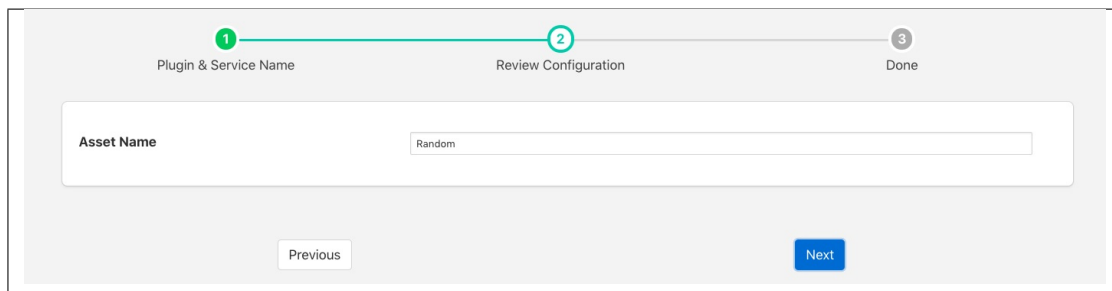
Then create a new connection between the 4 and 3 side of this jumper block. This is probably best done with a solder bridge.

17.1.22 Random

The *fledge-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 three-step configuration process. Step 1, 'Plugin & Service Name', is active and highlighted with a green circle. It contains a text input field labeled 'Asset Name' with the value 'Random' entered. Step 2, 'Review Configuration', is indicated by a green circle and a line. Step 3, 'Done', is indicated by a grey circle. At the bottom, there are 'Previous' and 'Next' buttons. The 'Next' button is blue and highlighted.

- Configure the plugin
 - **Asset name:** The name of the asset that will be created
- Click *Next*
- Enable the service and click on *Done*

17.1.23 Random Walk

The *fledge-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*

The configuration form is titled 'randomwalk'. It includes the following fields and values:

Field	Value
Asset name	randomwalk
Minimum Value	10
Maximum Value	100

Navigation buttons: 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*

17.1.24 OPC/UA Safe & Secure South Plugin

The *fledge-south-s2opcua* plugin allows Fledge 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 *fledge-south-opcua* plugin but is implemented using a different underlying OPC/UA open source library, from Syserel. 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 Fledge.

- 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 Plugin & Service Name 2 Review Configuration 3 Done

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: cacert

Server public key: OPCUAServer

Client public key: clientcert

Client private key: clientkey

Certificate revocation list: cacrl

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 Fledge. 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 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 key:** The name of the public key of the OPC/UA server specified in the *OPCUA Server URL*. This should be a DER format certificate signed by the certificate authority.
- **Client public key:** The name of the public key of the client application, i.e. the key to use for this plugin. This should be a DER format certificate signed by the certificate authority.
- **Client private key:** The name of the private key of the client application, i.e. the private key the plugin will use. This should be a PEM format key.
- **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.

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

The configuration described above uses the names of certificates that will be used by the plugin, these certificates must be loaded into the Fledge certificate store as a manual process and named to match the names used in the configuration before the plugin is started.

Typically the certificate authority certificate is retrieved and uploaded to the 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 Fledge certificate store.

Openssl may be used to generate and convert the keys and certificates required, and to do this is available as part of the underlying library.

17.1.25 SenseHAT



The *fledge-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 Fledge 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	<input type="text" value="sensehat/"/>
Pressure Sensor	<input checked="" type="checkbox"/>
Pressure Sensor Name	<input type="text" value="pressure"/>
Temperature Sensor	<input checked="" type="checkbox"/>
Temperature Sensor Name	<input type="text" value="temperature"/>
Humidity Sensor	<input checked="" type="checkbox"/>
Humidity Sensor Name	<input type="text" value="humidity"/>
Gyroscope Sensor	<input checked="" type="checkbox"/>
Gyroscope Sensor Name	<input type="text" value="gyroscope"/>
Accelerometer Sensor	<input checked="" type="checkbox"/>
Accelerometer Sensor Name	<input type="text" value="accelerometer"/>
Magnetometer Sensor	<input checked="" type="checkbox"/>
Magnetometer Sensor Name	<input type="text" value="magnetometer"/>
Joystick Sensor	<input checked="" type="checkbox"/>
Joystick Sensor Name	<input type="text" value="joystick"/>

- 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*

17.1.26 Sinusoid

The *fledge-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 Fledge system.





























The image shows a three-step configuration wizard. Step 1 is 'Plugin & Service Name', Step 2 is 'Review Configuration' (highlighted with a green circle), and Step 3 is 'Done'. In the 'Review Configuration' step, there is a text input field labeled 'Asset name' containing the value 'sinusoid'. At the bottom, there are two buttons: 'Previous' and 'Next'.

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.

The image shows the 'Sine South Service' configuration window. It has a title bar with a close button. The configuration is organized into sections. The 'Asset name' section contains a text input field with 'sinusoid' and a link 'Hide Advanced Config'. The 'Maximum Reading Latency (mS)' section has a text input field with '5000'. The 'Maximum buffered Readings' section has a text input field with '100'. The 'Reading Rate' section has a spin box with '10'. The 'Throttle' section has a checkbox that is unchecked. The 'Reading Rate Per' section has a dropdown menu with 'second' selected. The 'Minimum Log Level' section has a dropdown menu with 'warning' selected. The 'Enabled' section has a checkbox that is checked. At the bottom, there is a section for 'Applications' with a plus icon. At the very bottom right, there are 'Cancel' and 'Save' buttons.

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.

17.1.27 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 *fledge-south-systeminfo* plugin implements a that collects data about the machine that the Fledge 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*

