SNMP Graphing with Telegraph, InfluxDB, and Graphana

Eli-v6.4.84

Table of Contents

1 Introduction

SNMP aids with network management by exposing information about networked devices to a manager, which can observe them remotely. If SNMP v3 is in use, observations will remain private and secure. The *Telegraf-InfluxDB-Grafana*, TIG, stack is a combination of 3 programs which allow you to set up automatic querying and graphing of SNMP devices.

This guide walks you through setting up a TIG stack which graphs information from a Bridgeworks unit's SNMP agent.

2 Definitions

TIG *Telegraf-InfluxDB-Grafana* stack - Abbreviation for these 3 programs used together.

SNMP Simple Network Management Protocol - Can be enabled on your unit so you can integrate it with your existing network manager. Read our SNMP guide for help with this.

Telegraf An agent that collects and exports data about your system. It has many input and output plugins that allow it to collect a variety of data inputs and output it to many external systems. This guide is limited to SNMP input and InfluxDB output plugins. Other systems are outside the scope of this guide.

InfluxDB A time series database used for storing information collected by Telegraf.

Grafana A program for visualising and monitoring time series data.

MIB Management Information Base - These are structured database entries that contain information about the hierarchical layout of the Objects in the database. Where applicable, they also contain specific information about how to describe the value the object links to.

OID Object Identifier - These are addresses to a specific entry in the hierarchical database tree. They are the literal address that is interpreted by MIB objects.

Enterprise MIBs MIBs reserved by individuals and companies to describe data they need in SNMP which isn't provided by default SNMP. For example Bridgeworks' Enterprise MIB describes accelerated traffic stats.

3 Setup walkthrough

This section will walk you through: installing all necessary tools, booting them up, checking that they were installed correctly, and then reconfiguring them to display a graph of SNMP data.

Before following this guide, you'll need to have a Bridgeworks unit running SNMP. Bridgeworks' SNMP guide can help you with this.

The system used for this guide was a Ubuntu Server 17.04 OS installation on a virtualised machine. The programs used are also available on other operating Systems but their installation will be slightly different. Online installation tutorials are available for other operating systems.

3.1 Install

The first step will be to install the necessary programs: *Telegraf*, *InfluxDB*, *Grafana*, *Net-SNMP*, and *SNMP-MIBs-downloader*.

3.1.1 InfluxDB

The first program we'll install and start is the database, *InfluxDB*.

First, add the influxdata key and repository with the commands below.

```
> sudo curl -sL https://repos.influxdata.com/influxdb.key | sudo apt-key add -
> source /etc/lsb-release
> echo "deb https://repos.influxdata.com/${DISTRIB_ID,,} ${DISTRIB_CODENAME}
stable" | sudo tee /etc/apt/sources.list.d/influxdb.list
```

Next install the packages using apt:

```
> sudo apt update
> sudo apt install influxdb -y
```

Lastly, start InfluxDB and check it's running correctly.

```
> sudo systemctl start influxdb
> sudo netstat -plntu
```

The output of the netstat command shows all the open listen sockets currently available on this system. The rightmost column is labelled *PID/Program Name*, this column shows the program that has opened the listen socket described in that row. If an entry in this column ends with `influxd`, InfluxDB is installed and running.

The only set up needed for InfluxDB is to create a database for Telegraf to write to, and a user for Grafana to use. This is done with:

```
> influx
> create database telegraf
> show databases
> create user telegraf with password 'telegraf'
> exit
```

You can use a different username and password, you'll need to remember them for when you add InfluxDB as a Grafana data source.

