
kwapi-g5k Documentation

Release 1.0.0

Clement Parisot

August 10, 2015

1	What is the purpose of the project and vision for it?	3
2	Table of contents	5
2.1	Installing	5
2.2	System Architecture	6
2.3	Configuration Options	13
2.4	Kwapi Development	16
2.5	Glossary	17
3	Indices and tables	19

Kwapi is a framework designed for acquiring energy consumption and network metrics. It allows to import metrics from various sources and expose them in different ways.

Its architecture is based on a layer of drivers, which retrieve measurements from wattmeters or network switches, and a layer of plugins that collect and process them. The communication between these two layers goes through a bus. In the case of a distributed architecture, a plugin can listen to several drivers at remote locations.

Drivers and plugins are easily extensible to support other types of sources, and provide other services and metrics.

What is the purpose of the project and vision for it?

Kwapi could be used to do:

- Energy monitoring of data centers
- Usage-based billing
- Efficient scheduling
- Network traffic visualisation
- Long-term storage of measurements

It aims at supporting various wattmeters and switches, being scalable and easily extensible.

This documentation offers information on how Kwapi works and how to contribute to the project.

Table of contents

2.1 Installing

2.1.1 Installing Kwapi from source

1. Clone the Kwapi git repository to the management server:

```
$ git clone https://github.com/lpouillo/kwapi-g5k.git
```

2. Data management use numpy, which cannot be installed vi pip. On Debian/Ubuntu, use:

```
$ apt-get install python-numpy
```

3. Run the Kwapi installer and copy the configuration files:

```
$ pip install kwapi-g5k  
$ cp -r kwapi-g5k/etc/kwapi /etc/
```

2.1.2 Installing Kwapi on Grid'5000

1. Create a VM on the site you want to monitor. You can create a domU with Xen or KVM. The command should be similar to this one:

```
$ http_proxy="http://proxy:3128/" xen-create-image --hostname=kwapi --ip=ip_address --dist wheezy
```

You can find more informations about this procedure here: https://www.grid5000.fr/mediawiki/index.php/Services_Xen_DomU

2. Start your VM:

```
$ xm create /etc/xen/kwapi.cfg
```

3. Install configuration tool on the VM. Puppet manifests and files are available in **puppet-repo**. You must install the right version of Puppet used in Grid'5000. Installation procedure can be find here: [Puppet](#). After the certificate signing procedure, you should have a new Puppet node named **kwapi.site.grid5000.fr**.
4. Configure Kwapi with Puppet. You have to add additional classes on your new Puppet node. Use puppetplay to add **kwapi-g5k** class on the node:

```
$ puppetplay node kwapi.site.grid5000.fr --edit --add-classes kwapi-g5k
```

5. Wait one hour or use [Capistrano tool](#) to force Puppet execution on the VM.:

```
$ cap puppet:production HOST=kwapi.site.grid5000.fr MODULE=all SITE=site
```

6. Your VM is now configured with latest Grid'5000 version of Kwapi. You can connect on the node to check Kwapi service status.:

```
$ ssh kwapi.site.g5kadmin  
$ sudo service kwapi status
```

2.1.3 Running Kwapi modules as daemon

Kwapi can be started, stoped, restarted with the service command:

```
$ sudo service kwapi start|stop|restart
```

This command will start kwapi as a daemon and run the modules indicated in `/etc/kwapi/daemon.conf` file.

2.1.4 Running Kwapi modules in foreground (debugging)

If you want to manage each Kwapi module individually (drivers and plugins), you can run the following commands.

- Start the drivers on all the configured machines:

```
$ kwapi-drivers
```

- Start the forwarder on this machine and a remote machine (optional):

```
$ kwapi-forwarder
```

- Start the API plugin if you want to access metrics with the API:

```
$ kwapi-api
```

- Start the RRD plugin if you want to store data as RRD (mandatory to display graphs in a web browser):

```
$ kwapi-rrd
```

- Start the HDF5 plugin if you want to store fine grained datas:

```
$ kwapi-hdf5
```

- Start Live plugin to active Web visualisation of your mesures:

```
$ kwapi-live
```