17.1.28 Advantech USB-4704



The *fledge-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: usb4704

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"

```

Previous Next

- 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*

17.2 Fledge North Plugins

17.2.1 OMF

The *OMF* north plugin is included in all distributions of the Fledge core and provides the north bound interface to the OSIsoft data historians in all it 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 Fledge 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
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	<input type="checkbox"/>
Asset Framework hierarchies tree	/fledge/data_piwebapi/default
Asset Framework hierarchies rules	<div><div>1</div><div>{ }</div></div>
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

Select PI Web API from the Endpoint options.

- **Basic Information**

- **Endpoint:** Select what you wish to connect to, in this case PI Web API.
- **Send full structure:** Used to control if AF 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 within the PI Server. See [Naming Scheme](#).
- **Server hostname:** The hostname or address of 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. The readings or Fledge'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 Fledge server.

- **Asset Framework**

- **Asset Framework Hierarchies Tree:** The location in the Asset Framework into which the data will be inserted. All data will be inserted at this point in the Asset Framework unless a later rule overrides this.
- **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 equates to no authentication, basic authentication requires a user name and password and Kerberos allows integration with your single sign on environment.
- **PI Web API User Id:** The user name to authenticate with the PI Web API.
- **PI Web API Password:** 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 (Fledge 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 Fledge will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match Fledge 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 Fledge 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.

EDS OMF Endpoint

To use the OSISoft Edge Data Store first install Edge Data Store on the same machine as your Fledge 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 Fledge 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
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	<input type="checkbox"/>
Asset Framework hierarchies tree	/fledge/data_piwebapi/default
Asset Framework hierarchies rules	<div><div>1</div><div>{ }</div></div>
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
17.2. Fledge North Plugins	
OCS Namespace	name_space
OCS Tenant ID	ocs_tenant_id

Select Edge Data Store from the Endpoint options.

- **Basic Information**

- **Endpoint:** Select what you wish to connect to, in this case 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:** The hostname or address of the PI Server. This must be the localhost for EDS.
- **Server port:** The port the Edge Datastore is listening on. Leave as 0 if you are using the default port.
- **Data Source:** Defines which data is sent to the PI Server. The readings or Fledge’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 Fledge 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 (Fledge 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 Fledge will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match Fledge 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 Fledge data types to the data type configured in PI. The defaults 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.

OCS OMF Endpoint

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

omf ✕

Endpoint	OSIsoft Cloud Services ▼		
Server hostname	pi-server		
Server port, 0=use the default	0		
Producer Token	uid=5ced49c3-3a55-40e7-983f-c6cdcd5c5fd1&crt=20180620084'		
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	<input type="checkbox"/>		
Asset Framework hierarchies tree	/fledge/data_piwebapi/default		
Asset Framework hierarchies rules	<table><tr><td>1</td><td>{ }</td></tr></table>	1	{ }
1	{ }		
PI Web API Authentication Method	anonymous ▼		
PI Web API User Id	user_id		
PI Web API Password		
PI Web API Kerberos	piwebapi_kerberos_https.keytab		

Select OSIsoft Cloud Services from the Endpoint options.

- **Basic Information**

- **Endpoint:** Select what you wish to connect to, in this case 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 the PI Server. The readings or Fledge'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 Fledge server.

- **Authentication**

- **OCS Namespace:** Your namespace within the 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 OCS 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 (Fledge 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 Fledge will time out an HTTP connection attempt.

- **Other (Rarely changed)**

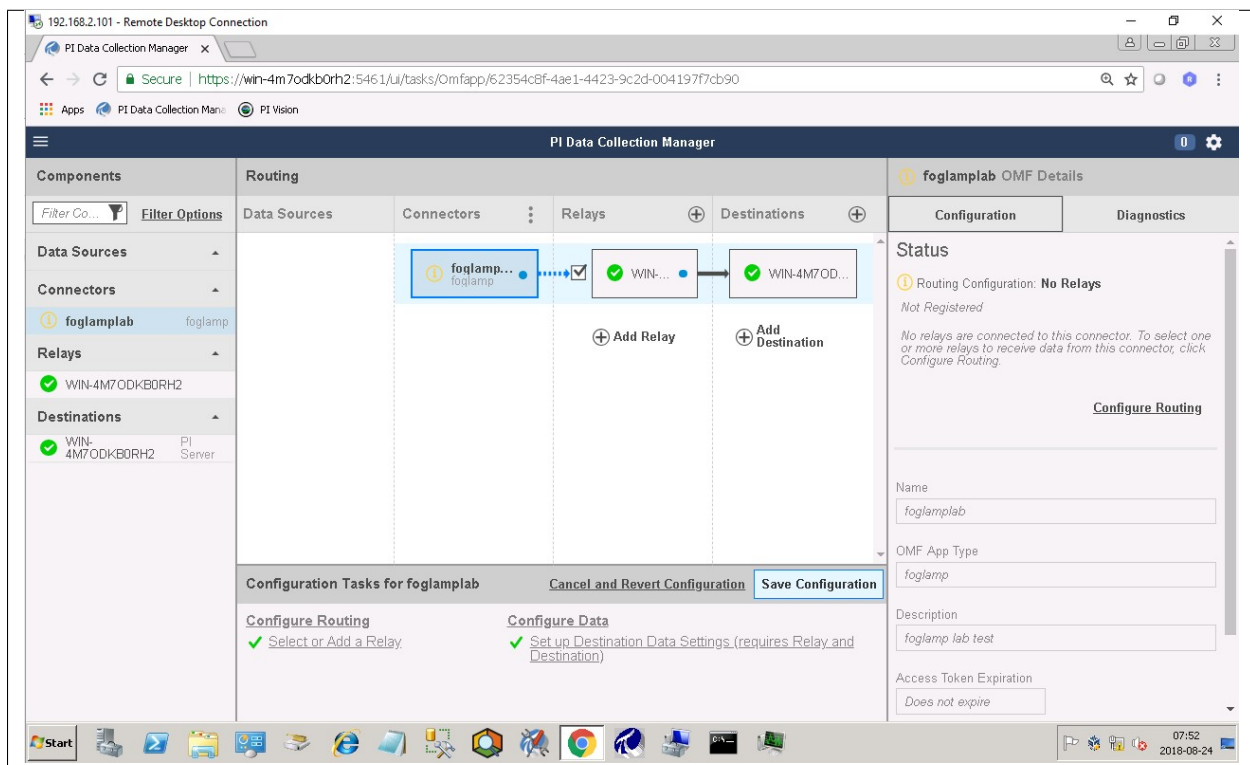
- **Integer Format:** Used to match Fledge data types to the data type configured in PI. This defaults to int64 but may be set to any OMF data type compatible with integer data, e.g. int32.
- **Number Format:** Used to match Fledge data types to the data type configured in PI. The default is float64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

PI Connector Relay


The PI Connector Relay was the original mechanism by which OMF data could be ingested into a PI Server, this has since been replaced by the PI Web API OMF endpoint. It is recommended that all new deployments should use the PI Web API endpoint as the Connector Relay has now been discontinued by OSIsoft. To use the Connector Relay, open and sign into the PI Relay Data Connection Manager.



To add a new connector for the Fledge 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 OMF Connector Relay by selecting the new Fledge application, clicking the check box for the OMF Connector Relay and then clicking “Save Configuration”.



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

Now go to the Fledge 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		
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	<input type="checkbox"/>		
Asset Framework hierarchies tree	/fledge/data_piwebapi/default		
Asset Framework hierarchies rules	<table><tr><td>1</td><td>{ }</td></tr></table>	1	{ }
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		

- **Basic Information**

- **Endpoint:** Select what you wish to connect to, in this case the Connector Relay.
- **Server hostname:** The hostname or address of the Connector Relay.
- **Server port:** The port the Connector Relay is listening on. Leave as 0 if you are using the default port.
- **Producer Token:** The Producer Token provided by PI
- **Data Source:** Defines which data is sent to the PI Server. The readings or Fledge'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 Fledge 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 (Fledge 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 Fledge will time out an HTTP connection attempt.

- **Other (Rarely changed)**

- **Integer Format:** Used to match Fledge 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 Fledge 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, in your environment you have objects that are liable to change, i.e. the types of attributes change or the number of attributes change between readings, then you may wish to take a different naming approach to if they do not.

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 for the AF Objects. OMF lets a type be defined, but once defined it can not be changed. A new type must be created rather than changing the existing type. This means a new asset framework object 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 Fledge. 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 data points within each reading within Fledge. 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 the objects in the AF framework and Attributes in the PI server. However if the structure of an asset changes a new AF Object will be created which will have the suffix `-type*x*` appended to it.

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

Use Attribute Hash 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 Fledge, both type suffices and attribute hashes will be applied to the naming.

Asset Framework Hierarchy Rules

The asset framework rules allow the location of specific assets within the PI Asset Framework to be controlled. There are two basic type 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"
      }
    }
  }
}
```

The name type rules are simply a set of asset name and AF 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 data point having a given value.
- **notequal:** This test looks for a name data point 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 data point called *temperature* in will be stored in the AF hierarchy *temperatures*, if the asset had a data point called *power* the asset will be placed in the AF hierarchy */Electrical/Power*.

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

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

In this case if the asset has a data point 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 data points within the reading. For example of the reading has a data point called *room* that contains the room in which the readings was taken, an AF location of */BuildingA/\${room}* would put the reading in the asset framework using the value of the room data point. 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 data point is missing. Therefore in our example above if we used the location `/BuildingA/${room:unknown}` a reading without a `room` data point would be place in `/BuildingA/unknown`. If no default is given and the data point is missing then the level in the hierarchy is ignore. E.g. if we use our original location `/BuildingA/${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 name data point 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 insert data into the PI Server as. 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 insert data into the PI Server as. 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 it's value, this 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 that asset within the Asset Framework. An asset framework hint would be as follow

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

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 into the processing of the data, however a specific plugin is available for adding the hints, the .

17.2.2 Google Cloud Platform North Plugin

The *fledge-north-gcp* plugin provide connectivity from a Fledge 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 Fledge 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 Fledge certificate store

