Graphite Documentation

Release 0.9.11

Chris Davis

August 20, 2013

CONTENTS

CHAPTER

ONE

OVERVIEW

Graphite does two things:

1. Store numeric time-series data
2. Render graphs of this data on demand

What Graphite does not do is collect data for you, however there are some *tools* out there that know how to send data to graphite. Even though it often requires a little code, *sending data* to Graphite is very simple.

1.1 About the project

Graphite is an enterprise-scale monitoring tool that runs well on cheap hardware. It was originally designed and written by Chris Davis at Orbitz in 2006 as side project that ultimately grew to be a foundational monitoring tool. In 2008, Orbitz allowed Graphite to be released under the open source Apache 2.0 license. Since then Chris has continued to work on Graphite and has deployed it at other companies including Sears, where it serves as a pillar of the e-commerce monitoring system. Today many large *companies* use it.

1.2 The architecture in a nutshell

Graphite consists of 3 software components:

1. **carbon** - a Twisted daemon that listens for time-series data
2. **whisper** - a simple database library for storing time-series data (similar in design to RRD)
3. **graphite webapp** - A Django webapp that renders graphs on-demand using Cairo

Feeding in your data is pretty easy, typically most of the effort is in collecting the data to begin with. As you send datapoints to Carbon, they become immediately available for graphing in the webapp. The webapp offers several ways to create and display graphs including a simple *URL API* for rendering that makes it easy to embed graphs in other webpages.

CHAPTER

TWO

INSTALLING GRAPHITE

2.1 Dependencies

Graphite renders graphs using the Cairo graphics library. This adds dependencies on several graphics-related libraries not typically found on a server. If you're installing from source you can use the `check-dependencies.py` script to see if the dependencies have been met or not.

Basic Graphite requirements:

- Python 2.4 or greater (2.6+ recommended)
- Pycairo
- Django 1.0 or greater
- django-tagging 0.3.1 or greater
- Twisted 8.0 or greater (10.0+ recommended)
- zope-interface (often included in Twisted package dependency)
- fontconfig and at least one font package (a system package usually)
- A WSGI server and web server. Popular choices are: - Apache with mod_wsgi and mod_python - gunicorn with nginx - uWSGI with nginx

Python 2.4 and 2.5 have extra requirements:

- simplejson
- python-sqlite2 or another Django-supported database module

Additionally, the Graphite webapp and Carbon require the whisper database library which is part of the Graphite project.

There are also several other dependencies required for additional features:

- Render caching: memcached and python-memcache
- LDAP authentication: python-ldap (for LDAP authentication support in the webapp)
- AMQP support: txamqp
- RRD support: python-rrdtool
- Dependant modules for additional database support (MySQL, PostgreSQL, etc). See Django database install instructions and the Django database documentation for details

See Also:

On some systems it is necessary to install fonts for Cairo to use. If the webapp is running but all graphs return as broken images, this may be why.

- https://answers.launchpad.net/graphite/+question/38833
- https://answers.launchpad.net/graphite/+question/133390
- https://answers.launchpad.net/graphite/+question/127623

2.2 Fulfilling Dependencies

Most current Linux distributions have all of the requirements available in the base packages. RHEL based distributions may require the EPEL repository for requirements. Python module dependencies can be install with pip rather than system packages if desired or if using a Python version that differs from the system default. Some modules (such as Cairo) may require library development headers to be available.

2.3 Default Installation Layout

Graphite defaults to an installation layout that puts the entire install in its own directory: `/opt/graphite`

2.3.1 Whisper

Whisper is installed Python's system-wide site-packages directory with Whisper's utilities installed in the bin dir of the system's default prefix (generally `/usr/bin/`).

2.3.2 Carbon and Graphite-web

Carbon and Graphite-web are installed in `/opt/graphite/` with the following layout:

- `bin/`
- `conf/`
- `lib/`

  Carbon `PYTHONPATH`
- `storage/`
  - `log`

    Log directory for Carbon and Graphite-web
  - `rrd`

    Location for RRD files to be read
  - `whisper`

    Location for Whisper data files to be stored and read
- `webapp/`

  Graphite-web `PYTHONPATH`

  - `graphite/`

    Location of `manage.py` and `local_settings.py`

  - `content/`

    Graphite-web static content directory

2.4 Installing Graphite

Several installation options exist:

2.4.1 Installing From Source

The latest source tarballs for Graphite-web, Carbon, and Whisper may be fetched from the Graphite project download page or the latest development branches may be cloned from the Github project page:

- Graphite-web: `git clone https://github.com/graphite-project/graphite-web.git`
- Carbon: `git clone https://github.com/graphite-project/carbon.git`
- Whisper: `git clone https://github.com/graphite-project/whisper.git`

Installing in the Default Location

To install Graphite in the *default location*, `/opt/graphite/`, simply execute `python setup.py install` as root in each of the project directories for Graphite-web, Carbon, and Whisper.

Installing Carbon in a Custom Location

Carbon's `setup.py` installer is configured to use a `prefix` of `/opt/graphite` and an `install-lib` of `/opt/graphite/lib`. Carbon's lifecycle wrapper scripts and utilities are installed in `bin`, configuration within `conf`, and stored data in `storage` all within `prefix`. These may be overridden by passing parameters to the `setup.py install` command.

The following parameters influence the install location:

- `--prefix`

  Location to place the `bin/` and `storage/` and `conf/` directories (defaults to `/opt/graphite/`)

- `--install-lib`

  Location to install Python modules (default: `/opt/graphite/lib`)

- `--install-data`

  Location to place the `storage` and `conf` directories (default: value of `prefix`)

- `--install-scripts`

  Location to place the scripts (default: `bin/` inside of `prefix`)

For example, to install everything in `/srv/graphite/`:

```
python setup.py install --prefix=/srv/graphite --install-lib=/srv/graphite/lib
```

To install Carbon into the system-wide site-packages directory with scripts in `/usr/bin` and storage and configuration in `/usr/share/graphite`:

```
python setup.py install --install-scripts=/usr/bin --install-lib=/usr/lib/python2.6/site-packages --i
```

Installing Graphite-web in a Custom Location

Graphite-web's `setup.py` installer is configured to use a `prefix` of `/opt/graphite` and an `install-lib` of `/opt/graphite/webapp`. Utilities are installed in `bin`, and configuration in `conf` within the `prefix`. These may be overridden by passing parameters to `setup.py install`

The following parameters influence the install location:

- `--prefix`

  Location to place the `bin/` and `conf/` directories (defaults to `/opt/graphite/`)

- `--install-lib`

  Location to install Python modules (default: `/opt/graphite/webapp`)

- `--install-data`

  Location to place the `webapp/content` and `conf` directories (default: value of `prefix`)

- `--install-scripts`

  Location to place scripts (default: `bin/` inside of `prefix`)

For example, to install everything in `/srv/graphite/`:

```
python setup.py install --prefix=/srv/graphite --install-lib=/srv/graphite/webapp
```

To install the Graphite-web code into the system-wide site-packages directory with scripts in `/usr/bin` and storage configuration, and content in `/usr/share/graphite`:

```
python setup.py install --install-scripts=/usr/bin --install-lib=/usr/lib/python2.6/site-packages --i
```

2.4.2 Installing From Pip

Versioned Graphite releases can be installed via pip. When installing with pip, Installation of dependencies will automatically be attempted.

Installing in the Default Location

To install Graphite in the *default location*, `/opt/graphite/`, simply execute as root:

```
pip install whisper
pip install carbon
pip install graphite-web
```

Note: On RedHat-based systems using the `python-pip` package, the pip executable is named `pip-python`

Installing Carbon in a Custom Location

Installation of Carbon in a custom location with *pip* is similar to doing so from a source install. Arguments to the underlying `setup.py` controlling installation location can be passed through *pip* with the `--install-option` option.

See *Installing Carbon in a Custom Location* for details of locations and available arguments

For example, to install everything in `/srv/graphite/`:

```
pip install carbon --install-option="--prefix=/srv/graphite" --install-option="--install-lib=/srv/gra
```

To install Carbon into the system-wide site-packages directory with scripts in `/usr/bin` and storage and configuration in `/usr/share/graphite`:

```
pip install carbon --install-option="--install-scripts=/usr/bin" --install-option="--install-lib=/usr
```

Installing Graphite-web in a Custom Location

Installation of Graphite-web in a custom location with *pip* is similar to doing so from a source install. Arguments to the underlying `setup.py` controlling installation location can be passed through *pip* with the `--install-option` option.

See *Installing Graphite-web in a Custom Location* for details on default locations and available arguments

For example, to install everything in `/srv/graphite/`:

```
pip install graphite-web --install-option="--prefix=/srv/graphite" --install-option="--install-lib=/s
```

To install the Graphite-web code into the system-wide site-packages directory with scripts in `/usr/bin` and storage configuration, and content in `/usr/share/graphite`:

```
pip install graphite-web --install-option="--install-scripts=/usr/bin" install-option="--install-lib=
```

2.4.3 Installing in Virtualenv

Virtualenv provides an isolated Python environment to run Graphite in.

Installing in the Default Location

To install Graphite in the *default location*, `/opt/graphite/`, create a virtualenv in `/opt/graphite` and activate it:

```
virtualenv /opt/graphite
source /opt/graphite/bin/activate
```

Once the virtualenv is activated, Graphite and Carbon can be installed *from source* or *via pip*. Note that dependencies will need to be installed while the virtualenv is activated unless –system-site-packages is specified at virtualenv creation time.

Installing in a Custom Location

To install from source activate the virtualenv and see the instructions for *graphite-web* and *carbon*

Running Carbon Within Virtualenv

Carbon may be run within Virtualenv by activating virtualenv before Carbon is started

Running Graphite-web Within Virtualenv

Running Django's manage.py within virtualenv requires activating virtualenv before executing as normal.

The method of running Graphite-web within Virtualenv depends on the WSGI server used:

Apache mod_wsgi

Note: The version Python used to compile mod_wsgi must match the Python installed in the virtualenv (generally the system Python)

To the Apache mod_wsgi config, add the root of the virtualenv as `WSGIPythonHome`, `/opt/graphite` in this example:

```
WSGIPythonHome /opt/graphite
```

and add the virtualenv's python site-packages to the `graphite.wsgi` file, python 2.6 in `/opt/graphite` in this example:

```
site.addsitedir('/opt/graphite/lib/python2.6/site-packages')
```

See the *mod_wsgi documentation on Virtual Environments <http://code.google.com/p/modwsgi/wiki/VirtualEnvironments>* for more details.