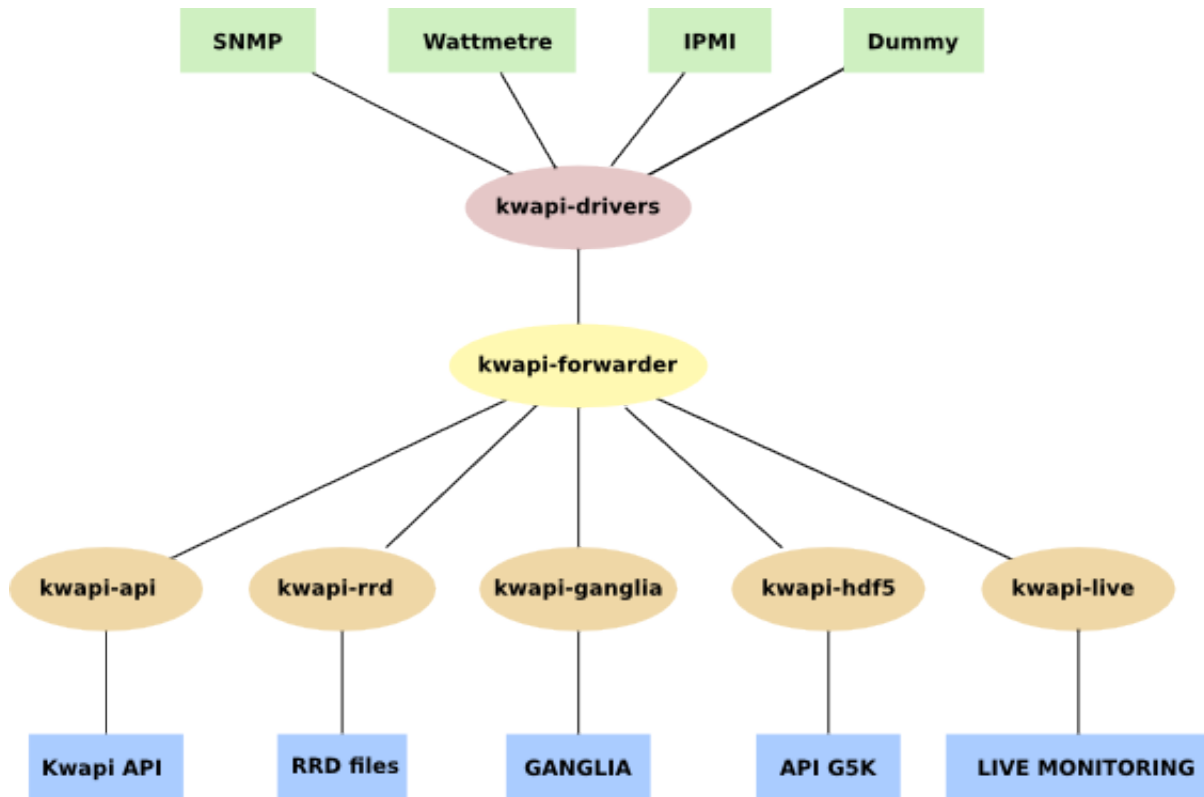
- Start Ganglia plugin to push data to the remote Ganglia server:

```
$ kwapi-ganglia
```

Warning: Don't forget to **stop** Kwapi daemon service before activating modules in foreground or it will result on conflict problems and data corruption !

2.2 System Architecture

Overview of the global layered architecture:



2.2.1 Kwapi drivers

Kwapi drivers are derived from a `Driver` superclass, itself derived from `Thread`. So drivers are threads. At least one driver thread is instantiated for each wattmeter or switch you want to monitor. Their constructors takes as arguments a list of probe IDs, and `kwargs` (specific arguments).

Driver threads roles are:

1. Setting up a wattmeter or a switch.
2. Listening and decoding received data.
3. Calling a driver superclass method with measurements as argument. This method appends signature to the measurements, and publishes them on the bus.

Message format:

Probe id	Name	Timestamp	Measure	Parameters
Key	Metric name	Seconds since epoch	Raw data from driver	Depends on metric type

There is several types of drivers already implemented in Kwapi. You can of course implement your own driver.

SNMP drivers

Using the SNMP protocol and the right OIDs, you can retrieve a subtree of values corresponding, for a wattmeter, at the power consumption, for a switch, at the network traffic of ihis interfaces.

Of course, this driver works with any device that implement SNMP protocol. So you can retrieve other metrics depending on what you want to monitor.