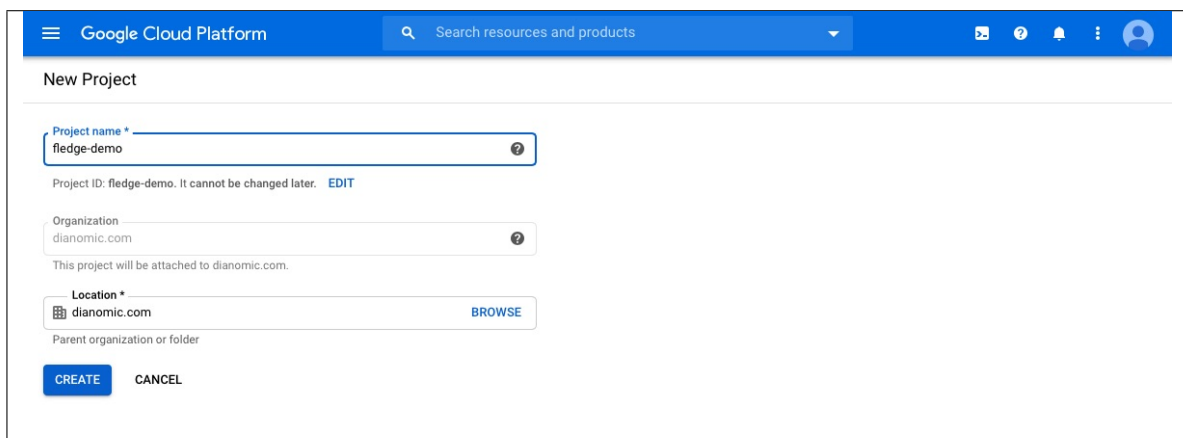
Create GCP IoT Core Project

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

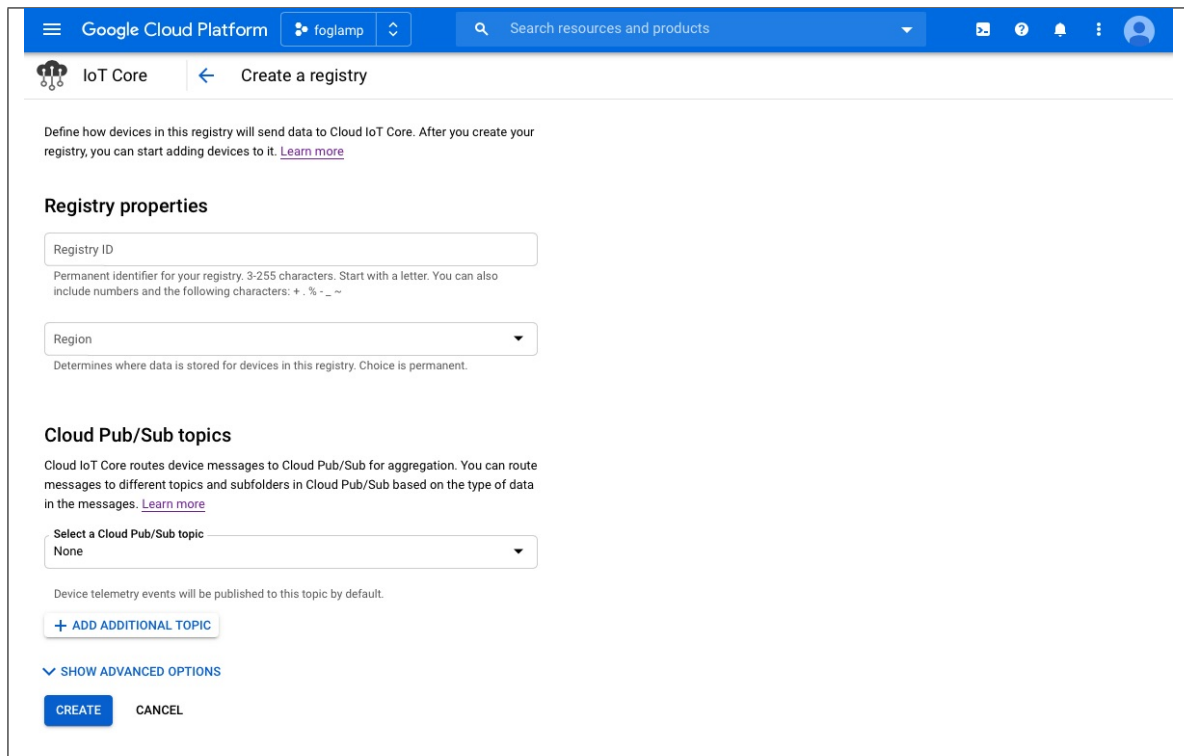
To create a registry in your project

- 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_fledge.pem -pkeyopt rsa_keygen_bits:2048
openssl rsa -in rsa_fledge.pem -pubout -out rsa_fledge.pem
```

- Goto the
- In the left pane of the IoT Core page in the Cloud Console, click Devices

Google Cloud Platform | foglamp | Search resources and products

IoT Core

- Registry details
- Devices**
- Gateways
- Monitoring

Devices | + CREATE A DEVICE | DELETE

Registry ID: flreg1
us-central1

Devices are things that connect to the internet directly or through a gateway. [Learn more](#)

Enter exact device ID

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

[Cloud IoT Core documentation](#)

- At the top of the Devices page, click *Create a device*

Google Cloud Platform | foglamp | Search resources and products

IoT Core

- Registry details
- Devices**
- Gateways
- Monitoring

← Create a device

Device ID

Permanent identifier for your device. 3-255 characters. Start with a letter. You can also include numbers and the following characters: + . % - _ ~

Device metadata (optional)

You can set custom metadata, such as manufacturer, location, etc. for the device. These can be used to query devices in this registry. [Learn more](#)

+ ADD ATTRIBUTE

✓ COMMUNICATION, STACKDRIVER LOGGING, AUTHENTICATION

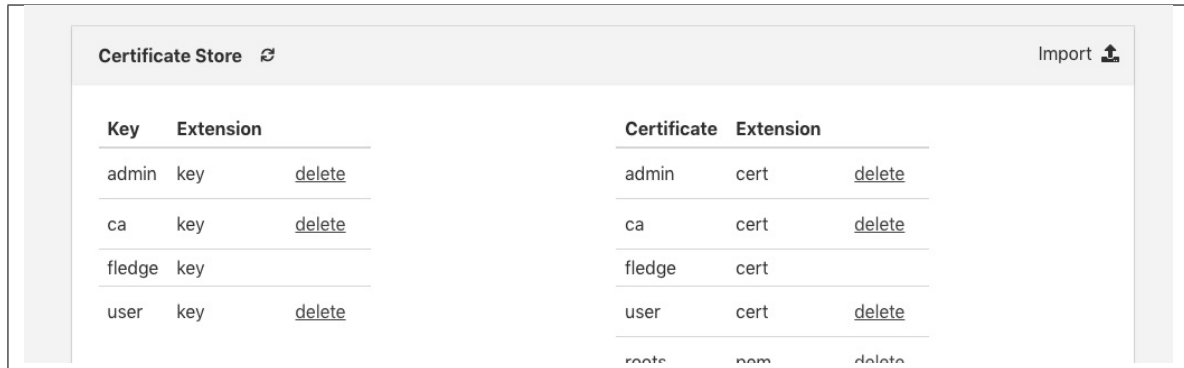
CREATE CANCEL

- 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.

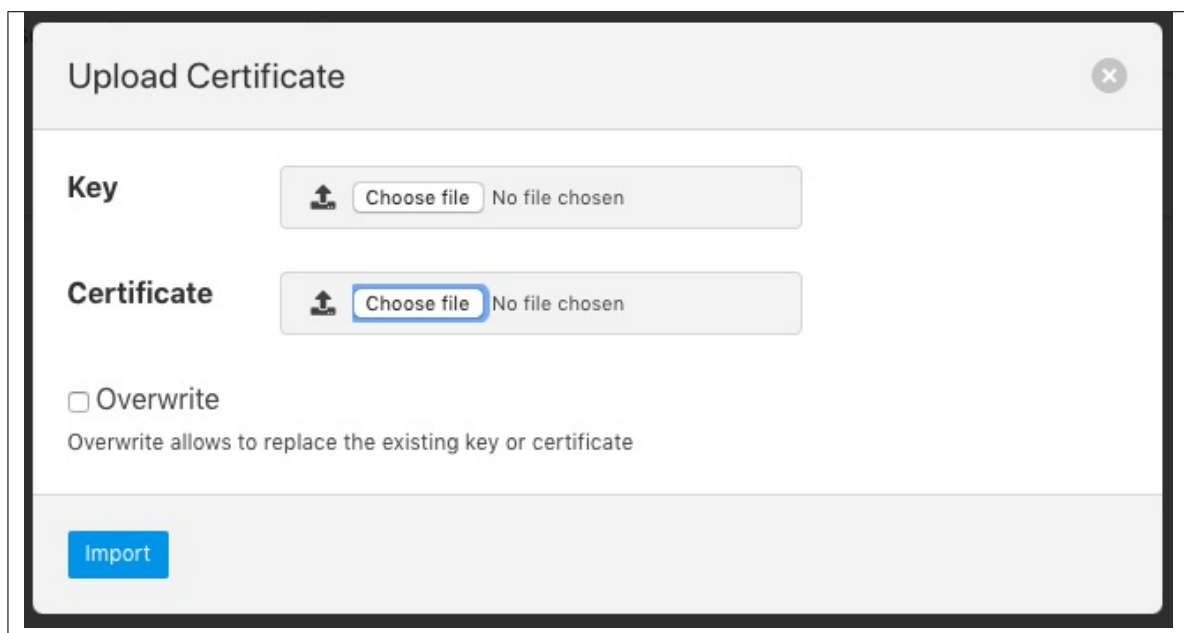
Upload Your Certificates

You should upload your certificates to Fledge

- From the Fledge user interface select the *Certificate Store* from the left-hand menu bar



- Click on the Import option in the top left corner



- 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 the 'Review Configuration' step of the Fledge plugin setup. The progress bar at the top indicates the current step is 2 of 3. The configuration form is as follows:

Field	Value
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 are 'Back' and 'Next' buttons.

- 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 Fledge statistics
- Click on *Next*
- Enable your plugin and click on *Done*

17.2.3 North HTTP

The *fledge-north-http* plugin allows data to be sent from the north of one Fledge instance into the south of another Fledge instance. It allows hierarchies of Fledge instances to be built. The Fledge to which the data is sent must run the corresponding in order for data to flow between the two Fledge instances. The plugin supports both HTTP and HTTPS transport protocols and sends a JSON payload of reading data in the internal Fledge format.

The plugin may also be used to send data from Fledge 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 Fledge.

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 Fledge 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 Fledge configuration process. A progress bar at the top indicates the sequence: 1. Plugin & Name, 2. Review Configuration (current step), and 3. Done. The configuration form contains the following fields:

- URL:** `http://localhost:6683/sensor-reading`
- Source:** `readings` (selected from a dropdown)
- Verify SSL:** ☐
- Apply Filter:** ☐
- Filter Rule:** `.[]`

At the bottom of the form are two buttons: 'Back' and 'Next'.

- **URL:** The URL of the receiving , the address and port should match the service in the up stream Fledge. 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 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.
- **Apply Filter:** This allows a simple jq format filter rule to be applied to the connection. This should not be confused with Fledge filters and exists for backward compatibility reason only.
- **Filter Rule:** A jq filter rule to apply. Since the introduction of Fledge 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 Fledge.

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
    }
  }
]
```

17.2.4 North HTTP-C

The *fledge-north-http-c* plugin allows data to be sent from the north of one Fledge instance into the south of another Fledge instance. It allows hierarchies of Fledge instances to be built. The Fledge to which the data is sent must run the corresponding in order for data to flow between the two Fledge instances. The plugin supports both HTTP and HTTPS transport protocols and sends a JSON payload of reading data in the internal Fledge 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 Fledge 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 Fledge.