Gunicorn

Ensure Gunicorn is installed in the activated virtualenv and execute as normal. If gunicorn is installed system-wide, it may be necessary to execute it from the virtualenv's bin path

uWSGI

Execute uWSGI using the `-H` option to specify the virtualenv root. See the uWSGI documentation on virtualenv for more details.

2.5 Initial Configuration

2.6 Help! It didn't work!

If you run into any issues with Graphite, please to post a question to our Questions forum on Launchpad or join us on IRC in #graphite on FreeNode

2.7 Post-Install Tasks

Configuring Carbon Once you've installed everything you will need to create some basic configuration. Initially none of the config files are created by the installer but example files are provided. Simply copy the `.example` files and customize.

`Administering Carbon` Once Carbon is configured, you need to start it up.

Feeding In Your Data Once it's up and running, you need to feed it some data.

Configuring The Webapp With data getting into carbon, you probably want to look at graphs of it. So now we turn our attention to the webapp.

Administering The Webapp Once its configured you'll need to get it running.

Using the Composer Now that the webapp is running, you probably want to learn how to use it.

CHAPTER

THREE

THE CARBON DAEMONS

When we talk about "Carbon" we mean one or more of various daemons that make up the storage backend of a Graphite installation. In simple installations, there is typically only one daemon, `carbon-cache.py`. This document gives a brief overview of what each daemon does and how you can use them to build a more sophisticated storage backend.

All of the carbon daemons listen for time-series data and can accept it over a common set of *protocols*. However, they differ in what they do with the data once they receive it.

3.1 carbon-cache.py

`carbon-cache.py` accepts metrics over various protocols and writes them to disk as efficiently as possible. This requires caching metric values in RAM as they are received, and flushing them to disk on an interval using the underlying *whisper* library.

`carbon-cache.py` requires some basic configuration files to run:

carbon.conf The `[cache]` section tells `carbon-cache.py` what ports (2003/2004/7002), protocols (newline delimited, pickle) and transports (TCP/UDP) to listen on.

storage-schemas.conf Defines a retention policy for incoming metrics based on regex patterns. This policy is passed to *whisper* when the `.wsp` file is pre-allocated, and dictates how long data is stored for.

As the number of incoming metrics increases, one `carbon-cache.py` instance may not be enough to handle the I/O load. To scale out, simply run multiple `carbon-cache.py` instances (on one or more machines) behind a `carbon-aggregator.py` or `carbon-relay.py`.

Warning: If clients connecting to the `carbon-cache.py` are experiencing errors such as *connection refused* by the daemon, a common reason is a shortage of file descriptors.
In the `console.log` file, if you find presence of:
`Could not accept new connection (EMFILE)`
or
`exceptions.IOError: [Errno 24] Too many open files: '/var/lib/graphite/whisper/systems/`
the number of files `carbon-cache.py` can open will need to be increased. Many systems default to a max of 1024 file descriptors. A value of 8192 or more may be necessary depending on how many clients are simultaneously connecting to the `carbon-cache.py` daemon.
In Linux, the system-global file descriptor max can be set via sysctl. Per-process limits are set via ulimit. See documentation for your operating system distribution for details on how to set these values.

3.2 carbon-relay.py

`carbon-relay.py` serves two distinct purposes: replication and sharding.

When running with `RELAY_METHOD = rules`, a `carbon-relay.py` instance can run in place of a `carbon-cache.py` server and relay all incoming metrics to multiple backend `carbon-cache.py`'s running on different ports or hosts.

In `RELAY_METHOD = consistent-hashing` mode, a `DESTINATIONS` setting defines a sharding strategy across multiple `carbon-cache.py` backends. The same consistent hashing list can be provided to the graphite webapp via `CARBONLINK_HOSTS` to spread reads across the multiple backends.

`carbon-relay.py` is configured via:

carbon.conf The `[relay]` section defines listener host/ports and a `RELAY_METHOD`

relay-rules.conf In `RELAY_METHOD = rules`, pattern/servers tuples define what servers metrics matching certain regex rules are forwarded to.

3.3 carbon-aggregator.py

`carbon-aggregator.py` can be run in front of `carbon-cache.py` to buffer metrics over time before reporting them into *whisper*. This is useful when granular reporting is not required, and can help reduce I/O load and whisper file sizes due to lower retention policies.

`carbon-aggregator.py` is configured via:

carbon.conf The `[aggregator]` section defines listener and destination host/ports.

aggregation-rules.conf Defines a time interval (in seconds) and aggregation function (sum or average) for incoming metrics matching a certain pattern. At the end of each interval, the values received are aggregated and published to `carbon-cache.py` as a single metric.

CHAPTER

FOUR

CONFIGURING CARBON

Carbon's config files all live in `/opt/graphite/conf/`. If you've just installed Graphite, none of the `.conf` files will exist yet, but there will be a `.conf.example` file for each one. Simply copy the example files, removing the .example extension, and customize your settings.

4.1 carbon.conf

This is the main config file, and defines the settings for each Carbon daemon.

Each setting within this file is documented via comments in the config file itself. The settings are broken down into sections for each daemon - carbon-cache is controlled by the `[cache]` section, carbon-relay is controlled by `[relay]` and carbon-aggregator by `[aggregator]`. However, if this is your first time using Graphite, don't worry about anything but the `[cache]` section for now.

Tip: Carbon-cache and carbon-relay can run on the same host! Try swapping the default ports listed for `LINE_RECEIVER_PORT` and `PICKLE_RECEIVER_PORT` between the `[cache]` and `[relay]` sections to prevent having to reconfigure your deployed metric senders. When setting `DESTINATIONS` in the `[relay]` section, keep in mind your newly-set `PICKLE_RECEIVER_PORT` in the `[cache]` section.

4.2 storage-schemas.conf

This configuration file details retention rates for storing metrics. It matches metric paths to patterns, and tells whisper what frequency and history of datapoints to store.

Important notes before continuing:

- There can be many sections in this file.
- The sections are applied in order from the top (first) and bottom (last).
- The patterns are regular expressions, as opposed to the wildcards used in the URL API.
- The first pattern that matches the metric name is used.
- This retention is set at the time the first metric is sent.
- Changing this file will not affect already-created .wsp files. Use whisper-resize.py to change those.

A given rule is made up of 3 lines:

- A name, specified inside square brackets.

- A regex, specified after "pattern="
- A retention rate line, specified after "retentions="

The retentions line can specify multiple retentions. Each retention of `frequency:history` is separated by a comma.

Frequencies and histories are specified using the following suffixes:

- s - second
- m - minute
- h - hour
- d - day
- y - year

Here's a simple, single retention example:

```
[garbage_collection]
pattern = garbageCollections$
retentions = 10s:14d
```

The name `[garbage_collection]` is mainly for documentation purposes, and will show up in `creates.log` when metrics matching this section are created.

The regular expression pattern will match any metric that ends with `garbageCollections`. For example, `com.acmeCorp.instance01.jvm.memory.garbageCollections` would match, but `com.acmeCorp.instance01.jvm.memory.garbageCollections.full` would not.

The retention line is saying that each datapoint represents 10 seconds, and we want to keep enough datapoints so that they add up to 14 days of data.

Here's a more complicated example with multiple retention rates:

```
[apache_busyWorkers]
pattern = ^servers\.www.*\.workers\.busyWorkers$
retentions = 15s:7d,1m:21d,15m:5y
```

In this example, imagine that your metric scheme is `servers.<servername>.<metrics>`. The pattern would match server names that start with 'www', followed by anything, that are sending metrics that end in '.workers.busyWorkers' (note the escaped '.' characters).

Additionally, this example uses multiple retentions. The general rule is to specify retentions from most-precise:least-history to least-precise:most-history – whisper will properly downsample metrics (averaging by default) as thresholds for retention are crossed.

By using multiple retentions, you can store long histories of metrics while saving on disk space and I/O. Because whisper averages (by default) as it downsamples, one is able to determine totals of metrics by reversing the averaging process later on down the road.

Example: You store the number of sales per minute for 1 year, and the sales per hour for 5 years after that. You need to know the total sales for January 1st of the year before. You can query whisper for the raw data, and you'll get 24 datapoints, one for each hour. They will most likely be floating point numbers. You can take each datapoint, multiply by 60 (the ratio of high-precision to low-precision datapoints) and still get the total sales per hour.

Additionally, whisper supports a legacy retention specification for backwards compatibility reasons - `seconds-per-datapoint:count-of-datapoints`

```
retentions = 60:1440
```

60 represents the number of seconds per datapoint, and 1440 represents the number of datapoints to store. This required some unnecessarily complicated math, so although it's valid, it's not recommended.

4.3 storage-aggregation.conf

This file defines how to aggregate data to lower-precision retentions. The format is similar to `storage-schemas.conf`. Important notes before continuing:

- This file is optional. If it is not present, defaults will be used.
- There is no `retentions` line. Instead, there are `xFilesFactor` and/or `aggregationMethod` lines.
- `xFilesFactor` should be a floating point number between 0 and 1, and specifies what fraction of the previous retention level's slots must have non-null values in order to aggregate to a non-null value. The default is 0.5.
- `aggregationMethod` specifies the function used to aggregate values for the next retention level. Legal methods are `average`, `sum`, `min`, `max`, and `last`. The default is `average`.
- These are set at the time the first metric is sent.
- Changing this file will not affect .wsp files already created on disk. Use whisper-set-aggregation-method.py to change those.

Here's an example:

```
[all_min]
pattern = \.min$
xFilesFactor = 0.1
aggregationMethod = min
```

The pattern above will match any metric that ends with `.min`.

The `xFilesFactor` line is saying that a minimum of 10% of the slots in the previous retention level must have values for next retention level to contain an aggregate. The `aggregationMethod` line is saying that the aggregate function to use is `min`.

If either `xFilesFactor` or `aggregationMethod` is left out, the default value will be used.

The aggregation parameters are kept separate from the retention parameters because the former depends on the type of data being collected and the latter depends on volume and importance.