Wattmeters via SNMP

Kwapi supports different kinds of wattmeters (IPMI, Eaton PDU, Wattsup, etc). Wattmeters communicate via IP networks or serial links. Each wattmeter has one or more sensors (probes). Wattmeters send their values quite often (each second), and they are listen by wattmeter drivers.

Network drivers

Second type of drivers is the network drivers. Giving a switch address and the correct protocol, you can retrieve incoming and outgoing traffic. You have just to configure the right OIDs and allow SNMP requests on the switch. Every port of the given switch is scanned and the current counter value of the interface is assigned to the configured neighbor. This counter is a 64-bits number that correspond to the total number of octets received on the interface, including framing characters.

Dummy drivers

Added as a testing feature, dummy drivers just send the value you ask them to send. They are used to simulate a probe. You can configure them as you wish. Kwapi implements dummynet and dummywatt drivers to simulate network and wattmeter drivers.

IPMI drivers

IPMI drivers use the command line *ipmitool* to retrieve information from IPMI sensors.

JSON url drivers

This driver can be used to get informations directly from JSON structured text of Grid'5000 API. It can be usefull when you want to import in Kwapi remote metrics from the Metrologie API.

Driver manager

The driver manager is the loader and the checker of driver threads. It loads all drivers according the configuration file, and checks regularly that driver threads are alive. In case of crash, the event is logged and the driver thread is reloaded. We can imagine that a driver will crash if a technician unplug a wattmeter, for example.

Bus

Currently, the internal Kwapi bus is ZeroMQ. Publishers are driver threads, and subscribers are plugins.

2.2.2 Kwapi plugins

Kwapi API plugin

Kwapi API offers a REST API. This API is linked in Grid'5000 API and adopts its standard to expose live measures to the users. Such data can then be imported in experiments by just sending a request to the API. This plugin contains a collector that computes kWh for power, interface counters for network and an API based on Flask.

Collector

The collector stores these values for each probe:

Probe id Key	Timestamp Seconds since epoch	Integrated Power only	Value Instantaneous consumption	Unit Metric unit	Type Gauge or cumulative values
------------------------	---	---------------------------------	---	----------------------------	---

Fields:

- Probe id: could be the hostname of the monitored machine. But it is a bit more complicated because a probe can monitor several machines (PDU).
- Timestamp: is updated when a new value is received.
- Integrated (power only): is computed by taking into account the new value in watt, and the elapsed time since the previous update.
- Value: offers the possibility to know instantaneous consumption or traffic of a device, without having to query two times a probe in a small interval to deduce it. This could be especially useful if a probe has a large refresh interval: there is no need to wait its next value.
- Unit: metric unit. For example 'W' stands for *watt* in power API.
- Type: metric type can be 'Gauge' or 'Cumulative'. It indicates if measures are retrieved as counter or not and if an integrated value can be calculated

No history is kept by this plugin. Storage is offered with other plugins. The collector is cleaned periodically to prevent a deleted probe from being stored indefinitely in the collector. So when a probe has not been updated for a long time, it is deleted.

API

Verb	URL	Parameters	Expected result
GET	/probe-ids/		Returns all known probe IDs.
GET	/probes/		Returns all information about all known probes.
GET	/probes/<probe>/	probe id	Returns all information about this probe (id, timestamp, value, type, integrated).
GET	/probes/<probe>/<metric>/	probe id metric {power, network}	Returns the probe's metric value.

Kwapi RRD plugin

It stores information from the drivers directly in RRD files. The advantage of such files is that they permits to render efficiently Graphs with various scales. The size of the databases are constant. One database per probe and per metric is created. By default, RRD files are stored in `/var/lib/kwapi/kwapi-rrd`.

This plugin create and update automatically the RRD files, depending on the values he receives from the drivers.