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 Fledge 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 displays the 'Review Configuration' step of the Fledge plugin configuration process. The interface is divided into a left sidebar with labels for 'URL', 'Secondary URL', 'Source', and 'Headers', and a main configuration area on the right. The 'URL' field is populated with 'http://localhost:6683/sensor-reading'. The 'Secondary URL' field is empty. The 'Source' dropdown menu is set to 'readings'. The 'Headers' section shows a single header with a value of '{}'. Below these fields, there are input fields for 'Sleep Time Retry' (1), 'Maximum Retry' (3), and 'Http Timeout (in seconds)' (10). There are also checkboxes for 'Verify SSL' and 'Apply Filter', both of which are currently unchecked. A 'Filter Rule' field at the bottom is empty. At the top of the configuration window, a progress bar indicates the current step is '2 Review Configuration', with '1 Plugin & Name' and '3 Done' as previous and next steps respectively. At the bottom of the window, there are 'Back' and 'Next' buttons.

- **URL:** The URL of the receiving , the address and port should match the service in the up stream Fledge. 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
 - **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.
 - **Sleep Time Retry:** A tuning parameter used to control how often a connection is retried to the up stream Fledge 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 Fledge. 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.
 - **Apply Filter:** This allows a simple jq format filter rule to be applied to the connection. This should not be confused with Fledge filters and exists for backward compatibility reason only.
 - **Filter Rule:** A jq filter rule to apply. Since the introduction of Fledge filters in the north task this has become deprecated and should not be used.
- 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
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 Fledge.

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
    }
  }
]
```

17.2.5 Kafka Producer

The *fledge-north-kafka* plugin sends data from Fledge to the an Apache Kafka. Fledge 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 Fledge data to send to Kafka; Readings or Fledge Statistics.

17.2.6 OPCUA Server

The *fledge-north-opcua* plugin is a rather unusual north plugin as it does not send data to a system, but rather acts as a server from which other systems can pull data from Fledge. This is slightly at odds with the concept of short running tasks for sending north and does require a little more configuration when creating the North OPCUA server.

The process of creating a North OPCUA Server start 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 *opcua* option.

The screenshot shows a configuration wizard with three steps: 1. Plugin & Name, 2. Review Configuration, and 3. Done. In the 'Plugin & Name' step, the 'North Plugin' dropdown menu is open, displaying a list of available plugins: ocs_v2, OMF, opcua (highlighted), pi_server, and pi_server_v2. To the right of the dropdown, the text 'OPCUA Server' is visible. Below the dropdown, there is a link that says 'available plugins'. The 'Name' field contains the text 'OPCUA Server'. The 'Repeat (Interval)' section has two input fields: the first contains '0' and the second contains '01:00:00'. At the bottom of the form, there are two buttons: 'Back' and 'Next'.

In addition to setting a name for this task it is recommended to run the OPCUA North as a service rather than a task. Running as a periodically restarted task will cause clients to be disconnected at regular intervals, when run as a service the disconnections do not occur. If run as a task set the *Repeat* interval to a higher value than the 30 second default as we will be later setting the maximum run time of the north task to a higher value. Once complete click on *Next* and move on to the configuration of the plugin itself.

1 Plugin & Name 2 Review Configuration 3 Done

Server Name: Fledge OPCUA

URL: opc.tcp://localhost:4840/fledge/server

URI: urn:/fledge.dianomic.com

Namespace: http://fledge.dianomic.com

Source: readings

Object Root:

Hierarchy:

1	{}
---	----

Back Next

This second page allows for the setting of the configuration within the OPCUA server.

- **Server Name:** The name the OPCUA server will report itself as to any client that connects to it.
- **URL:** The URL that any client application will use to connect to the OPCUA server. This should always start `opc.tcp://`
- **URI:** The URI you wish to associate to your data, this is part of the OPCUA specification and may be set to any option you wish or can be left as default.
- **Namespace:** This defines the namespace that you wish to use for your OPCUA objects. If you are not employing a client that does namespace checking this is best left as the default.
- **Source:** What data is being made available via this OPCUA server. You may chose to make the reading data available or the Fledge statistics
- **Object Root:** This item can be used to define a root within the OPCUA server under which all objects are stored. If left empty then the objects will be created under the root.
- **Hierarchy:** This allows you to define a hierarchy for the OPCUA objects that is based on the meta data within the readings. See below for the definition of hierarchies.

Once you have completed your configuration click *Next* to move to the final page and then enable your north task and click *Done*.

The only step left is to modify the duration for which the task runs. This can only be done **after** it has been run for the first time. Enter your *North* task list again and select the OPCUA North that you just created. This will show the configuration of your North task. Click on the *Show Advanced Config* option to display your advanced configuration.

OPCUA Server

Server Name

Fledge OPCUA

URL

opc.tcp://localhost:4840/fledge/server

URI

urn://fledge.dianomic.com

Namespace

http://fledge.dianomic.com

Source

readings

Duration

60

Readings Block Size

500

Sleep Interval

1

Enabled

☒

Exclusive

☒

Interval

0

01:00:00

Hide Advanced Config

Applications

The *Duration* option controls how long the north task will run before stopping. Each time it stops any client connected to the Fledge OPCUA server will be disconnected, in order to reduce the disconnect/reconnect volumes it is advisable to set this to a value greater than the 60 second default. In our example here we set the repeat interval to one hour, so ideally we should set the duration to an hour also such that there is no time when an OPCUA server is not running. *Duration* is set in seconds, so should be 3600 in our example.

Hierarchy Definition

The hierarchy definition is a JSON document that defines where in the object hierarchy data is placed. The placement is controlled by meta data attached to the readings.

Assuming that we attach meta data to each of the assets we read that give a plant name and building to each asset using the names *plant* and *building* on those assets. If we wanted to store all data for the same plant in a single location in the OPCUA object hierarchy and have each building under the plant, then we can define a hierarchy as follows

```
{
  "plant" :
```

(continues on next page)

(continued from previous page)

```

    {
        "building" : ""
    }
}

```

If we had the following 4 assets with the metadata as defined

```

{
    "asset_code" : "A",
    "plant"      : "Bolton",
    "building"   : "10"
    ....
}
{
    "asset_code" : "B",
    "plant"      : "Bolton",
    "building"   : "7"
    ....
}
{
    "asset_code" : "C",
    "plant"      : "Milan",
    "building"   : "A"
    ....
}
{
    "asset_code" : "D",
    "plant"      : "Milan",
    "building"   : "C"
    ....
}
{
    "asset_code" : "General",
    "plant"      : "Milan",
    ....
}

```

The data would be shown in the OPCUA server in the following structure

```

Bolton
    10
        A
    7
        B
Milan
    A
        C
    C
        D
    General

```

Any data that does not fit this structure will be stored at the root.

17.2.7 ThingSpeak

The *fledge-north-thingspeak* plugin provides a mechanism to , allowing an easy route to send data from an Fledge 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 shows the ThingSpeak web interface. At the top is a navigation bar with 'Channels', 'Apps', and 'Support' menus. The main heading is 'My Channels'. Below it is a 'New Channel' button and a search bar. A table lists the channels:

Name	Created	Updated
sinusoid	2018-08-08	2019-07-26 15:04

Below the table are buttons for 'Private', 'Public', 'Settings', 'Sharing', 'API Keys', and 'Data Import / Export'. On the right, the 'Help' section explains how to collect data and create channels. The 'Examples' section lists various hardware devices. The 'Upgrade' section offers options for faster data sending or commercial use.

- 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 Fledge.
- 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

ThingSpeak™

Channels ▾

Apps ▾

Support ▾

Commercial Use

How to Buy

MR

sinusoid

Channel ID: 556345 | Test channel

Author: markdianomic

Access: Private

Private View

Public View

Channel Settings

Sharing

API Keys

Data Import / Export

Write API Key

Key

XXXXXXXXXXXXX

Generate New Write API Key

Read API Keys

Key

XXXXXXXXXXXXX

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=XXXXXXXXXXXXX&field1=0
```