4.4 relay-rules.conf

Relay rules are used to send certain metrics to a certain backend. This is handled by the carbon-relay system. It must be running for relaying to work. You can use a regular expression to select the metrics and define the servers to which they should go with the servers line.

Example:

```
[example]
pattern = ^mydata\.foo\..+
servers = 10.1.2.3, 10.1.2.4:2004, myserver.mydomain.com
```

You must define at least one section as the default.

4.5 aggregation-rules.conf

Aggregation rules allow you to add several metrics together as the come in, reducing the need to sum() many metrics in every URL. Note that unlike some other config files, any time this file is modified it will take effect automatically. This requires the carbon-aggregator service to be running.

The form of each line in this file should be as follows:

```
output_template (frequency) = method input_pattern
```

This will capture any received metrics that match 'input_pattern' for calculating an aggregate metric. The calculation will occur every 'frequency' seconds and the 'method' can specify 'sum' or 'avg'. The name of the aggregate metric will be derived from 'output_template' filling in any captured fields from 'input_pattern'.

For example, if your metric naming scheme is:

```
<env>.applications.<app>.<server>.<metric>
```

You could configure some aggregations like so:

```
<env>.applications.<app>.all.requests (60) = sum <env>.applications.<app>.*.requests
<env>.applications.<app>.all.latency (60) = avg <env>.applications.<app>.*.latency
```

As an example, if the following metrics are received:

```
prod.applications.apache.www01.requests
prod.applications.apache.www02.requests
prod.applications.apache.www03.requests
prod.applications.apache.www04.requests
prod.applications.apache.www05.requests
```

They would all go into the same aggregation buffer and after 60 seconds the aggregate metric 'prod.applications.apache.all.requests' would be calculated by summing their values.

4.6 whitelist and blacklist

The whitelist functionality allows any of the carbon daemons to only accept metrics that are explicitly whitelisted and/or to reject blacklisted metrics. The functionality can be enabled in carbon.conf with the `USE_WHITELIST` flag. This can be useful when too many metrics are being sent to a Graphite instance or when there are metric senders sending useless or invalid metrics.

`GRAPHITE_CONF_DIR` is searched for `whitelist.conf` and `blacklist.conf`. Each file contains one regular expressions per line to match against metric values. If the whitelist configuration is missing or empty, all metrics will be passed through by default.

CHAPTER

FIVE

FEEDING IN YOUR DATA

Getting your data into Graphite is very flexible. There are three main methods for sending data to Graphite: Plaintext, Pickle, and AMQP.

It's worth noting that data sent to Graphite is actually sent to the *Carbon and Carbon-Relay*, which then manage the data. The Graphite web interface reads this data back out, either from cache or straight off disk.

Choosing the right transfer method for you is dependent on how you want to build your application or script to send data:

- For a singular script, or for test data, the plaintext protocol is the most straightforward method.
- For sending large amounts of data, you'll want to batch this data up and send it to Carbon's pickle receiver.
- Finally, Carbon can listen to a message bus, via AMQP.

5.1 The plaintext protocol

The plaintext protocol is the most straightforward protocol supported by Carbon.

The data sent must be in the following format: `<metric path> <metric value> <metric timestamp>`. Carbon will then help translate this line of text into a metric that the web interface and Whisper understand.

On Unix, the `nc` program can be used to create a socket and send data to Carbon (by default, 'plaintext' runs on port 2003):

```
PORT=2003
SERVER=graphite.your.org
echo "local.random.diceroll 4 `date +%s`" | nc ${SERVER} ${PORT};
```

5.2 The pickle protocol

The pickle protocol is a much more efficient take on the plaintext protocol, and supports sending batches of metrics to Carbon in one go.

The general idea is that the pickled data forms a list of multi-level tuples:

```
[(path, (timestamp, value)), ...]
```

Once you've formed a list of sufficient size (don't go too big!), send the data over a socket to Carbon's pickle receiver (by default, port 2004). You'll need to pack your pickled data into a packet containing a simple header:

```
payload = pickle.dumps(listOfMetricTuples)
header = struct.pack("!L", len(payload))
message = header + payload
```

You would then send the `message` object through a network socket.

5.3 Using AMQP

...

CHAPTER

SIX

CONFIGURING THE WEBAPP

CHAPTER

SEVEN

ADMINISTERING THE WEBAPP

CHAPTER

EIGHT

USING THE COMPOSER

...

CHAPTER

NINE

THE RENDER URL API

The graphite webapp provides a `/render` endpoint for generating graphs and retreiving raw data. This endpoint accepts various arguments via query string parameters. These parameters are separated by an ampersand (`&`) and are supplied in the format:

```
&name=value
```

To verify that the api is running and able to generate images, open `http://GRAPHITE_HOST:GRAPHITE_PORT/render` in a browser. The api should return a simple 330x250 image with the text "No Data".

Once the api is running and you've begun *feeding data into carbon*, use the parameters below to customize your graphs and pull out raw data. For example:

```
# single server load on large graph
http://graphite/render?target=server.web1.load&height=800&width=600

# average load across web machines over last 12 hours
http://graphite/render?target=averageSeries(server.web*.load)&from=-12hours

# number of registered users over past day as raw json data
http://graphite/render?target=app.numUsers&format=json

# rate of new signups per minute
http://graphite/render?target=summarize(deriviative(app.numUsers),"1min")&title=New_Users_Per_Minute
```

Note: Most of the functions and parameters are case sensitive. For example `&linewidth=2` will fail silently. The correct parameter in this case is `&lineWidth=2`

9.1 Graphing Metrics

To begin graphing specific metrics, pass one or more target parameters and specify a time window for the graph via from / until.

9.1.1 target

This will draw one or more metrics

Example:

```
&target=company.server05.applicationInstance04.requestsHandled
(draws one metric)
```

Let's say there are 4 identical application instances running on each server.

```
&target=company.server05.applicationInstance*.requestsHandled
(draws 4 metrics / lines)
```

Now let's say you have 10 servers.

```
&target=company.server*.applicationInstance*.requestsHandled
(draws 40 metrics / lines)
```

You can also run any number of *functions* on the various metrics before graphing.

```
&target=averageSeries(company.server*.applicationInstance.requestsHandled)
(draws 1 aggregate line)
```

The target param can also be repeated to graph multiple related metrics.

```
&target=company.server1.loadAvg&target=company.server1.memUsage
```

Note: If more than 10 metrics are drawn the legend is no longer displayed. See the hideLegend parameter for details.

9.1.2 from / until

These are optional parameters that specify the relative or absolute time period to graph. `from` specifies the beginning, `until` specifies the end. If `from` is omitted, it defaults to 24 hours ago. If `until` is omitted, it defaults to the current time (now).

There are multiple formats for these functions:

```
&from=-RELATIVE_TIME
&from=ABSOLUTE_TIME
```

RELATIVE_TIME is a length of time since the current time. It is always preceded my a minus sign (-) and follow by a unit of time. Valid units of time:

Abbreviation	Unit
s	Seconds
min	Minutes
h	Hours
d	Days
w	Weeks
mon	30 Days (month)
y	365 Days (year)

ABSOLUTE_TIME is in the format HH:MM_YYMMDD, YYYYMMDD, MM/DD/YY, or any other `at(1)`-compatible time format.

Abbreviation	Meaning
HH	Hours, in 24h clock format. Times before 12PM must include leading zeroes.
MM	Minutes
YYYY	4 Digit Year.
MM	Numeric month representation with leading zero
DD	Day of month with leading zero

`&from` and `&until` can mix absolute and relative time if desired.

Examples:

```
&from=-8d&until=-7d
(shows same day last week)

&from=04:00_20110501&until=16:00_20110501
(shows 4AM-4PM on May 1st, 2011)

&from=20091201&until=20091231
(shows December 2009)

&from=noon+yesterday
(shows data since 12:00pm on the previous day)

&from=6pm+today
(shows data since 6:00pm on the same day)

&from=january+1
(shows data since the beginning of the current year)

&from=monday
(show data since the previous monday)
```

9.2 Data Display Formats

Along with rendering an image, the api can also generate SVG with embedded metadata or return the raw data in various formats for external graphing, analysis or monitoring.

9.2.1 format

Controls the format of data returned. Affects all `&targets` passed in the URL.

Examples:

```
&format=png
&format=raw
&format=csv
&format=json
&format=svg
```

png

Renders the graph as a PNG image of size determined by width and height

raw

Renders the data in a custom line-delimited format. Targets are output one per line and are of the format `<target name>,<start timestamp>,<end timestamp>,<series step>|[data]*`

```
entries,1311836008,1311836013,1|1.0,2.0,3.0,5.0,6.0
```

csv

Renders the data in a CSV format suitable for import into a spreadsheet or for processing in a script

```
entries,2011-07-28 01:53:28,1.0
entries,2011-07-28 01:53:29,2.0
entries,2011-07-28 01:53:30,3.0
entries,2011-07-28 01:53:31,5.0
entries,2011-07-28 01:53:32,6.0
```

json

Renders the data as a json object. The jsonp option can be used to wrap this data in a named call for cross-domain access

```
[{
  "target": "entries",
  "datapoints": [
    [1.0, 1311836008],
    [2.0, 1311836009],
    [3.0, 1311836010],
    [5.0, 1311836011],
    [6.0, 1311836012]
  ]
}]
```

svg

Renders the graph as SVG markup of size determined by width and height. Metadata about the drawn graph is saved as an embedded script with the variable `metadata` being set to an object describing the graph

```
<script>
  <![CDATA[
    metadata = {
      "area": {
        "xmin": 39.195507812499997,
        "ymin": 33.96875,
        "ymax": 623.794921875,
        "xmax": 1122
      },
      "series": [
        {
          "start": 1335398400,
          "step": 1800,
          "end": 1335425400,
          "name": "summarize(test.data, \"30min\", \"sum\")",
          "color": "#859900",
          "data": [null, null, 1.0, null, 1.0, null, 1.0, null, 1.0, null, 1.0, null, null, null, nu
          "options": {},
          "valuesPerPoint": 1
        }
      ],
      "y": {
        "labelValues": [0, 0.25, 0.5, 0.75, 1.0],
        "top": 1.0,
        "labels": ["0 ", "0.25 ", "0.50 ", "0.75 ", "1.00  "],
```

```
        "step": 0.25,
        "bottom": 0
      },
      "x": {
        "start": 1335398400,
        "end": 1335423600
      },
      "font": {
        "bold": false,
        "name": "Sans",
        "italic": false,
        "size": 10
      },
      "options": {
        "lineWidth": 1.2
      }
    }
  ]]>
</script>
```

pickle

Returns a Python pickle (serialized Python object). The response will have the MIME type 'application/pickle'. The pickled object is a list of dictionaries with the keys: `name`, `start`, `end`, `step`, and `values` as below:

```
[
  {
    'name' : 'summarize(test.data, "30min", "sum")',
    'start': 1335398400,
    'end'  : 1335425400,
    'step' : 1800,
    'values' : [None, None, 1.0, None, 1.0, None, 1.0, None, 1.0, None, 1.0, None, None, None, None],
  }
]
```

9.2.2 rawData

Deprecated since version 0.9.9. Used to get numerical data out of the webapp instead of an image. Can be set to true, false, csv. Affects all `&targets` passed in the URL.