As each plugin, he needs:

- Probe id: probe identifier (could be different than probe name)
- Name: metric name
- Timestamp: time of the measure, given by the driver, unix format timestamp
- Measure: measure
- Parameters: other informations about the metrics

Kwapi Live plugin

Web interface

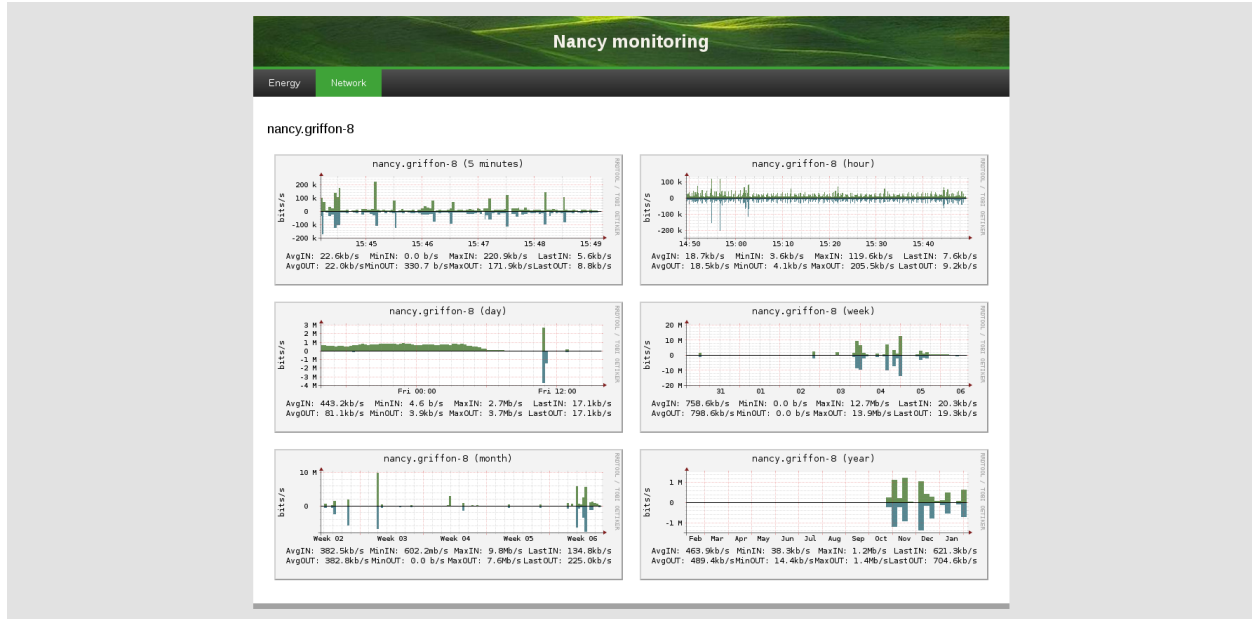
The visualization plugin provides a web interface with power consumption and network traffic graphs. It is based on Flask and RRDtool.

Verb	URL	Parameters	Expected result
GET	/<metric>/last/<period>/	metric { power, network } period { minute, hour, day, week, month, year }	Returns a webpage with a summary graph and all probe graphs.
GET	/<metric>/probe/<probe>/	metric { power, network } probe id	Returns a webpage with all graphs about this probe (all periods).
GET	/<metric>/summary-graph/<start>/<end>	metric { power, network } start timestamp end timestamp	Return a summary graph of the metric evolution about this period
GET	/<metric>/graph/<probe>/<start timestamp>/<end timestamp>	metric { power, network } probe id start timestamp end timestamp	Returns a graph about this probe.

Webpage with a summary graph and all probe graphs:



Webpage with scales summaries of a probe:



In the menu bar, you can choose the period the metric you want to display. For each metric you can select a timescale (last minutes, hour, day, week, month or year). By clicking on a probe, you can display all graphs available for this probe, with different resolutions.

You can select several probes and display a stacked summary of their consumption. Use the job field to automatically monitor probes corresponding to your job (select the correct probes and adapt timescale)

Graphs

The summary graph shows the total measurements for the selected metric (sum of all the probes). Each colour corresponds to a probe.

The legend contains:

- Minimum, maximum, average and last measures.
- Integrated measure (energy consumed (kWh) or network traffic (Kb/s)).
- Cost if any.