Read a Channel Feed

```
GET https://api.thingspeak.com/channels/556345/feeds.json?api_key=XXXXXXXXXXXXX&results=2
```

Read a Channel Field

```
GET https://api.thingspeak.com/channels/556345/fields/1.json?api_key=XXXXXXXXXXXXX&results=2
```

Read Channel Status Updates

```
GET https://api.thingspeak.com/channels/556345/status.json?api_key=XXXXXXXXXXXXX
```

[Learn More](#)

Once you have created your channel on you may create your north task on Fledge 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 displays the 'Review Configuration' step of the Fledge setup process. The configuration details are as follows:

- URL:** `https://api.thingspeak.com/channels`
- API Key:** (Empty field)
- Source:** `readings`
- Fields:** A JSON document:


```

1 {
2   "elements": [
3     {
4       "asset": "sinusoid",
5       "reading": "sinusoid"
6     }
7   ]
8 }
      
```
- Channel ID:** `0`

Navigation buttons at the bottom: `Back` and `Next`.

- **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 Fledge 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*

17.3 Fledge Filter Plugins

17.3.1 Asset Filter

The *fledge-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"
}

```

17.3.2 Change Filter

The *fledge-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: ☒

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

17.3.3 Delta Filter

The *fledge-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

The screenshot shows the 'Review Configuration' step of the Fledge configuration interface for the delta plugin. The progress bar at the top indicates two steps: '1 Plugin Name' and '2 Review Configuration'. The configuration form contains the following fields:

- Tolerance %:** A text input field with the value '0'.
- Minimum Rate:** A text input field with the value '0'.
- Minimum Rate Units:** A dropdown menu with 'per second' selected.
- Individual Tolerances:** A JSON editor showing a single entry: '1 {}'.
- Enabled:** A checkbox that is currently unchecked.

At the bottom of the configuration panel are 'Previous' and 'Done' buttons.

- 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.

17.3.4 Expression Filter

The *fledge-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 Fledge 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

The screenshot shows the configuration interface for the 'expression' plugin. At the top, a progress bar indicates two steps: '1 Plugin Name' and '2 Review Configuration'. The main configuration area contains three fields: 'Datapoint Name' with the value 'calculated', 'Expression to apply' with the value 'log(x)', and 'Enabled' which is a checked checkbox. At the bottom, there are two buttons: 'Previous' and 'Done'.

- 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 *fledge-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)
```

17.3.5 Fast Fourier Transform Filter

The *fledge-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 Fledge 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

17.3.6 Flir Validity Filter

The *fledge-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 Plugin Name 2 Review Configuration

Area Labels

Box	Label
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.

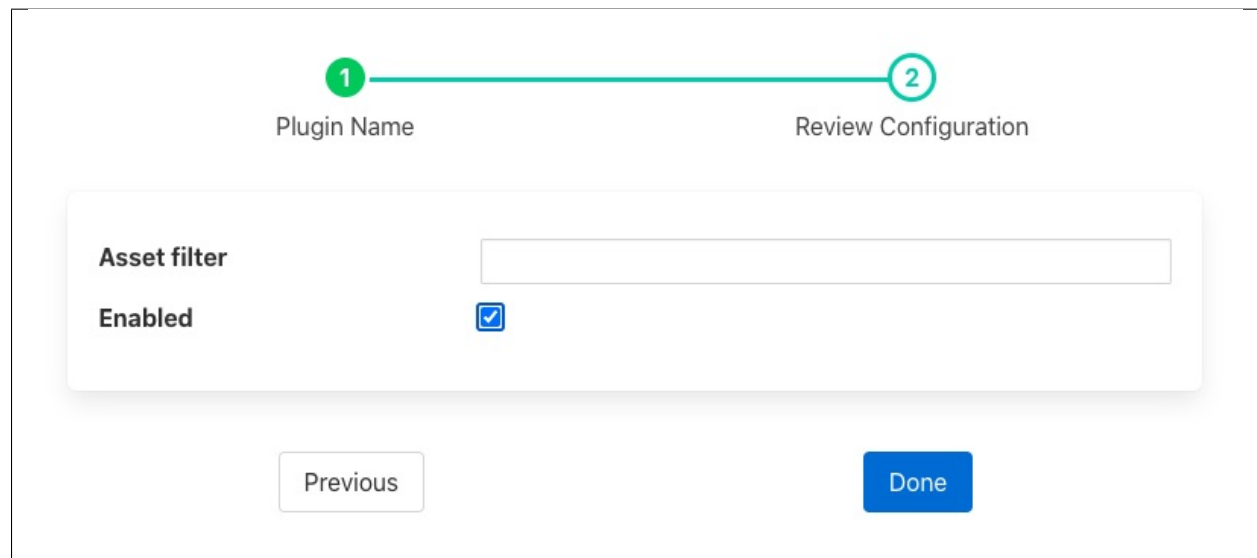
```
{
  "areas" : [
    "Motor",
    "Bearing",
    "Tool",
    "4",
    "5",
    "6",
    "7",
    "8",
    "9",
    "10"
  ]
}
```

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*.

17.3.7 Log Filter

The *fledge-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 configuration interface for the *fledge-filter-log* plugin is shown. It features a progress bar at the top with two steps: 1. Plugin Name and 2. Review Configuration. Below the progress bar, there is a configuration form with the following elements:

- Asset filter:** A text input field for entering a regular expression.
- Enabled:** A checkbox that is currently checked.
- Navigation:** Two buttons at the bottom, '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.

17.3.8 Metadata Filter

The *fledge-filter-metadata* filter allows data to be added to assets within Fledge. 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 Fledge is sending to the up stream system, in which case the metadata would probably related to that particular Fledge instance. Adding metadata in the north is particularly useful when a hierarchy of Fledge systems is used and an audit trail is required with the data or the individual Fledge 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 Fledge instance in each building to collect building data. We also install a Fledge instance in each site to collect the data from each individual Fledge instance per building, this allows us to then send the site data to the head office without having to allow each building Fledge to have access to the corporate network. Only the site Fledge 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 Fledge, for example the “Pearson” building, we add the following metadata

```
{
  "value" : {
    "building": "Pearson"
  }
}
```

Likewise to the “Lawrence” building Fledge 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 Fledge 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.

17.3.9 OMF Hint Filter

The *fledge-filter-omfhint* filter allows hints to be added to assets within Fledge 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 Fledge is sending to the up stream 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 Fledge that then forwards to OMF.

To add a 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.

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"
    }
  }
}
```

In this case all numeric data except for *frequency* will be stored as float32 and *frequency* will be stored as uint16.

17.3.10 Python 2.7 Filter

The *fledge-filter-python27* filter allows snippets of Python to be easily written that can be used as filters in Fledge. 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 Fledge 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

1
2

Plugin Name
Review Configuration

Python script

```

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}

```

ema.py

Choose Files
No file chosen

Configuration

```

1  {"rate" : 0.75}

```

Enabled
☐

Previous
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`.

17.3.11 Python 3.5 Filter

The *fledge-filter-python35* filter allows snippets of Python to be easily written that can be used as filters in Fledge. A similar filter exists that uses Python 2.7 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 Fledge 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
```

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

1 Plugin Name 2 Review Configuration

Python script

```

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}

```

ema.py

Choose Files No file chosen

Configuration

```

1 {"rate" : 0.75}

```

Enabled ☐

Previous 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*.

17.3.12 Rate Filter

The *fledge-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.