Example:

```
&target=carbon.agents.graphiteServer01.cpuUsage&from=-5min&rawData=true
```

Returns the following text:

```
carbon.agents.graphiteServer01.cpuUsage,1306217160,1306217460,60|0.0,0.00666666520965,0.006666666242
```

9.3 Graph Parameters

9.3.1 areaAlpha

Default: 1.0

Takes a floating point number between 0.0 and 1.0 Sets the alpha (transparency) value of filled areas when using an areaMode

9.3.2 areaMode

Default: none

Enables filling of the area below the graphed lines. Fill area is the same color as the line color associated with it. See areaAlpha to make this area transparent. Takes one of the following parameters which determines the fill mode to use:

`none` Disables areaMode

`first` Fills the area under the first target and no other

`all` Fills the areas under each target

`stacked` Creates a graph where the filled area of each target is stacked on one another. Each target line is displayed as the sum of all previous lines plus the value of the current line.

9.3.3 bgcolor

Default: value from the [default] template in graphTemplates.conf

Sets the background color of the graph.

Color Names	RGB Value
black	0,0,0
white	255,255,255
blue	100,100,255
green	0,200,0
red	200,0,50
yellow	255,255,0
orange	255, 165, 0
purple	200,100,255
brown	150,100,50
aqua	0,150,150
gray	175,175,175
grey	175,175,175
magenta	255,0,255
pink	255,100,100
gold	200,200,0
rose	200,150,200
darkblue	0,0,255
darkgreen	0,255,0
darkred	255,0,0
darkgray	111,111,111
darkgrey	111,111,111

RGB can be passed directly in the format #RRGGBB where RR, GG, and BB are 2-digit hex vaules for red, green and blue, respectively.

Examples:

```
&bgcolor=blue
&bgcolor=#2222FF
```

9.3.4 cacheTimeout

Default: The value of DEFAULT_CACHE_DURATION from local_settings.py

The time in seconds for the rendered graph to be cached (only relevant if memcached is configured)

9.3.5 colorList

Default: value from the [default] template in graphTemplates.conf

Takes one or more comma-separated color names or RGB values (see bgcolor for a list of color names) and uses that list in order as the colors of the lines. If more lines / metrics are drawn than colors passed, the list is reused in order.

Example:

```
&colorList=green,yellow,orange,red,purple,#DECAFF
```

9.3.6 drawNullAsZero

Default: false

Converts any None (null) values in the displayed metrics to zero at render time.

9.3.7 fgcolor

Default: value from the [default] template in graphTemplates.conf

Sets the foreground color. This only affects the title, legend text, and axis labels.

See majorGridLineColor, and minorGridLineColor for further control of colors.

See bgcolor for a list of color names and details on formatting this parameter.

9.3.8 fontBold

Default: value from the [default] template in graphTemplates.conf

If set to true, makes the font bold.

Example:

```
&fontBold=true
```

9.3.9 fontItalic

Default: value from the [default] template in graphTemplates.conf

If set to true, makes the font italic / oblique. Default is false.

Example:

```
&fontItalic=true
```

9.3.10 fontName

Default: value from the [default] template in graphTemplates.conf

Change the font used to render text on the graph. The font must be installed on the Graphite Server.

Example:

```
&fontName=FreeMono
```

9.3.11 fontSize

Default: value from the [default] template in graphTemplates.conf

Changes the font size. Must be passed a positive floating point number or integer equal to or greater than 1. Default is 10

Example:

```
&fontSize=8
```

9.3.12 format

See: Data Display Formats

9.3.13 from

See: from / until

9.3.14 graphOnly

Default: False

Display only the graph area with no grid lines, axes, or legend

9.3.15 graphTypes

Default: line

Sets the type of graph to be rendered. Currently there are only two graph types:

line A line graph displaying metrics as lines over time

pie A pie graph with each slice displaying an aggregate of each metric calculated using the function specified by pieMode

9.3.16 hideLegend

Default: <unset>

If set to `true`, the legend is not drawn. If set to `false`, the legend is drawn. If unset, the `LEGEND_MAX_ITEMS` settings in `local_settings.py` is used to determine whether or not to display the legend.

Hint: If set to `false` the `&height` parameter may need to be increased to accommodate the additional text.

Example:

```
&hideLegend=false
```

9.3.17 hideAxes

Default: False

If set to `true` the X and Y axes will not be rendered Example:

```
&hideAxes=true
```

9.3.18 hideYAxis

Default: False

If set to `true` the Y Axis will not be rendered

9.3.19 hideGrid

Default: False

If set to `true` the grid lines will not be rendered

Example:

```
&hideGrid=true
```

9.3.20 height

Default: 250

Sets the height of the generated graph image in pixels.

See also: width

Example:

```
&width=650&height=250
```

9.3.21 jsonp

Default: <unset>

If set and combined with `format=json`, wraps the JSON response in a function call named by the parameter specified.

9.3.22 leftColor

Default: color chosen from colorList

In dual Y-axis mode, sets the color of all metrics associated with the left Y-axis.

9.3.23 leftDashed

Default: False

In dual Y-axis mode, draws all metrics associated with the left Y-axis using dashed lines

9.3.24 leftWidth

Default: value of the parameter lineWidth

In dual Y-axis mode, sets the line width of all metrics associated with the left Y-axis

9.3.25 lineMode

Default: slope

Sets the line drawing behavior. Takes one of the following parameters:

`slope` Slope line mode draws a line from each point to the next. Periods will Null values will not be drawn

`staircase` Staircase draws a flat line for the duration of a time period and then a vertical line up or down to the next value

`connected` Like a slope line, but values are always connected with a slope line, regardless of whether or not there are Null values between them

Example:

```
&lineMode=staircase
```

9.3.26 lineWidth

Default: 1.2

Takes any floating point or integer (negative numbers do not error but will cause no line to be drawn). Changes the width of the line in pixels.

Example:

```
&lineWidth=2
```

9.3.27 logBase

Default: <unset>

If set, draws the graph with a logarithmic scale of the specified base (e.g. 10 for common logarithm)

9.3.28 localOnly

Default: False

Set to prevent fetching from remote Graphite servers, only returning metrics which are accessible locally

9.3.29 majorGridLineColor

Default: value from the [default] template in graphTemplates.conf

Sets the color of the major grid lines.

See bgcolor for valid color names and formats.

Example:

```
&majorGridLineColor=#FF22FF
```

9.3.30 margin

Default: 10 Sets the margin around a graph image in pixels on all sides.

Example:

```
&margin=20
```

9.3.31 max

Deprecated since version 0.9.0: See yMax

9.3.32 minorGridLineColor

Default: value from the [default] template in graphTemplates.conf

Sets the color of the minor grid lines.

See bgcolor for valid color names and formats.

Example:

```
&minorGridLineColor=darkgrey
```

9.3.33 minorY

Sets the number of minor grid lines per major line on the y-axis.

Example:

```
&minorY=3
```

9.3.34 min

Deprecated since version 0.9.0: See yMin

9.3.35 minXStep

Default: 1

Sets the minimum pixel-step to use between datapoints drawn. Any value below this will trigger a point consolidation of the series at render time. The default value of `1` combined with the default lineWidth of `1.2` will cause a minimal amount of line overlap between close-together points. To disable render-time point consolidation entirely, set this to `0` though note that series with more points than there are pixels in the graph area (e.g. a few month's worth of per-minute data) will look very 'smooshed' as there will be a good deal of line overlap. In response, one may use lineWidth to compensate for this.

9.3.36 noCache

Default: False

Set to disable caching of rendered images

9.3.37 pickle

Deprecated since version 0.9.10: See Data Display Formats

9.3.38 pieMode

Default: average

The type of aggregation to use to calculate slices of a pie when `graphType=pie`. One of:

`average` The average of non-null points in the series

`maximum` The maximum of non-null points in the series

`minimum` THe minimum of non-null points in the series

9.3.39 rightColor

Default: color chosen from colorList

In dual Y-axis mode, sets the color of all metrics associated with the right Y-axis.

9.3.40 rightDashed

Default: False

In dual Y-axis mode, draws all metrics associated with the right Y-axis using dashed lines

9.3.41 rightWidth

Default: value of the parameter lineWidth

In dual Y-axis mode, sets the line width of all metrics associated with the right Y-axis

9.3.42 template

Default: default

Used to specify a template from `graphTemplates.conf` to use for default colors and graph styles.

Example:

```
&template=plain
```

9.3.43 thickness

Deprecated since version 0.9.0: See: lineWidth

9.3.44 title

Default: <unset>

Puts a title at the top of the graph, center aligned. If unset, no title is displayed.

Example:

```
&title=Apache Busy Threads, All Servers, Past 24h
```

9.3.45 tz

Default: The timezone specified in local_settings.py

Time zone to convert all times into.

Examples:

```
&tz=America/Los_Angeles
&tz=UTC
```

Note: To change the default timezone, edit `webapp/graphite/local_settings.py`.

9.3.46 uniqueLegend

Default: False

Display only unique legend items, removing any duplicates

9.3.47 until

See: from / until

9.3.48 vtitle

Default: <unset>

Labels the y-axis with vertical text. If unset, no y-axis label is displayed.