File sizes:

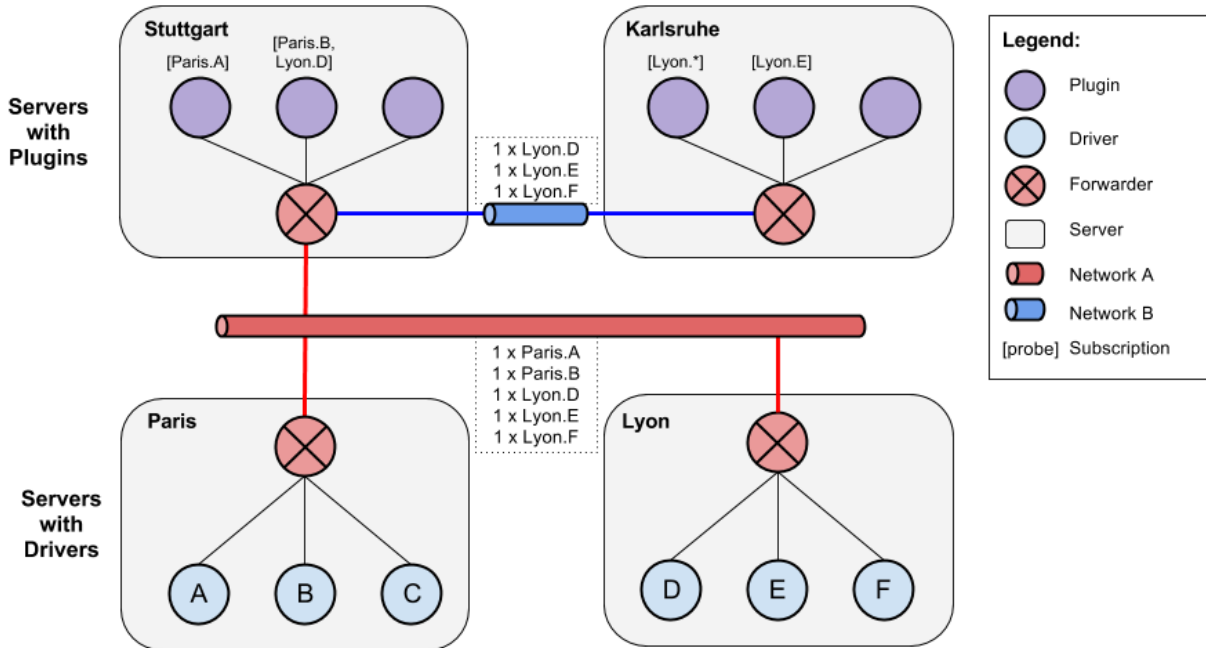
- RRD file: 10 Ko.
- Probe graph: 12 Ko.
- Summary graph: 24 Ko.

A cache mechanism prevents graphs from being rebuilt uselessly.

2.2.3 Kwapi forwarder

The forwarder aims at decreasing the network traffic: if multiple plugins listen the same probe, the metric is sent once on the network, and the forwarder duplicate it and sends a copy to each listeners. The forwarder can also be installed on a gateway machine, in order to connect isolated networks.

The following diagram shows these two features:



Using the forwarder is optional, and the plugins can be configured to subscribe directly to the drivers. Direct subscribing without using the forwarder is recommended if the drivers and the plugins are running on the same machine.

2.2.4 Kwapi HDF5

Kwapi HDF5 is used to store fine grained metric with Kwapi. Each measure returned by the drivers are stored in an HDF5 file on the server. The main advantages of this database are: * Very large datasets: store several months of power consumption of numerous probes * Fast access * Hierarchical store: data can be grouped by site or cluster * Parallel writing * Compressed file for low storage cost * Heterogeneous data support

You can configure the split period of your HDF5 files in the configuration file (1 file per month or less) depending on how much data you want to save.

Collector

The HDF5 Collector is composed of one Writer by metric with their proper buffers and a single queue per metric where all the

- The update function put the new received value in the queue that correspond to his metric.
- Each Collector iterate on his corresponding queue and for each new measure, writes an entry in his internal buffer
- When a Collector buffer is full, it writes his values to the database on the disk depending on the current date
- If the plugin is stopped, a **STOP** flag his added in all the queues
- When the Collector receive a **STOP** flag, buffers are flushed on the disk and Collector exits

Writes on the database are made with a fixed *chunk_size* that can be set in the configuration file.

API

REST API permits to retrieve measures from those databases. Unlike RRD database, HDFStore store raw measures and datas are not altered. API is very similar to Kwapi-API

Verb	URL	Parameters	Expected result
GET	/<metric>/	metric { power, network }	Returns all known probe IDs for the metric.
POST	/<metric>/timeseries/ [job_id probes]	metric { power, network } job_id in Grid 5000 probes: list of probes	Return all data for the selected probes and selected range. Selection is made with job_id or probe name given.