17.3.13 Rename Filter

The *fledge-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-b256-63370ccf5b",
  "ts": "2021-06-28 14:03:22.106562+00:00",
  "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"
}
```


17.3.14 Replace Filter

The *fledge-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 Replace Filter. At the top, a progress bar indicates two steps: 1. Plugin Name and 2. Review Configuration. The main configuration area contains three fields: 'Replace' with the value '[\:\?]', 'With' with the value '-', and an 'Enabled' checkbox which is currently unchecked. A help icon (?) is located to the right of the 'Replace' field. At the bottom, there are two buttons: 'Previous' and 'Done'.

- **Replace:** The set of reserved characters to be replaced.
- **With:** The character to replace each occurrence of the above characters with

17.3.15 Root Mean Squared (RMS) Filter

The *fledge-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.

17.3.16 Scale Filter

The *fledge-filter-scale* plugin is a simple filter that allows a scale factor and an offset to be applied to numerical data. 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.

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 window titled "Sine South Service". It features a progress bar at the top with two steps: "1 Plugin Name" and "2 Review Configuration". The "Review Configuration" step is currently active. Below the progress bar, there is a configuration area with the following fields:

- Scale Factor**: A text input field containing the value "100.0".
- Constant Offset**: A text input field containing the value "0.0".
- Asset filter**: An empty text input field.
- Enabled**: A checkbox that is checked.

At the bottom of the configuration area, there are two buttons: "Previous" and "Done".

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/

17.3.17 Scale Set Filter

The *fledge-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

1 Plugin Name 2 Review Configuration

Scale factors

```

1 {
2   "factors": [
3     {
4       "asset": ".*",
5       "datapoint": ".*",
6       "scale": 1,
7       "offset": 0
8     }
9   ]
10 }

```

Enabled ☐

Previous Done

- 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" : [
    {
      "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
    }
  ]
}
```

17.3.18 Threshold Filter

The *fledge-filter-threshold* plugin is a filter that is used to control the forwarding of data within Fledge. 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 Fledge 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

1 Plugin Name 2 Review Configuration

Expression speed > 10

Enabled ☒

Previous Done

- Enter the expression to control forwarding in the box labeled *Expression*
- Enable the filter and click on *Done* to activate it

Expressions

The *fledge-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.

17.4 Fledge Notification Rule Plugins

17.4.1 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.

The screenshot displays the Fledge Notification Rule configuration interface. At the top, a progress bar indicates four steps: 1. Notification Instance, 2. Rule (current step), 3. Delivery Channel, and 4. Done. Below the progress bar, a configuration form is shown with the following fields:

- Asset name:** A text input field.
- Datapoint name:** A text input field.
- Condition:** A dropdown menu with the value ">" selected.
- Trigger value:** A text input field containing "0.0".
- Evaluation data:** A dropdown menu with the value "Single Item" selected.
- Window evaluation:** A dropdown menu with the value "Average" selected.
- Time window:** A text input field containing "30".

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

17.4.2 Moving Average Rule

The *fledge-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

Deviation %

Direction

Average

EMA Factor

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 %

Direction

- Above Average
- Below Average
- ✓ Both

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

EMA Factor

- ✓ Simple Moving Average
- Exponential Moving Average

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.

17.4.3 Expression Rule

The *fledge-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.

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Rule Plugin

Simple-Expression

Threshold

Generate a notification based on the evaluation of a user provided expression

[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 name

Expression to apply

Previous Next

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)

17.5 Fledge Notification Delivery Plugins

17.5.1 Amazon Alexa Notification

The *fledge-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 displays the configuration screen for the Alexa notification plugin. At the top, a progress bar indicates the current step is '3. Delivery Channel'. The main form contains the following fields:

- Access Code:** An empty text input field.
- Title:** A text input field containing the example notification: "The level in tank 15 is below 10%".
- Enabled:** A checkbox that is currently checked.

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

- 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”.

17.5.2 Asset Notification

The *fledge-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 Fledge. 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.

17.5.3 Email Notifications

The *fledge-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 displays the 'Delivery Channel' configuration step for an email notification. The progress bar at the top indicates the current step is 3 of 4. The configuration form contains the following fields and values:

- 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 of the dialog, 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.