3.1.1.1 InfluxDB Expected output

```
> user@sys:~$ sudo curl -sL https://repos.influxdata.com/influxdb.key | sudo apt-key add -

OK
> user@sys:~$ source /etc/lsb-release
> user@sys:~$ echo "deb https://repos.influxdata.com/${DISTRIB_ID,,} ${DISTRIB_CODENAME}
stable" |

> sudo tee /etc/apt/sources.list.d/influxdb.list

deb https://repos.influxdata.com/ubuntu bionic stable
> user@sys:~$ sudo apt update

Hit:1 http://archive.ubuntu.com/ubuntu bionic InRelease
Get:2 http://archive.ubuntu.com/ubuntu bionic-security InRelease [88.7 kB]
Get:3 http://archive.ubuntu.com/ubuntu bionic-updates InRelease [88.7 kB]
Get:4 https://repos.influxdata.com/ubuntu bionic InRelease [4,731 B]
Get:5 https://repos.influxdata.com/ubuntu bionic/stable amd64 Packages [923 B]
Fetched 183 kB in 1s (123 kB/s)
Reading package lists... Done
Building dependency tree
Reading state information... Done
182 packages can be upgraded. Run 'apt list --upgradable' to see them.
> user@sys:~$ sudo apt install influxdb -y

Reading package lists... Done
Building dependency tree
Reading state information... Done
The following NEW packages will be installed:
influxdb
0 upgraded, 1 newly installed, 0 to remove and 182 not upgraded.
Need to get 46.7 MB of archives.
After this operation, 107 MB of additional disk space will be used.
Get:1 https://repos.influxdata.com/ubuntu bionic/stable amd64 influxdb amd64 1.7.4-1
[46.7 MB]
Fetched 46.7 MB in 5s (10.2 MB/s)
Selecting previously unselected package influxdb.
(Reading database ... 102196 files and directories currently installed.)
Preparing to unpack .../influxdb_1.7.4-1_amd64.deb ...
Unpacking influxdb (1.7.4-1) ...
Setting up influxdb (1.7.4-1) ...
```

```
Created symlink /etc/systemd/system/influxd.service → /lib/systemd/system/influxdb.service.
Created symlink /etc/systemd/system/multi-user.target.wants/influxdb.service → /lib/
systemd/system
/influxdb.service.
Processing triggers for man-db (2.8.3-2) ...
> user@sys:~$ sudo systemctl start influxdb
> user@sys:~$ sudo netstat -plntu

Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address      Foreign Address    State     PID/Program name
tcp         0      0 127.0.0.53:53      0.0.0.0:*          LISTEN    888/systemd-resolve
tcp         0      0 0.0.0.0:22         0.0.0.0:*          LISTEN    1353/sshd
tcp         0      0 127.0.0.1:8088     0.0.0.0:*          LISTEN    2811/influxd
tcp6        0      0 :::8086            :::*               LISTEN    2811/influxd
tcp6        0      0 :::22              :::*               LISTEN    1353/sshd
udp         0      0 127.0.0.53:53      0.0.0.0:*                    888/systemd-resolve
udp         0      0 10.10.72.27:68     0.0.0.0:*                    869/systemd-network
> user@sys:~$ influx

Connected to http://localhost:8086 version 1.7.4
InfluxDB shell version: 1.7.4
Enter an InfluxQL query
> create database telegraf
> create user telegraf with password 'telegraf'
> show databases

name: databases
name
----
_internal
telegraf
> exit

user@sys:~$
```

3.1.2   Net-SNMP

Telegraf uses Net-SNMP to get data from the unit. It also includes `snmpwalk` which is useful for checking your device can connect to the unit. To install Net-SNMP, use:

```
> sudo apt install snmp
```

To check it's installed correctly, run a snmpwalk to the unit. For instructions on how to use snmpwalk, see Net-SNMP's documentation (http://www.net-snmp.org/docs/man/snmpwalk.html).

3.1.3 SNMP-MIBs-downloader

SNMP stores it's data with an index system of OIDs. This is a large tree which means SNMP data is usually preceded by long number chains like 1.3.6.1.2.1.31.1.1.1.17.4. MIBs are used to associate these long OID chains with more easily human readable names. E.g. IF-MIB::ifConnectorPresent.4. Telegraf can use these when you specify which data to import from the SNMP agent.

The command below installs SNMP-MIBs-downloader.

```
> sudo apt install snmp-mibs-downloader
```

To query acceleration data, Bridgeworks' Enterprise MIBs will also be needed. These describe Bridgeworks specific data and can be downloaded from your unit's SNMP page. For more instruction see the Bridgeworks SNMP guide.