Example:

```
&vtitle=Threads
```

9.3.49 vtitleRight

Default: <unset>

In dual Y-axis mode, sets the title of the right Y-Axis (See: vtitle)

9.3.50 width

Default: 330

Sets the width of the generated graph image in pixels.

See also: height

Example:

```
&width=650&height=250
```

9.3.51 xFormat

Default: Determined automatically based on the time-width of the X axis

Sets the time format used when displaying the X-axis. See datetime.date.strftime() for format specification details.

9.3.52 yAxisSide

Default: left

Sets the side of the graph on which to render the Y-axis. Accepts values of `left` or `right`

9.3.53 yDivisor

Default: 4,5,6

Supplies the preferred number of intermediate values for the Y-axis to display (Y values between the min and max). Note that Graphite will ultimately choose what values (and how many) to display based on a set of 'pretty' values. To explicitly set the Y-axis values, see yStep

9.3.54 yLimit

Reserved for future use See: yMax

9.3.55 yLimitLeft

Reserved for future use See: yMaxLeft

9.3.56 yLimitRight

Reserved for future use See: yMaxRight

9.3.57 yMin

Default: The lowest value of any of the series displayed

Manually sets the lower bound of the graph. Can be passed any integer or floating point number.

Example:

```
&yMin=0
```

9.3.58 yMax

Default: The highest value of any of the series displayed

Manually sets the upper bound of the graph. Can be passed any integer or floating point number.

Example:

```
&yMax=0.2345
```

9.3.59 yMaxLeft

In dual Y-axis mode, sets the upper bound of the left Y-Axis (See: yMax)

9.3.60 yMaxRight

In dual Y-axis mode, sets the upper bound of the right Y-Axis (See: yMax)

9.3.61 yMinLeft

In dual Y-axis mode, sets the lower bound of the left Y-Axis (See: yMin)

9.3.62 yMinRight

In dual Y-axis mode, sets the lower bound of the right Y-Axis (See: yMin)

9.3.63 yStep

Default: Calculated automatically

Manually set the value step between Y-axis labels and grid lines

9.3.64 yStepLeft

In dual Y-axis mode, Manually set the value step between the left Y-axis labels and grid lines (See: yStep)

9.3.65 yStepRight

In dual Y-axis mode, Manually set the value step between the right Y-axis labels and grid lines (See: yStep)

9.3.66 yUnitSystem

Default: si

Set the unit system for compacting Y-axis values (e.g. 23,000,000 becomes 23M). Value can be one of:

`si` Use si units (powers of 1000) - K, M, G, T, P

`binary` Use binary units (powers of 1024) - Ki, Mi, Gi, Ti, Pi

`none` Dont compact values, display the raw number

CHAPTER

TEN

FUNCTIONS

Functions are used to transform, combine, and perform computations on *series* data. Functions are applied using the Composer interface or by manipulating the `target` parameters in the *Render API*.

10.1 Usage

Most functions are applied to one *series list*. Functions with the parameter `*seriesLists` can take an arbitrary number of series lists. To pass multiple series lists to a function which only takes one, use the `group()` function.

10.2 List of functions

absolute(*seriesList*)
Takes one metric or a wildcard seriesList and applies the mathematical abs function to each datapoint transforming it to its absolute value.

Example:

```
&target=absolute(Server.instance01.threads.busy)
&target=absolute(Server.instance*.threads.busy)
```

alias(*seriesList*, *newName*)
Takes one metric or a wildcard seriesList and a string in quotes. Prints the string instead of the metric name in the legend.

```
&target=alias(Sales.widgets.largeBlue,"Large Blue Widgets")
```

aliasByMetric(*seriesList*)
Takes a seriesList and applies an alias derived from the base metric name.

```
&target=aliasByMetric(carbon.agents.graphite.creates)
```

aliasByNode(*seriesList*, **nodes*)
Takes a seriesList and applies an alias derived from one or more "node" portion/s of the target name. Node indices are 0 indexed.

```
&target=aliasByNode(ganglia.*.cpu.load5,1)
```

aliasSub(*seriesList*, *search*, *replace*)
Runs series names through a regex search/replace.

```
&target=aliasSub(ip.*TCP*,"^.*TCP(\d+)","\1")
```

alpha (*seriesList*, *alpha*)
Assigns the given alpha transparency setting to the series. Takes a float value between 0 and 1.

areaBetween (*seriesList*)
Draws the area in between the two series in seriesList

asPercent (*seriesList*, *total=None*)
Calculates a percentage of the total of a wildcard series. If *total* is specified, each series will be calculated as a percentage of that total. If *total* is not specified, the sum of all points in the wildcard series will be used instead.

The *total* parameter may be a single series or a numeric value.

Example:

```
&target=asPercent(Server01.connections.{failed,succeeded}, Server01.connections.attempted)
&target=asPercent(apache01.threads.busy,1500)
&target=asPercent(Server01.cpu.*.jiffies)
```

averageAbove (*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N. Out of all metrics passed, draws only the metrics with an average value above N for the time period specified.

Example:

```
&target=averageAbove(server*.instance*.threads.busy,25)
```

Draws the servers with average values above 25.

averageBelow (*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N. Out of all metrics passed, draws only the metrics with an average value below N for the time period specified.

Example:

```
&target=averageBelow(server*.instance*.threads.busy,25)
```

Draws the servers with average values below 25.

averageSeries (**seriesLists*)
Short Alias: avg()

Takes one metric or a wildcard seriesList. Draws the average value of all metrics passed at each time.

Example:

```
&target=averageSeries(company.server.*.threads.busy)
```

averageSeriesWithWildcards (*seriesList*, **position*)
Call averageSeries after inserting wildcards at the given position(s).

Example:

```
&target=averageSeriesWithWildcards(host.cpu-[0-7].cpu-{user,system}.value, 1)
```

This would be the equivalent of `target=averageSeries(host.*.cpu-user.value)&target=averageSeries(`

cactiStyle (*seriesList*, *system=None*)
Takes a series list and modifies the aliases to provide column aligned output with Current, Max, and Min values in the style of cacti. Optonally takes a "system" value to apply unit formatting in the same style as the Y-axis. NOTE: column alignment only works with monospace fonts such as terminus.

```
&target=cactiStyle(ganglia.*.net.bytes_out,"si")
```

color (*seriesList*, *theColor*)
Assigns the given color to the seriesList

Example:

```
&target=color(collectd.hostname.cpu.0.user, 'green')
&target=color(collectd.hostname.cpu.0.system, 'ff0000')
&target=color(collectd.hostname.cpu.0.idle, 'gray')
&target=color(collectd.hostname.cpu.0.idle, '6464ffaa')
```

consolidateBy (*seriesList*, *consolidationFunc*)
Takes one metric or a wildcard seriesList and a consolidation function name.

Valid function names are 'sum', 'average', 'min', and 'max'

When a graph is drawn where width of the graph size in pixels is smaller than the number of datapoints to be graphed, Graphite consolidates the values to to prevent line overlap. The consolidateBy() function changes the consolidation function from the default of 'average' to one of 'sum', 'max', or 'min'. This is especially useful in sales graphs, where fractional values make no sense and a 'sum' of consolidated values is appropriate.

```
&target=consolidateBy(Sales.widgets.largeBlue, 'sum')
&target=consolidateBy(Servers.web01.sda1.free_space, 'max')
```

constantLine (*value*)
Takes a float F.

Draws a horizontal line at value F across the graph.

Example:

```
&target=constantLine(123.456)
```

countSeries (**seriesLists*)
Draws a horizontal line representing the number of nodes found in the seriesList.

```
&target=countSeries(carbon.agents.*.*)
```

cumulative (*seriesList*)
Takes one metric or a wildcard seriesList.

Sets the consolidation function to 'sum' for the given metric seriesList.

Alias for `consolidateBy(series, 'sum')`

```
&target=cumulative(Sales.widgets.largeBlue)
```

currentAbove (*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N. Out of all metrics passed, draws only the metrics whose value is above N at the end of the time period specified.

Example:

```
&target=currentAbove(server*.instance*.threads.busy,50)
```

Draws the servers with more than 50 busy threads.

currentBelow (*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N. Out of all metrics passed, draws only the metrics whose value is below N at the end of the time period specified.

Example:

```
&target=currentBelow(server*.instance*.threads.busy,3)
```

Draws the servers with less than 3 busy threads.

dashed(**seriesList*)
Takes one metric or a wildcard seriesList, followed by a float F.

Draw the selected metrics with a dotted line with segments of length F If omitted, the default length of the segments is 5.0

Example:

```
&target=dashed(server01.instance01.memory.free,2.5)
```

derivative(*seriesList*)
This is the opposite of the integral function. This is useful for taking a running total metric and calculating the delta between subsequent data points.

This function does not normalize for periods of time, as a true derivative would.

Example:

```
&target=derivative(company.server.application01.ifconfig.TXPackets)
```

Each time you run ifconfig, the RX and TXPackets are higher (assuming there is network traffic.) By applying the derivative function, you can get an idea of the packets per minute sent or received, even though you're only recording the total.

diffSeries(**seriesLists*)
Can take two or more metrics, or a single metric and a constant. Subtracts parameters 2 through n from parameter 1.

Example:

```
&target=diffSeries(service.connections.total,service.connections.failed)
&target=diffSeries(service.connections.total,5)
```

divideSeries(*dividendSeriesList*, *divisorSeriesList*)
Takes a dividend metric and a divisor metric and draws the division result. A constant may *not* be passed. To divide by a constant, use the scale() function (which is essentially a multiplication operation) and use the inverse of the dividend. (Division by 8 = multiplication by 1/8 or 0.125)

Example:

```
&target=divideSeries(Series.dividends,Series.divisors)
```

drawAsInfinite(*seriesList*)
Takes one metric or a wildcard seriesList. If the value is zero, draw the line at 0. If the value is above zero, draw the line at infinity. If the value is null or less than zero, do not draw the line.

Useful for displaying on/off metrics, such as exit codes. (0 = success, anything else = failure.)

Example:

```
drawAsInfinite(Testing.script.exitCode)
```

events(**tags*)
Returns the number of events at this point in time. Usable with drawAsInfinite.