2.2.5 Kwapi Ganglia

This plugin his pretty simple. You first have to run a Ganglia server somewhere. It have to accept data from the remote Kwapi server. For this, edit the configuration file according to Ganglia Documentation. Check for the Ganglia **multicast address**.

The single parameter of this plugin is the *ganglia_server* address. Edit this field in the configuration file to point your remote Ganglia server. All data received from your drivers will be send to the server. Actual configuration just send single power probes consumption to Ganglia.

Metric	Remote name	Parameters
power	pdu2	<ul style="list-style-type: none"> units: Watts type: uint16 value: int(metrics) hostname: ip:hostname (ex: 192.168.1.1:griffon-54.nancy.grid5000.fr) spoof: True
network_in	None	
network_out	None	

2.3 Configuration Options

2.3.1 Kwapi drivers specific

The following table lists the Kwapi drivers specific options in the drivers configuration file. For information we are listing the configuration elements that we use after the Kwapi drivers specific elements.

Parameter	Default	Note
probes_endpoint	ipc:///tmp/kwapi-drivers	Endpoint where the drivers send their measurements ipc:///<file> or tcp:///<host>:<port>
enable_signing	true	Enable message signing between drivers and plugins
metering_secret	change this or be hacked	Secret value for signing metering messages
check_drivers_interval	60	Check drivers at the specified interval and restart them if they are crashed

The configuration file contains a section for each wattmeter.

A sample configuration file can be found in [drivers.conf](#).

2.3.2 Kwapi plugin API specific

The following table lists the Kwapi API specific options in the API configuration file. For information we are listing the configuration elements that we use after the Kwapi API specific elements.

Parameter	Default	Note
api_port	5000	API port
probes_endpoint	ipc:///tmp/kwapi-forwarder	Endpoint where the measurements are received
signature_checking	true	Enable the verification of signed metering messages
driver_metering_secret	change this or be hacked	Secret value for verifying signed metering messages
cleaning_interval	300	Delete the probes that have not been updated during the specified interval

A sample configuration file can be found in [api.conf](#).

2.3.3 Kwapi plugin RRD specific

The following table lists the Kwapi RRD specific options in the RRD configuration file. For information we are listing the configuration elements that we use after the Kwapi RRD specific elements.

Parameter	Default	Note
probes_endpoint	ipc:///tmp/kwapi-forwarder	Endpoint where the measurements are received
signature_checking	true	Enable the verification of signed metering messages
driver_metering_secret	change this or be hacked	Secret value for verifying signed metering messages
rrd_dir	/var/lib/kwapi/kwapi-rrd	The directory where are stored RRD files

A sample configuration file can be found in [rrd.conf](#).

2.3.4 Kwapi plugin Live specific

The following table lists the Kwapi Live specific options in the Live configuration file. For information we are listing the configuration elements that we use after the Kwapi Live specific elements.

Parameter	Default	Note
live_port	8080	Port used to display webpages
probes_endpoint	ipc:///tmp/kwapi-forwarder	Endpoint where the measurements are received
signature_checking	true	Enable the verification of signed metering messages
driver_metering_secret	change this or be hacked	Secret value for verifying signed metering messages
png_dir	/var/lib/kwapi/kwapi-png	The directory where are stored PNG files
rrd_dir	/var/lib/kwapi/kwapi-rrd	The directory where are stored RRD files
currency	€	The currency symbol used in graphs
kwh_price	0.125	The kWh price used in graphs
hue	100	The hue of the graphs
max_watts	400	The maximum value of the summary graph
refresh_interval	5	The webpage auto-refresh interval

A sample configuration file can be found in [live.conf](#).

Warning: Be sure that <i>rrd_dir</i> directory is the same in RRD Plugin and Live plugin

2.3.5 Kwapi plugin Ganglia specific

The following table lists the Kwapi Ganglia specific options in the Ganglia configuration file. For information we are listing the configuration elements that we use after the Kwapi API specific elements.

Parameter	Default	Note
ganglia_server	udp://239.2.11.71:8649	Ganglia server address
probes_endpoint	ipc:///tmp/kwapi-forwarder	Endpoint where the measurements are received
signature_checking	true	Enable the verification of signed metering messages
driver_metering_secret	change this or be hacked	Secret value for verifying signed metering messages

A sample configuration file can be found in [ganglia.conf](#).