3.1.4 Telegraf

Telegraf gathers data and outputs to the InfluxDB database. Install and start it with:

```
> sudo apt install telegraf -y
> sudo systemctl start telegraf
```

You can then test it is running with:

```
> sudo systemctl status telegraf
```

If you see the result below, Telegraf has been installed correctly.

```
> user@sys:~$ sudo systemctl status telegraf

  telegraf.service - The plugin-driven server agent for reporting metrics into InfluxDB
Loaded: loaded (/lib/systemd/system/telegraf.service; enabled; vendor preset: enabled)
Active: active (running) since Wed 2019-03-27 10:20:22 UTC; 1min 42s ago
Docs: https://github.com/influxdata/telegraf
Main PID: 81754 (telegraf)
Tasks: 10 (limit: 1084)
```

3.1.5 Grafana

The last tool to install is *Grafana*, which provides a GUI for graphing data stored in InfluxDB. Similar to InfluxDB, you need to add a key and repository then install. The `echo 'deb...` command will only run if you have write privileges in /etc/apt/. On Ubuntu, this can be changed with `chmod`.

```
> sudo curl https://packagecloud.io/gpg.key | sudo apt-key add -
```

```
> echo 'deb https://packagecloud.io/grafana/stable/debian/ stretch main' >
/etc/apt/sources.list.d/grafana.list

> sudo apt update
> sudo apt install grafana -y
```

If the install is successful, you can start and test Grafana using:

```
> sudo systemctl start grafana-server
> sudo netstat -plntu
```

There should be a `grafana-server` running on port `3000`.

```
Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address      Foreign Address    State      PID/Program name
tcp        0      0 127.0.0.53:53      0.0.0.0:*          LISTEN     779/systemd-resolve
tcp        0      0 0.0.0.0:22         0.0.0.0:*          LISTEN     1051/sshd
tcp        0      0 127.0.0.1:8088     0.0.0.0:*          LISTEN     44367/influxd
tcp6       0      0 :::8086            :::*              LISTEN     44367/influxd
tcp6       0      0 :::22             :::*              LISTEN     1051/sshd
tcp6       0      0 :::3000            :::*              LISTEN     43929/grafana-server
udp        0      0 127.0.0.53:53      0.0.0.0:*                     779/systemd-resolve
udp        0      0 10.10.72.19:68     0.0.0.0:*                     761/systemd-network
```

3.2 Adding SNMP

3.2.1 Telegraf.conf

Once everything is installed, Telegraf needs reconfiguring to use SNMP as its input. This means changing the telegraf.conf file in /etc/telegraf/. Here's an example telegraf.conf file.

```
[[inputs.snmp]]
agents = [ "10.10.72.18:161" ]
version = 2
community = "public"
name = "snmp"

[[inputs.snmp.field]]
name = "hostname"
oid = "RFC1213-MIB::sysName.0"
is_tag = true

[[inputs.snmp.table]]
name = "snmp"
inherit_tags = [ "hostname" ]
oid = "BW-SVC-STATS-MIB::svcListenersTable"
```
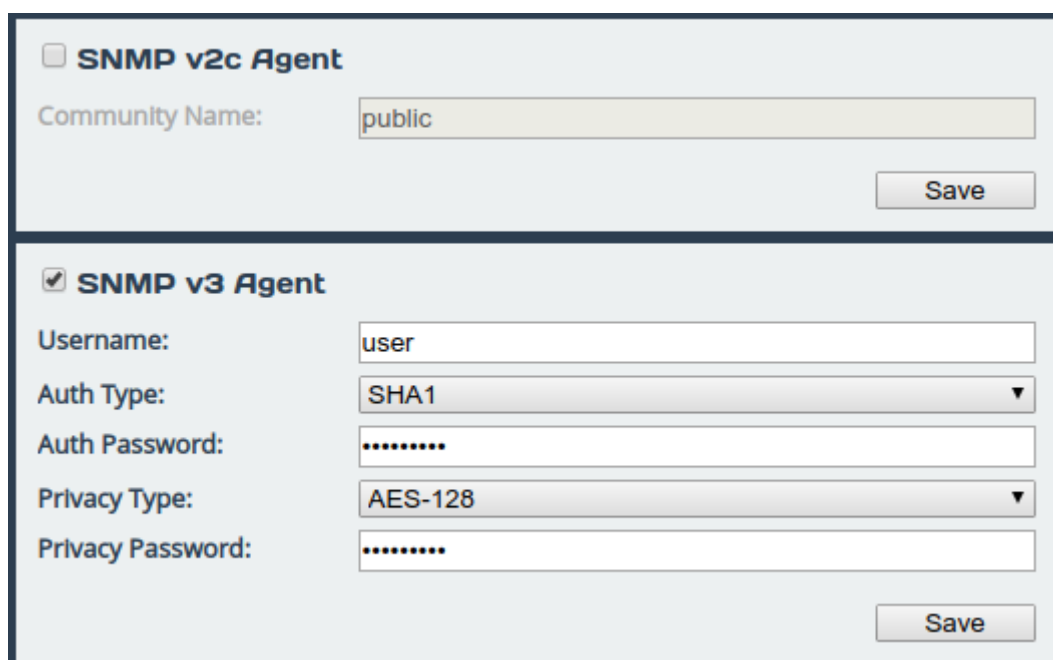
```
[outputs.influxdb]
url = "http://localhost:8086"
database = "telegraf"
```