Example:

```
&target=events("tag-one", "tag-two")
&target=events("*")
```

Returns all events tagged as "tag-one" and "tag-two" and the second one returns all events.

exclude (*seriesList*, *pattern*)
Takes a metric or a wildcard seriesList, followed by a regular expression in double quotes. Excludes metrics that match the regular expression.

Example:

```
&target=exclude(servers*.instance*.threads.busy,"server02")
```

group (**seriesLists*)
Takes an arbitrary number of seriesLists and adds them to a single seriesList. This is used to pass multiple seriesLists to a function which only takes one

groupByNode (*seriesList*, *nodeNum*, *callback*)
Takes a serieslist and maps a callback to subgroups within as defined by a common node

```
&target=groupByNode(ganglia.by-function.*.*.cpu.load5,2,"sumSeries")

Would return multiple series which are each the result of applying the "sumSeries" function
to groups joined on the second node (0 indexed) resulting in a list of targets like
sumSeries(ganglia.by-function.server1.*.cpu.load5),sumSeries(ganglia.by-function.server2.*.cpu.l
```

highestAverage (*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N. Out of all metrics passed, draws only the top N metrics with the highest average value for the time period specified.

Example:

```
&target=highestAverage(server*.instance*.threads.busy,5)
```

Draws the top 5 servers with the highest average value.

highestCurrent (*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N. Out of all metrics passed, draws only the N metrics with the highest value at the end of the time period specified.

Example:

```
&target=highestCurrent(server*.instance*.threads.busy,5)
```

Draws the 5 servers with the highest busy threads.

highestMax (*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N.

Out of all metrics passed, draws only the N metrics with the highest maximum value in the time period specified.

Example:

```
&target=highestMax(server*.instance*.threads.busy,5)
```

Draws the top 5 servers who have had the most busy threads during the time period specified.

hitcount (*seriesList*, *intervalString*, *alignToInterval=False*)
Estimate hit counts from a list of time series.

This function assumes the values in each time series represent hits per second. It calculates hits per some larger interval such as per day or per hour. This function is like summarize(), except that it compensates automatically

for different time scales (so that a similar graph results from using either fine-grained or coarse-grained records) and handles rarely-occurring events gracefully.

holtWintersAberration (*seriesList*, *delta=3*)
Performs a Holt-Winters forecast using the series as input data and plots the positive or negative deviation of the series data from the forecast.

holtWintersConfidenceArea (*seriesList*, *delta=3*)
Performs a Holt-Winters forecast using the series as input data and plots the area between the upper and lower bands of the predicted forecast deviations.

holtWintersConfidenceBands (*seriesList*, *delta=3*)
Performs a Holt-Winters forecast using the series as input data and plots upper and lower bands with the predicted forecast deviations.

holtWintersForecast (*seriesList*)
Performs a Holt-Winters forecast using the series as input data. Data from one week previous to the series is used to bootstrap the initial forecast.

identity (*name*)
Identity function: Returns datapoints where the value equals the timestamp of the datapoint. Useful when you have another series where the value is a timestamp, and you want to compare it to the time of the datapoint, to render an age

Example:

```
&target=identity("The.time.series")
```

This would create a series named "The.time.series" that contains points where x(t) == t.

integral (*seriesList*)
This will show the sum over time, sort of like a continuous addition function. Useful for finding totals or trends in metrics that are collected per minute.

Example:

```
&target=integral(company.sales.perMinute)
```

This would start at zero on the left side of the graph, adding the sales each minute, and show the total sales for the time period selected at the right side, (time now, or the time specified by '&until=').

keepLastValue (*seriesList*, *limit=inf*)
Takes one metric or a wildcard seriesList, and optionally a limit to the number of 'None' values to skip over. Continues the line with the last received value when gaps ('None' values) appear in your data, rather than breaking your line.

Example:

```
&target=keepLastValue(Server01.connections.handled)
&target=keepLastValue(Server01.connections.handled, 10)
```

legendValue (*seriesList*, **valueTypes*)
Takes one metric or a wildcard seriesList and a string in quotes. Appends a value to the metric name in the legend. Currently one or several of: *last*, *avg*, *total*, *min*, *max*. The last argument can be *si* (default) or *binary*, in that case values will be formatted in the corresponding system.

&target=legendValue(Sales.widgets.largeBlue, 'avg', 'max', 'si')

limit (*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N.

Only draw the first N metrics. Useful when testing a wildcard in a metric.

Example:

```
&target=limit(server*.instance*.memory.free,5)
```

Draws only the first 5 instance's memory free.

lineWidth(*seriesList*, *width*)
Takes one metric or a wildcard seriesList, followed by a float F.

Draw the selected metrics with a line width of F, overriding the default value of 1, or the &lineWidth=X.X parameter.

Useful for highlighting a single metric out of many, or having multiple line widths in one graph.

Example:

```
&target=lineWidth(server01.instance01.memory.free,5)
```

logarithm(*seriesList*, *base=10*)
Takes one metric or a wildcard seriesList, a base, and draws the y-axis in logarithmic format. If base is omitted, the function defaults to base 10.

Example:

```
&target=log(carbon.agents.hostname.avgUpdateTime,2)
```

lowestAverage(*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N. Out of all metrics passed, draws only the bottom N metrics with the lowest average value for the time period specified.

Example:

```
&target=lowestAverage(server*.instance*.threads.busy,5)
```

Draws the bottom 5 servers with the lowest average value.

lowestCurrent(*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by an integer N. Out of all metrics passed, draws only the N metrics with the lowest value at the end of the time period specified.

Example:

```
&target=lowestCurrent(server*.instance*.threads.busy,5)
```

Draws the 5 servers with the least busy threads right now.

maxSeries(**seriesLists*)
Takes one metric or a wildcard seriesList. For each datapoint from each metric passed in, pick the maximum value and graph it.

Example:

```
&target=maxSeries(Server*.connections.total)
```

maximumAbove(*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by a constant n. Draws only the metrics with a maximum value above n.

Example:

```
&target=maximumAbove(system.interface.eth*.packetsSent,1000)
```

This would only display interfaces which sent more than 1000 packets/min.

maximumBelow(*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by a constant n. Draws only the metrics with a maximum value below n.

Example:

```
&target=maximumBelow(system.interface.eth*.packetsSent,1000)
```

This would only display interfaces which sent less than 1000 packets/min.

minSeries(**seriesLists*)
Takes one metric or a wildcard seriesList. For each datapoint from each metric passed in, pick the minimum value and graph it.

Example:

```
&target=minSeries(Server*.connections.total)
```

minimumAbove(*seriesList*, *n*)
Takes one metric or a wildcard seriesList followed by a constant n. Draws only the metrics with a minimum value above n.

Example:

```
&target=minimumAbove(system.interface.eth*.packetsSent,1000)
```

This would only display interfaces which sent more than 1000 packets/min.

mostDeviant(*n*, *seriesList*)
Takes an integer N followed by one metric or a wildcard seriesList. Draws the N most deviant metrics. To find the deviant, the average across all metrics passed is determined, and then the average of each metric is compared to the overall average.

Example:

```
&target=mostDeviant(5, server*.instance*.memory.free)
```

Draws the 5 instances furthest from the average memory free.

movingAverage(*seriesList*, *windowSize*)
Graphs the moving average of a metric (or metrics) over a fixed number of past points, or a time interval.

Takes one metric or a wildcard seriesList followed by a number N of datapoints or a quoted string with a length of time like '1hour' or '5min' (See `from / until` in the render_api_ for examples of time formats). Graphs the average of the preceeding datapoints for each point on the graph. All previous datapoints are set to None at the beginning of the graph.

Example:

```
&target=movingAverage(Server.instance01.threads.busy,10)
&target=movingAverage(Server.instance*.threads.idle,'5min')
```

movingMedian(*seriesList*, *windowSize*)
Graphs the moving median of a metric (or metrics) over a fixed number of past points, or a time interval.

Takes one metric or a wildcard seriesList followed by a number N of datapoints or a quoted string with a length of time like '1hour' or '5min' (See `from / until` in the render_api_ for examples of time formats). Graphs the median of the preceeding datapoints for each point on the graph. All previous datapoints are set to None at the beginning of the graph.

Example:

```
&target=movingMedian(Server.instance01.threads.busy,10)
&target=movingMedian(Server.instance*.threads.idle,'5min')
```

multiplySeries(**seriesLists*)
Takes two or more series and multiplies their points. A constant may not be used. To multiply by a constant, use the scale() function.

Example:

```
&target=multiplySeries(Series.dividends,Series.divisors)
```

nPercentile(*seriesList*, *n*)
Returns n-percent of each series in the seriesList.

nonNegativeDerivative(*seriesList*, *maxValue=None*)
Same as the derivative function above, but ignores datapoints that trend down. Useful for counters that increase for a long time, then wrap or reset. (Such as if a network interface is destroyed and recreated by unloading and re-loading a kernel module, common with USB / WiFi cards.

Example:

```
&target=nonNegativederivative(company.server.application01.ifconfig.TXPackets)
```

offset(*seriesList*, *factor*)
Takes one metric or a wildcard seriesList followed by a constant, and adds the constant to each datapoint.

Example:

```
&target=offset(Server.instance01.threads.busy,10)
```

percentileOfSeries(*seriesList*, *n*, *interpolate=False*)
percentileOfSeries returns a single series which is composed of the n-percentile values taken across a wildcard series at each point. Unless *interpolate* is set to True, percentile values are actual values contained in one of the supplied series.

randomWalkFunction(*name*)
Short Alias: randomWalk()

Returns a random walk starting at 0. This is great for testing when there is no real data in whisper.

Example:

```
&target=randomWalk("The.time.series")
```

This would create a series named "The.time.series" that contains points where x(t) == x(t-1)+random()-0.5, and x(0) == 0.

rangeOfSeries(**seriesLists*)
Takes a wildcard seriesList. Distills down a set of inputs into the range of the series

Example:

```
&target=rangeOfSeries(Server*.connections.total)
```

removeAbovePercentile(*seriesList*, *n*)
Removes data above the nth percentile from the series or list of series provided. Values above this percentile are assigned a value of None.

removeAboveValue(*seriesList*, *n*)
Removes data above the given threshold from the series or list of series provided. Values above this threshole are assigned a value of None

removeBelowPercentile (*seriesList*, *n*)
Removes data below the nth percentile from the series or list of series provided. Values below this percentile are assigned a value of None.

removeBelowValue (*seriesList*, *n*)
Removes data below the given threshold from the series or list of series provided. Values below this threshole are assigned a value of None

scale (*seriesList*, *factor*)
Takes one metric or a wildcard seriesList followed by a constant, and multiplies the datapoint by the constant provided at each point.