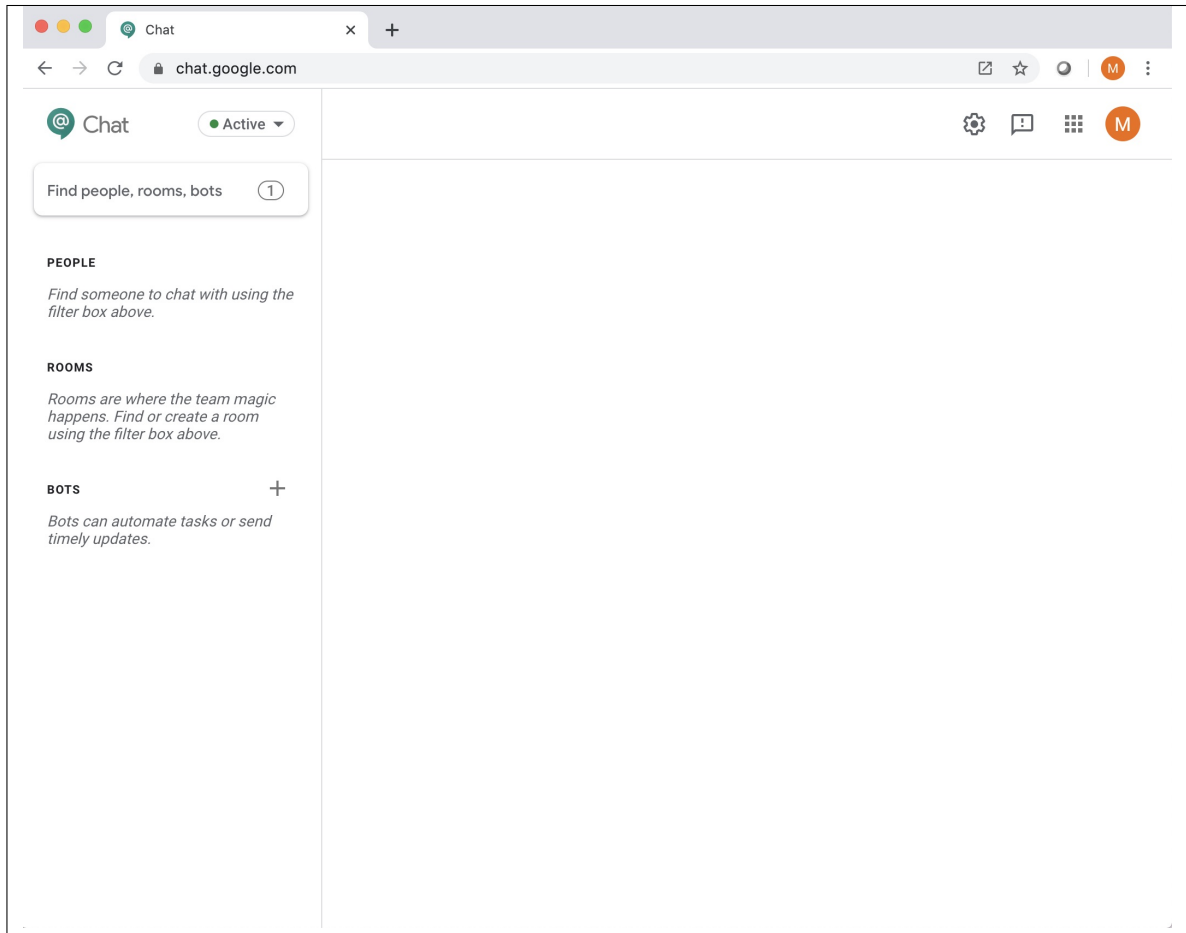
17.5.4 Google Chat

The *fledge-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

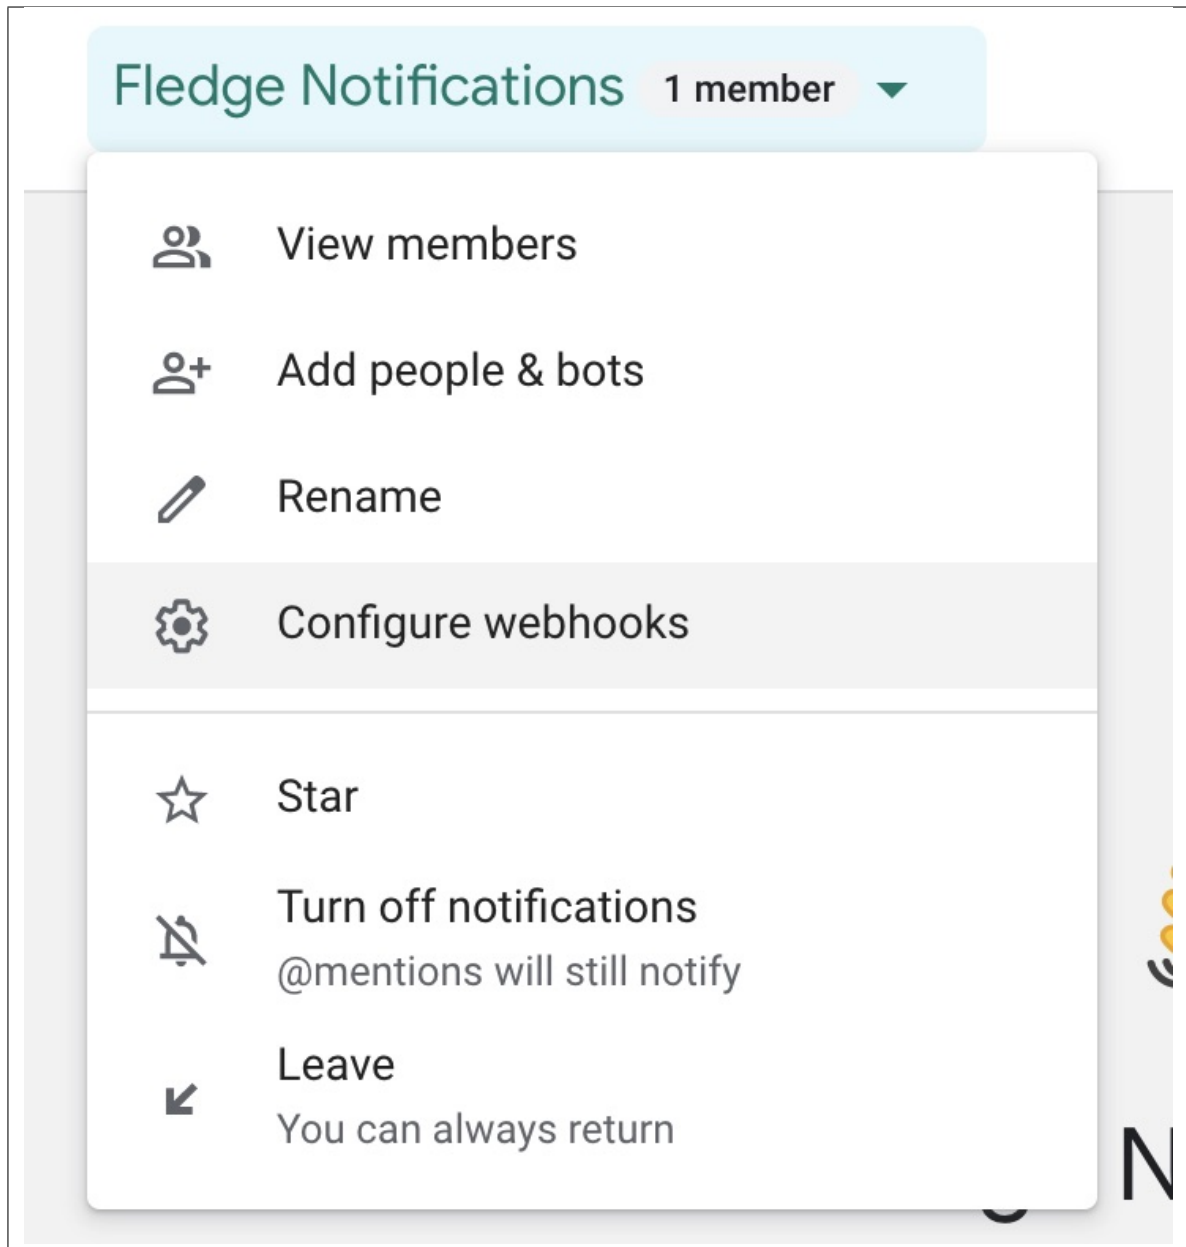
must create a webhook for sending data to that chatroom.

To create a webhook

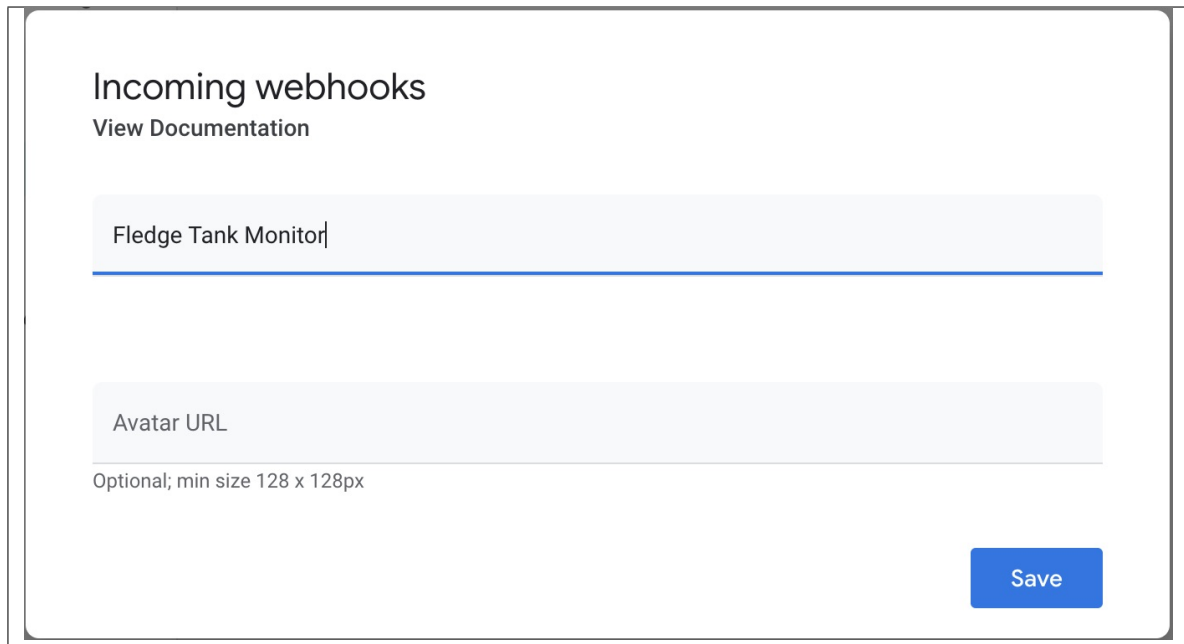
- Go to the page in your browser



- 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](#)

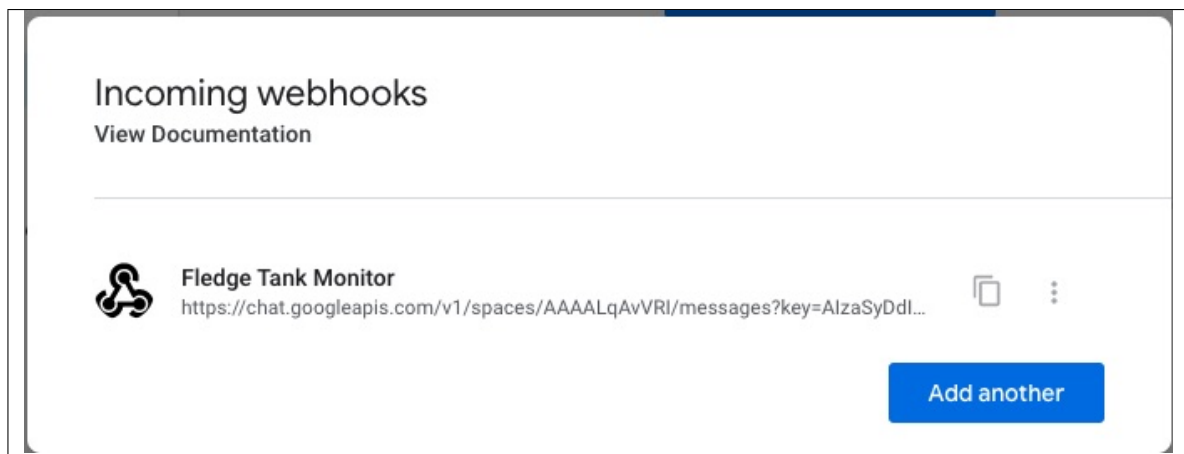
Fledge Tank Monitor

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...