The first 3 blocks are input plugins, which tell Telegraf where to fetch data from. The last is the output plugin, which is configured to send to the local InfluxDB database. This will be the same in your system.

The `[[inputs.snmp]]` block depends on your SNMP configuration. The field and table blocks dictate which information you poll from the agent. The above example is for a unit with IP `10.10.72.18` running *SNMPv2c* on the default port, `161`, with community string, `public`. Below is the agent configuration for a v3 agent.

```
[[inputs.snmp]]
agents = [ "10.10.72.18:161" ]
version = 3
sec_name = "user"
sec_level = "authPriv"
auth_protocol = "SHA"
auth_password = "password1"
priv_protocol = "AES"
priv_password = "password2"
name = "snmp"
```

SNMP v3 needs far more security information. If there is no auth or priv then only `agents`, `version`, and `sec_name` are needed as `sec_level` defaults to `noAuthNoPriv`. Other options for auth and priv protocols are "MD5" and "DES" respectively. If you have auth enabled and priv disabled, `sec_level` is `authNoPriv`. The screenshot below shows the units SNMP page corresponding to the above Telegraf configuration. SNMP v3 with auth and priv is recommended for security reasons.

Configuration of the field and table inputs varies depending on what information you wish to query the SNMP agent for. Our example gets a table of network service stats. The OID is the identifier

of information within the SNMP data structure. Changing this will change what data Telegraf gets from SNMP. Descriptive information about the data stored in particular OIDs can be looked up in relevant MIB files.
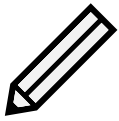
> Note: BW-SVC-STATS-MIB::svcListenersTable is only available if data acceleration is set up and a service exists. If you have not set up data acceleration yet, swap that OID for "IP-MIB::ipSystemStatsTable". This contains generic SNMP IP info.

Once you've saved your finished telegraf.conf file, restart the Telegraf process and test it using the commands below.

```
> sudo systemctl restart telegraf
> telegraf --test
```

If you have a valid configuration then you will see a block of SNMP data like the one below.

```
> user@sys:/etc/telegraf$ telegraf --test

2019-04-01T14:17:22Z I! Starting Telegraf 1.9.4
2019-04-01T14:17:22Z I! Using config file: /etc/telegraf/telegraf.conf
> snmp,agent_host=10.10.72.18,host=sys,hostname=bridgeworks,svcListenerPending=2i,
svcListenerProtocol="netapp",svcListenerRemote="564dbdca-eb84-d33d-a57d-58cb1b13a0e0",
svcListenerCountIn=2172983i,svcListenerCountOut=0i,svcListenerBytesIn=4947112063848i,
svcListenerBytesOut=0i,svcListenerUUID="33c28536-fdbd-5aac-89c8-9c2e1d166584",
svcListenerAddress="192.168.62.2",svcListenerPort="11104-11105" 1557829425000000000,
svcListenerVLANDomain=""
```

3.2.2 Grafana Dashboard

Telegraf should now be collecting the SNMP data and storing it in InfluxDB. The next step is to log into Grafana and create a graph with it.