2.3.6 General options

The following is the list of options that we use:

Parameter	Default	Note
log_file		Log output to a named file
verbose	true	Print more verbose output

2.3.7 Kwapi forwarder specific

The following table lists the Kwapi forwarder specific options in the forwarder configuration file. For information we are listing the configuration elements that we use after the Kwapi forwarder specific elements.

Parameter	Default	Note
for-forwarder_endpoint	ipc:///tmp/kwapi-forwarder	Endpoint where the measurements are forwarded and where the plugins subscriptions are received
probes_endpoint	ipc:///tmp/kwapi-drivers	Endpoint where the drivers send their measurements. ipc://<file> or tcp://<host>:<port>

A sample configuration file can be found in [forwarder.conf](#).

2.3.8 Kwapi Daemon specific

The following table lists the Kwapi service specific options in the daemon configuration file.

Set a parameter to **false** will not start the corresponding plugin/driver when you start the service.

Warning: Always run *service kwapi stop* **BEFORE** modifying any of the following parameters !

Parameter	Default	Note
KWAPI_DRIVERS	true	Start Kwapi drivers in kwapi service
KWAPI_FORWARDER	true	Start Kwapi forwarder in kwapi service
KWAPI_API	true	Start Kwapi api in kwapi service
KWAPI_RRD	true	Start Kwapi rrd in kwapi service
KWAPI_HDF5	true	Start Kwapi hdf5 in kwapi service
KWAPI_LIVE	true	Start Kwapi live in kwapi service
KWAPI_GANGLIA	true	Start Kwapi ganglia in kwapi service

A sample configuration file can be found in [daemon.conf](#).

2.4 Kwapi Development

2.4.1 Project Hosting Details

Bug tracker <https://intranet.grid5000.fr/bugzilla/>

Code Hosting <https://github.com/lpouillo/kwapi-g5k>

Mailing list g5k-dev team

Grid'5000 Doc <https://www.grid5000.fr/mediawiki/index.php/Monitoring>

See also:

- user

2.4.2 Areas to Contribute

Drivers

Kwapi aims at supporting various drivers. If you have a non-supported wattmeter, you can easily contribute by writing a new one.

Plugins

Kwapi plugins process the metrics. You can contribute by writing new plugins to bring new fonctionnalités.

Testing

The first version of Kwapi has not yet unit tests and has not seen much run-time in real environments.

2.4.3 Working with the Source

Setting up a Development Sandbox

1. Set up a server or virtual machine.
2. Clone the kwapi project to the machine:

```
$ git clone https://github.com/lpouillo/kwapi-g5k.git
$ cd ./kwapi-g5k
```

3. Once this is done, use develop option of *setup.py* file to install kwapi locally:

```
$ python setup.py develop
```

4. If some dependant packages are missing, fix them with *pip install*:

```
$ pip install -r requirments.txt
```

4. You can start to hack kwapi. If you are preparing a patch, create a topic branch and switch to it before making any changes:

```
$ git checkout -b TOPIC-BRANCH
```

5. Use git to push your changes and ask for a pull request.
6. Package your solution for Debian installation:

```
$ python setup.py --command-packages=stdeb.command bdist_deb
$ cd deb_dist/
```

All the deb archives are exported in this directory.

7. Import the new generated packages of kwapi-g5k on the remote apt repository.
8. Execute Puppet on the VM to install the latest version of Kwapi or simply run:

```
$ apt-get update && apt-get install python-kwapi-g5k
```

Code Reviews

Kwapi uses the GitHub to hos all code and developer documentation contributions. You can report an issue or a feature request on this repository.

Bugzilla can also be used for API related bugs or device configuration problems.

2.5 Glossary

driver Software thread running querying a wattmeter or switch and sending the results to the plugins.

forwarder Component that forwards plugins subscriptions and metrics. Used to minimize the network traffic, or to connect isolated networks through a gateway.

plugin An action triggered whenever a meter reaches a certain threshold.

probe A wattmeter sensor or network device. A wattmeter can have only one probe (usually the IPMI cards), or multiple probes (usually the PDUs). A network device usually have multiples probes that correspond to his network interfaces. One probe is defined for incoming traffic and one for outgoing traffic.

Indices and tables

- `genindex`
- `modindex`
- `search`

D

driver, [17](#)

F

forwarder, [17](#)

P

plugin, [17](#)

probe, [17](#)