Example:

```
&target=scale(Server.instance01.threads.busy,10)
&target=scale(Server.instance*.threads.busy,10)
```

scaleToSeconds (*seriesList*, *seconds*)
Takes one metric or a wildcard seriesList and returns "value per seconds" where seconds is a last argument to this functions.

Useful in conjunction with derivative or integral function if you want to normalize its result to a known resolution for arbitrary retentions

secondYAxis (*seriesList*)
Graph the series on the secondary Y axis.

sinFunction (*name*, *amplitude=1*)
Short Alias: sin()

Just returns the sine of the current time. The optional amplitude parameter changes the amplitude of the wave.

Example:

```
&target=sin("The.time.series", 2)
```

This would create a series named "The.time.series" that contains sin(x)*2.

smartSummarize (*seriesList*, *intervalString*, *func='sum'*, *alignToFrom=False*)
Smarter experimental version of summarize.

The alignToFrom parameter has been deprecated, it no longer has any effect. Alignment happens automatically for days, hours, and minutes.

sortByMaxima (*seriesList*)
Takes one metric or a wildcard seriesList.

Sorts the list of metrics by the maximum value across the time period specified. Useful with the &areaMode=all parameter, to keep the lowest value lines visible.

Example:

```
&target=sortByMaxima(server*.instance*.memory.free)
```

sortByMinima (*seriesList*)
Takes one metric or a wildcard seriesList.

Sorts the list of metrics by the lowest value across the time period specified.

Example:

```
&target=sortByMinima(server*.instance*.memory.free)
```

stacked(*seriesLists*, *stackName='__DEFAULT__'*)
Takes one metric or a wildcard seriesList and change them so they are stacked. This is a way of stacking just a couple of metrics without having to use the stacked area mode (that stacks everything). By means of this a mixed stacked and non stacked graph can be made

It can also take an optional argument with a name of the stack, in case there is more than one, e.g. for input and output metrics.

Example:

```
&target=stacked(company.server.application01.ifconfig.TXPackets, 'tx')
```

stddevSeries(**seriesLists*)
Takes one metric or a wildcard seriesList. Draws the standard deviation of all metrics passed at each time.

Example:

```
&target=stddevSeries(company.server.*.threads.busy)
```

stdev(*seriesList*, *points*, *windowTolerance=0.1*)
Takes one metric or a wildcard seriesList followed by an integer N. Draw the Standard Deviation of all metrics passed for the past N datapoints. If the ratio of null points in the window is greater than windowTolerance, skip the calculation. The default for windowTolerance is 0.1 (up to 10% of points in the window can be missing). Note that if this is set to 0.0, it will cause large gaps in the output anywhere a single point is missing.

Example:

```
&target=stdev(server*.instance*.threads.busy,30)
&target=stdev(server*.instance*.cpu.system,30,0.0)
```

substr(*seriesList*, *start=0*, *stop=0*)
Takes one metric or a wildcard seriesList followed by 1 or 2 integers. Assume that the metric name is a list or array, with each element separated by dots. Prints n - length elements of the array (if only one integer n is passed) or n - m elements of the array (if two integers n and m are passed). The list starts with element 0 and ends with element (length - 1).

Example:

```
&target=substr(carbon.agents.hostname.avgUpdateTime,2,4)
```

The label would be printed as "hostname.avgUpdateTime".

sumSeries(**seriesLists*)
Short form: sum()

This will add metrics together and return the sum at each datapoint. (See integral for a sum over time)

Example:

```
&target=sum(company.server.application*.requestsHandled)
```

This would show the sum of all requests handled per minute (provided requestsHandled are collected once a minute). If metrics with different retention rates are combined, the coarsest metric is graphed, and the sum of the other metrics is averaged for the metrics with finer retention rates.

sumSeriesWithWildcards(*seriesList*, **position*)
Call sumSeries after inserting wildcards at the given position(s).

Example:

```
&target=sumSeriesWithWildcards(host.cpu-[0-7].cpu-{user,system}.value, 1)
```

This would be the equivalent of `target=sumSeries(host.*.cpu-user.value)&target=sumSeries(host.*.cp`

summarize(*seriesList*, *intervalString*, *func='sum'*, *alignToFrom=False*)
Summarize the data into interval buckets of a certain size.

By default, the contents of each interval bucket are summed together. This is useful for counters where each increment represents a discrete event and retrieving a "per X" value requires summing all the events in that interval.

Specifying 'avg' instead will return the mean for each bucket, which can be more useful when the value is a gauge that represents a certain value in time.

'max', 'min' or 'last' can also be specified.

By default, buckets are caculated by rounding to the nearest interval. This works well for intervals smaller than a day. For example, 22:32 will end up in the bucket 22:00-23:00 when the interval=1hour.

Passing alignToFrom=true will instead create buckets starting at the from time. In this case, the bucket for 22:32 depends on the from time. If from=6:30 then the 1hour bucket for 22:32 is 22:30-23:30.

Example:

```
&target=summarize(counter.errors, "1hour") # total errors per hour
&target=summarize(nonNegativeDerivative(gauge.num_users), "1week") # new users per week
&target=summarize(queue.size, "1hour", "avg") # average queue size per hour
&target=summarize(queue.size, "1hour", "max") # maximum queue size during each hour
&target=summarize(metric, "13week", "avg", true)&from=midnight+20100101 # 2010 Q1-4
```

threshold(*value*, *label=None*, *color=None*)
Takes a float F, followed by a label (in double quotes) and a color. (See `bgcolor` in the render_api_ for valid color names & formats.)

Draws a horizontal line at value F across the graph.

Example:

```
&target=threshold(123.456, "omgwtfbbq", red)
```

timeFunction(*name*)
Short Alias: time()

Just returns the timestamp for each X value. T

Example:

```
&target=time("The.time.series")
```

This would create a series named "The.time.series" that contains in Y the same value (in seconds) as X.

timeShift(*seriesList*, *timeShift*, *resetEnd=True*)
Takes one metric or a wildcard seriesList, followed by a quoted string with the length of time (See `from / until` in the render_api_ for examples of time formats).

Draws the selected metrics shifted in time. If no sign is given, a minus sign (-) is implied which will shift the metric back in time. If a plus sign (+) is given, the metric will be shifted forward in time.

Will reset the end date range automatically to the end of the base stat unless resetEnd is False. Example case is when you timeshift to last week and have the graph date range set to include a time in the future, will limit this timeshift to pretend ending at the current time. If resetEnd is False, will instead draw full range including future time.

Useful for comparing a metric against itself at a past periods or correcting data stored at an offset.

Example:

```
&target=timeShift(Sales.widgets.largeBlue,"7d")
&target=timeShift(Sales.widgets.largeBlue,"-7d")
&target=timeShift(Sales.widgets.largeBlue,"+1h")
```

timeStack (*seriesList*, *timeShiftUnit*, *timeShiftStart*, *timeShiftEnd*)

Takes one metric or a wildcard seriesList, followed by a quoted string with the length of time (See `from / until` in the render_api_ for examples of time formats). Also takes a start multiplier and end multiplier for the length of time

create a seriesList which is composed the orginal metric series stacked with time shifts starting time shifts from the start multiplier through the end multiplier

Useful for looking at history, or feeding into seriesAverage or seriesStdDev

Example:

```
&target=timeStack(Sales.widgets.largeBlue,"1d",0,7)    # create a series for today and each of t
```

transformNull (*seriesList*, *default=0*)

Takes a metric or wild card seriesList and an optional value to transform Nulls to. Default is 0. This method compliments drawNullAsZero flag in graphical mode but also works in text only mode. Example:

```
&target=transformNull(webapp.pages.*.views,-1)
```

This would take any page that didn't have values and supply negative 1 as a default. Any other numeric value may be used as well.

useSeriesAbove (*seriesList*, *value*, *search*, *replace*)

Compares the maximum of each series against the given *value*. If the series maximum is greater than *value*, the regular expression search and replace is applied against the series name to plot a related metric

e.g. given useSeriesAbove(ganglia.metric1.reqs,10,'reqs','time'), the response time metric will be plotted only when the maximum value of the corresponding request/s metric is > 10

```
&target=useSeriesAbove(ganglia.metric1.reqs,10,"reqs","time")
```

CHAPTER

ELEVEN

THE DASHBOARD UI

...

CHAPTER

TWELVE

THE WHISPER DATABASE

Whisper is a fixed-size database, similar in design and purpose to RRD (round-robin-database). It provides fast, reliable storage of numeric data over time. Whisper allows for higher resolution (seconds per point) of recent data to degrade into lower resolutions for long-term retention of historical data.

12.1 Data Points

Data points in Whisper are stored on-disk as big-endian double-precision floats. Each value is paired with a timestamp in seconds since the UNIX Epoch (01-01-1970). The data value is parsed by the Python float() function and as such behaves in the same way for special strings such as `'inf'`. Maximum and minimum values are determined by the Python interpreter's allowable range for float values which can be found by executing:

```
python -c 'import sys; print sys.float_info'
```

12.2 Archives: Retention and Precision

Whisper databases contain one or more archives, each with a specific data resolution and retention (defined in number of points or max timestamp age). Archives are ordered from the highest-resolution and shortest retention archive to the lowest-resolution and longest retention period archive.

To support accurate aggregation from higher to lower resolution archives, the precision of a longer retention archive must be divisible by precision of next lower retention archive. For example, an archive with 1 data point every 60 seconds can have a lower-resolution archive following it with a resolution of 1 data point every 300 seconds because 60 cleanly divides 300. In contrast, a 180 second precision (3 minutes) could not be followed by a 600 second precision (10 minutes) because the ratio of points to be propagated from the first archive to the next would be 3 1/3 and Whisper will not do partial point interpolation.