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

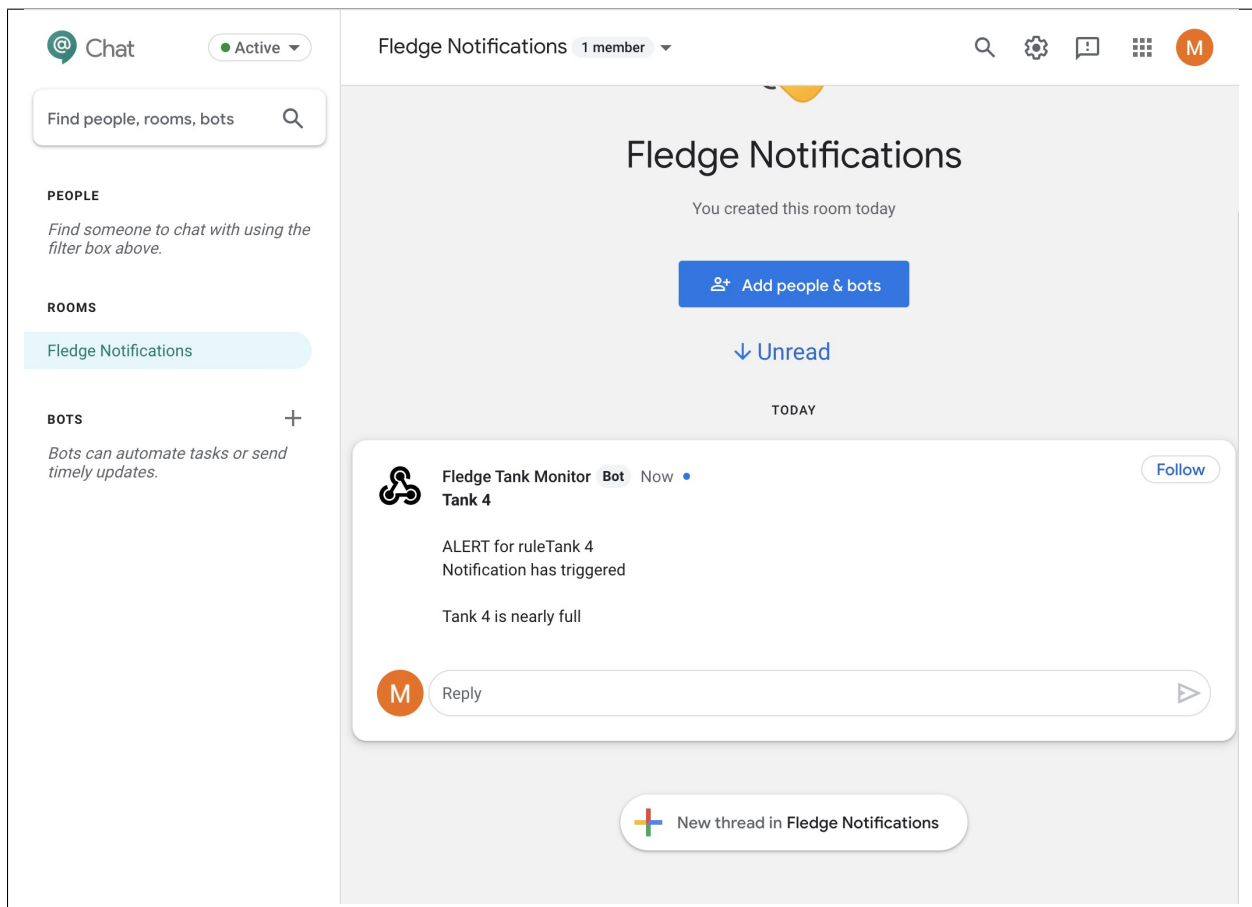
Message Text

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.




17.5.5 IFTTT Delivery Plugin

The *fledge-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 Fledge 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 Fledge 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 Fledge 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 shows the IFTTT plugin configuration page. At the top, there is a progress bar with four steps: 1. Notification Instance, 2. Rule, 3. Delivery Channel, and 4. Done. Step 3 is currently active. Below the progress bar, there are three fields: 'IFTTT Trigger' with the value 'button_press', 'IFTTT Key' with a long string of 'X's, and 'Enabled' with a checked checkbox. At the bottom, there are 'Previous' and 'Next' buttons.

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.

17.5.6 MQTT Notification

The *fledge-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*

The screenshot displays a four-step configuration process for a notification instance. The steps are: 1. Notification Instance, 2. Rule, 3. Delivery Channel (currently active), and 4. Done. The 'Delivery Channel' step contains a form with the following fields:

Field	Value
MQTT Broker	tcp://mqtt.local:1883
MQTT Topic	Fledge
Trigger Payload	Triggered
Clear Payload	Cleared
Enabled	<input checked="" type="checkbox"/>

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

17.5.7 Operation Notification

The *fledge-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

17.5.8 Python 3 Script

The *fledge-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*

Notification Instance Rule Delivery Channel Done

Python script

```
1 from time import sleep
2 from envirophat 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.

17.5.9 Set Point Control Notification

The *fledge-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*

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Service

Trigger Value

```
1 {
2   "values": {
3     "name": "value"
4   }
5 }
```

Cleared Value

```
1 {
2   "values": {
3     "name": "value"
4   }
5 }
```

Enabled ☐

- 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 *fledge-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, *fledge-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.

17.5.10 Slack Messages

The *fledge-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 Fledge 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 four-step configuration process: 1. Notification Instance, 2. Rule, 3. Delivery Channel (current step), and 4. Done. The 'Delivery Channel' step contains a form with the following fields:

- Slack Webhook URL:** A text input field containing the URL `https://hooks.slack.com/services/T2GBZ52AF/BGFNTP7NG/YxQwIjds5ZHMirFZqUjUc6z`.
- Message Text:** A text input field containing the message `The temperature of the coolant in pump4 is above normal operating levels`.
- Enabled:** A checkbox that is checked.


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

- Configure the delivery plugin
 - **Slack Webhook URL:** Paste the URL you obtain above from the page
 - **Message Test:** 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

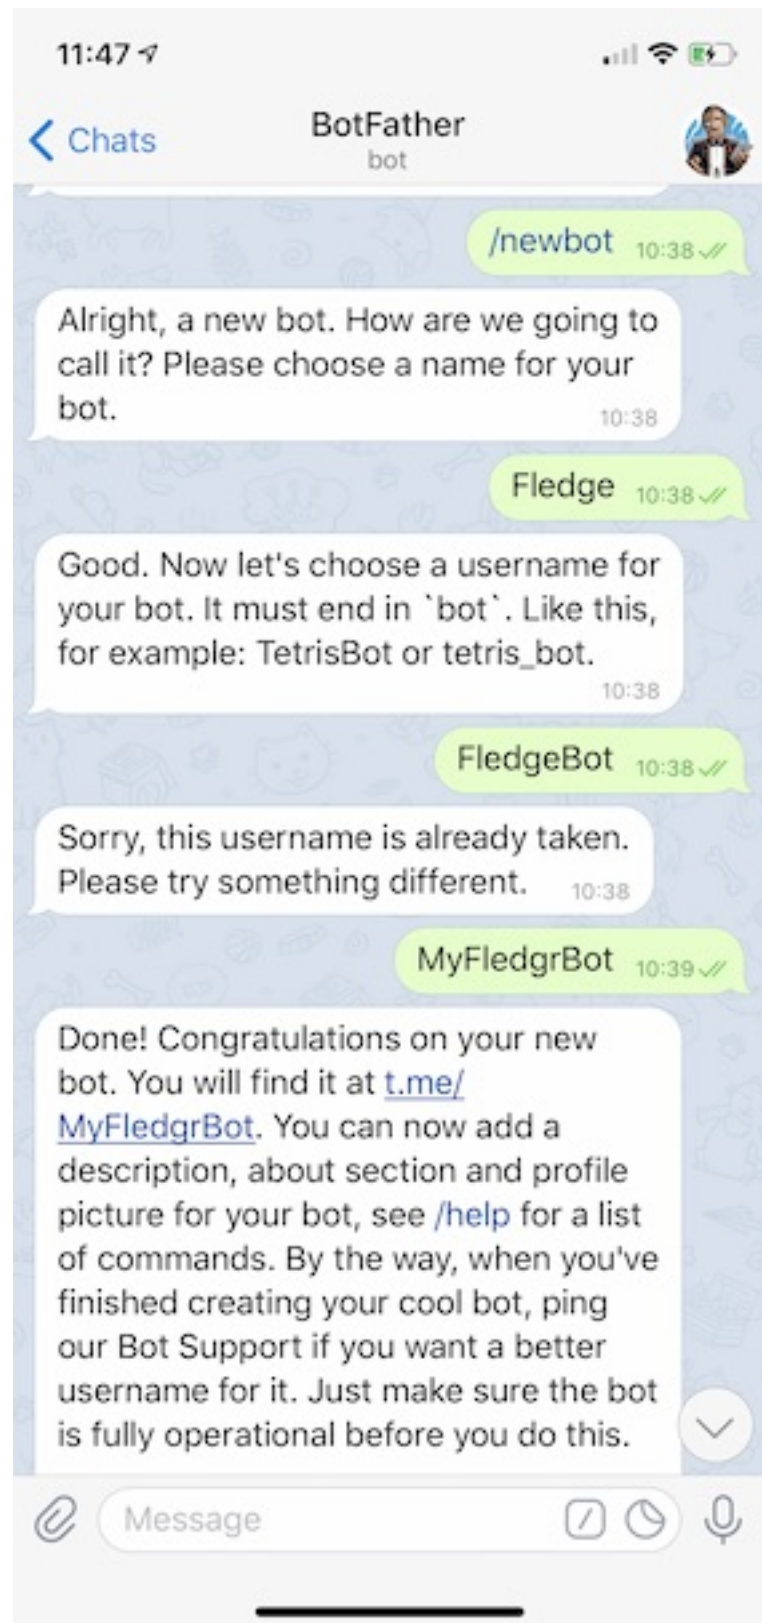


and/or your mobile devices



17.5.11 Telegram Messages

The *fledge-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*

1 Notification Instance 2 Rule 3 Delivery Channel 4 Done

Telegram BOT API token

Telegram user chat_id

Telegram BOT API url prefix

Enabled ☐

Previous 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