Open a web browser, Chromium 73 is used in this guide, and enter your device's IP address followed by :3000 into the URL bar. This will be the IP address of the device running Grafana, not of your unit. You should see the login page in the screenshot below.

On first login the *username* and *password* are both "admin", you will be prompted to change these. Once on the dashboard, left click *Add data source* followed by *InfluxDB* then enter the details of your InfluxDB database, as shown in the screenshot below, then left click *Save & Test*.

The final step is to create a graph. Navigate to *Home* and left click *Create your first dashboard* followed by *Graph* then press the "e" key to enter edit mode. Change to the metrics tab and you'll see data is collected for the graph with an InfluxDB query. Below is a simple example setup showing: *snmp* as a source, the *agent_host* selected as your SNMP agent's IP address and a *field* which can be any numerical data from the agent.

Once these settings are selected SNMP data should begin to appear on the graph.

4 Troubleshooting

This section outlines some issues that can occur while following this guide and suggests some likely causes.

4.1 Install Errors

If "sudo apt update" fails and the error says *NO PUBKEY*, then you need to acquire a public key. You can import one using the command below.

```
> sudo apt-key adv --keyserver keyserver.ubuntu.com --recv-keys <PUBKEY>
```

If "sudo apt install snmp-mibs-downloader" throws an error saying snmp-mibs-downloader can't be found, you may need to edit apt's source list, found at etc/apt/sources.list and ensure multiverse is enabled. It should include these lines.

```
deb http://archive.ubuntu.com/ubuntu <OSversion> main multiverse universe
deb http://archive.ubuntu.com/ubuntu <OSversion>-security main multiverse universe
deb http://archive.ubuntu.com/ubuntu <OSversion>-updates main multiverse universe
```

If it doesn't, add them in, then run the install command again.

4.2 Dataflow Errors

If, when choosing metrics on Grafana, there are none of the expected options in your field list, either Grafana's InfluxDB source is incorrectly configured, or InfluxDB has no data. If your source configuration looks correct, you can check whether InfluxDB is receiving data by opening it and querying the database with the commands below.

```
> influx
> use telegraf
> SELECT * FROM snmp
```

If no data is present, your telegraf configuration is likely incorrect. Test it with the command below.

```
> telegraf --test
```

If no SNMP data is displayed, your input plugins may be configured wrong. Make sure your device can connect to the agent using a snmpwalk and mirror the snmpwalk's options in the `[[inputs.snmp]]` block of telegraf.conf. If that looks correct there may be an issue with your OIDs.

If "telegraf –test" does display SNMP data. It means the data is being collected but not reaching InfluxDB. Double check telegraf.conf's output plugins.

If when you run a snmpwalk including -m ALL and the data only has OIDs, no names, you don't have "SNMP-MIBs-downloader package" installed correctly.

4.3   No Data on Grafana Graph

If you can select the data on Grafana but no data is being displayed on the graph, it's likely there just isn't data in the time currently being displayed. Change the time settings to "last 30 minutes".

5 Expanding your Stack

This guide walked through setting up the simplest possible TIG stack that connects to the SNMP agent on our unit. There are many ways you'll want to expand this, either by getting more data out of our agent, collecting data from other SNMP devices, or adding more complex graphs to your Grafana dashboard.

The creators of the 4 programs involved all provide comprehensive documentation which will instruct you in expanding your setup and solve any issues not addressed in Troubleshooting.

Net-SNMP `http://www.net-snmp.org/`

Telegraf `https://docs.influxdata.com/telegraf/v1.10/`

Telegraf's SNMP Plugin `https://github.com/influxdata/telegraf/tree/release-1.6/plugins/inputs/snmp`

InfluxDB `https://docs.influxdata.com/influxdb/v1.7/`

Grafana `http://docs.grafana.org`

6 Useful Links

The following section contains links to other guides and FAQs. Support is available through our website: `https://support.4bridgeworks.com/`

The following resources are available online:

- User Manuals

- Installation Guides

- General FAQ

- AWS FAQ

If your question is not answered in our documentation, please submit a ticket through our website.