The total retention time of the database is determined by the archive with the highest retention as the time period covered by each archive is overlapping (see Multi-Archive Storage and Retrieval Behavior). That is, a pair of archives with retentions of 1 month and 1 year will not provide 13 months of data storage as may be guessed. Instead, it will provide 1 year of storage - the length of it's longest archive.

12.3 Rollup Aggregation

Whisper databases with more than a single archive need a strategy to collapse multiple data points for when the data rolls up a lower precision archive. By default, an average function is used. Available aggregation methods are: * average * sum * last * max * min

12.4 Multi-Archive Storage and Retrieval Behavior

When Whisper writes to a database with multiple archives, the incoming data point is written to all archives at once. The data point will be written to the lowest resolution archive as-is, and will be aggregated by the configured aggregation method (see Rollup Aggregation) and placed into each of the higher-retention archives.

When data is retrieved (scoped by a time range), the first archive which can satisfy the entire time period is used. If the time period overlaps an archive boundary, the lower-resolution archive will be used. This allows for a simpler behavior while retrieving data as the data's resolution is consistent through an entire returned series.

12.5 Disk Space Efficiency

Whisper is somewhat inefficient in its usage of disk space because of certain design choices:

Each data point is stored with its timestamp Rather than a timestamp being inferred from its position in the archive, timestamps are stored with each point. The timestamps are used during data retrieval to check the validity of the data point. If a timestamp does not match the expected value for its position relative to the beginning of the requested series, it is known to be out of date and a null value is returned

Archives overlap time periods During the write of a data point, Whisper stores the same data in all archives at once (see Multi-Archive Storage and Retrieval Behavior). Implied by this behavior is that all archives store from now until each of their retention times. Because of this, lower-resolution archives should be configured to significantly lower resolution and higher retentions than their higher-resolution counterparts so as to reduce the overlap.

All time-slots within an archive take up space whether or not a value is stored While Whisper allows for reliable storage of irregular updates, it is most space efficient when data points are stored at every update interval. This behavior is a consequence of the fixed-size design of the database and allows the reading and writing of series data to be performed in a single contiguous disk operation (for each archive in a database).

12.6 Differences Between Whisper and RRD

RRD can not take updates to a time-slot prior to its most recent update This means that there is no way to back-fill data in an RRD series. Whisper does not have this limitation, and this makes importing historical data into Graphite much more simple and easy

RRD was not designed with irregular updates in mind In many cases (depending on configuration) if an update is made to an RRD series but is not followed up by another update soon, the original update will be lost. This makes it less suitable for recording data such as operational metrics (e.g. code pushes)

Whisper requires that metric updates occur at the same interval as the finest resolution storage archive This pushes the onus of aggregating values to fit into the finest precision archive to the user rather than the database. It also means that updates are written immediately into the finest precision archive rather than being staged first for aggregation and written later (during a subsequent write operation) as they are in RRD.

12.7 Performance

Whisper is fast enough for most purposes. It is slower than RRDtool primarily as a consequence of Whisper being written in Python, while RRDtool is written in C. The speed difference between the two in practice is quite small as much effort was spent to optimize Whisper to be as close to RRDtool's speed as possible. Testing has shown that update operations take anywhere from 2 to 3 times as long as RRDtool, and fetch operations take anywhere from 2 to

5 times as long. In practice the actual difference is measured in hundreds of microseconds (10^-4) which means less than a millisecond difference for simple cases.

12.8 Database Format

Whisper-File	*Header,Data*			
	Header	*Metadata,ArchiveInfo+*		
		Metadata	aggregationType,maxRetention,xFilesFactor,archiveCount	
		ArchiveInfo	Offset,SecondsPerPoint,Points	
	Data	*Archive+*		
		Archive	*Point+*	
			Point	timestamp,value

Data types in Python's struct format:

Metadata	`!2LfL`
ArchiveInfo	`!3L`
Point	`!Ld`

CHAPTER

THIRTEEN

GRAPHITE TERMINOLOGY

Graphite uses many terms that can have ambiguous meaning. The following definitions are what these terms mean in the context of Graphite.

datapoint A *value* stored at a *timestamp bucket*. If no value is recorded at a particular timestamp bucket in a *series*, the value will be None (null).

function A time-series function which transforms, combines, or performs computations on one or more *series*. See *Functions*

metric See *series*

metric series See *series*

precision See *resolution*

resolution The number of seconds per datapoint in a *series*. Series are created with a resolution which determines how often a *datapoint* may be stored. This resolution is represented as the number of seconds in time that each datapoint covers. A series which stores one datapoint per minute has a resolution of 60 seconds. Similarly, a series which stores one datapoint per second has a resolution of 1 second.

retention The number of datapoints retained in a *series*. Alternatively: The length of time datapoints are stored in a series.

series A named set of datapoints. A series is identified by a unique name, which is composed of elements separated by periods (`.`) which are used to display the collection of series into a heirarchical tree. A series storing system load average on a server called `apache02` in datacenter `metro_east` might be named as `metro_east.servers.apache02.system.load_average`

series list A series name or wildcard which matches one or more *series*. Series lists are received by *functions* as a list of matching series. From a user perspective, a series list is merely the name of a metric. For example, each of these would be considered a single series list:

- `metro_east.servers.apache02.system.load_average.1_min`,
- `metro_east.servers.apache0{1,2,3}.system.load_average.1_min`
- `metro_east.servers.apache01.system.load_average.*`

target A source of data used as input for a Graph. A target can be a single metric name, a metric wildcard, or either of these enclosed within one or more *functions*

timestamp A point in time in which *values* can be associated. Time in Graphite is represented as epoch time with a maximum resolution of 1-second.

timestamp bucket A *timestamp* after rounding down to the nearest multiple of a *series's resolution*.

value A numeric or null value. Values are stored as double-precision floats. Values are parsed using the python `float()` constructor and can also be None (null). The range and precision of values is system dependant and can be found by executing (with Python 2.6 or later):: python -c 'import sys; print sys.float_info'

CHAPTER

FOURTEEN

TOOLS THAT WORK WITH GRAPHITE

14.1 Backstop

Backstop is a simple endpoint for submitting metrics to Graphite. It accepts JSON data via HTTP POST and proxies the data to one or more Carbon/Graphite listeners.

14.2 Bucky

Bucky is a small service implemented in Python for collecting and translating metrics for Graphite. It can current collect metric data from CollectD daemons and from StatsD clients.

14.3 collectd

collectd is a daemon which collects system performance statistics periodically and provides mechanisms to store the values in a variety of ways, including RRD. To send collectd metrics into carbon/graphite, use collectd's write-graphite plugin (available as of 5.1). Other options include:

- Jordan Sissel's node collectd-to-graphite proxy
- Joe Miller's perl collectd-graphite plugin
- Gregory Szorc's python collectd-carbon plugin
- Paul J. Davis's Bucky service

Graphite can also read directly from collectd's RRD files. RRD files can simply be added to `STORAGE_DIR/rrd` (as long as directory names and files do not contain any `.` characters). For example, collectd's `host.name/load/load.rrd` can be symlinked to `rrd/collectd/host_name/load/load.rrd` to graph `collectd.host_name.load.load.{short,mid,long}term`.

14.4 Collectl

Collectl is a collection tool for system metrics that can be run both interactively and as a daemon and has support for collecting from a broad set of subsystems. Collectl includes a Graphite interface which allows data to easily be fed to Graphite for storage.

14.5 Charcoal

Charcoal is a simple Sinatra dashboarding frontend for Graphite or any other system status service which can generate images directly from a URL. Charcoal configuration is driven by a YAML config file.

14.6 Descartes

Descartes is a Sinatra-based dashboard that allows users to correlate multiple metrics in a single chart, review long-term trends across one or more charts, and to collaborate with other users through a combination of shared dashboards and rich layouts.

14.7 Diamond

Diamond is a Python daemon that collects system metrics and publishes them to Graphite. It is capable of collecting cpu, memory, network, I/O, load and disk metrics. Additionally, it features an API for implementing custom collectors for gathering metrics from almost any source.

14.8 Evenflow

Evenflow is a simple service for submitting sFlow datagrams to Graphite. It accepts sFlow datagrams from multiple network devices and proxies the data to a Carbon listener. Currently only Generic Interface Counters are supported. All other message types are discarded.

14.9 Ganglia

Ganglia is a scalable distributed monitoring system for high-performance computing systems such as clusters and Grids. It collects system performance metrics and stores them in RRD, but now there is an add-on that allows Ganglia to send metrics directly to Graphite. Further integration work is underway.

14.10 GDash

Gdash is a simple Graphite dashboard built using Twitters Bootstrap driven by a small DSL.

14.11 Giraffe

Giraffe is a Graphite real-time dashboard based on Rickshaw and requires no server backend. Inspired by Gdash, Tasseo and Graphene it mixes features from all three into a slightly different animal.

14.12 Graphitus

graphitus is a client side dashboard for graphite built using bootstrap and underscore.js.

14.13 Graph-Explorer

Graph-Explorer is a graphite dashboard which uses plugins to add tags and metadata to metrics and a query language with lets you filter through them and compose/manipulate graphs on the fly. Also aims for high interactivity using TimeseriesWidget and minimal hassle to set up and get running.

14.14 Graphene

Graphene is a Graphite dashboard toolkit based on D3.js and Backbone.js which was made to offer a very aesthetic realtime dashboard. Graphene provides a solution capable of displaying thousands upon thousands of datapoints all updated in realtime.

14.15 Graphite-relay

Graphite-relay is a fast Graphite relay written in Scala with the Netty framework.

14.16 Graphite-Tattle

Graphite-Tattle is a self-service dashboard frontend for Graphite and Ganglia.

14.17 Graphiti

Graphiti is a powerful dashboard front end with a focus on ease of access, ease of recovery and ease of tweaking and manipulation.

14.18 Graphitoid

Graphitoid is an Android app which allows one to browse and display Graphite graphs on an Android device.

14.19 Graphios

Graphios is a small Python daemon to send Nagios performance data (perfdata) to Graphite.

14.20 Graphitejs

Graphitejs is a jQuery plugin for easily making and displaying graphs and updating them on the fly using the Graphite URL api.

14.21 Graphsky

Graphsky is flexible and easy to configure PHP based dashboard. It uses JSON template files to build graphs and specify which graphs need to be displayed when, similar to Ganglia-web. Just like Ganglia, it uses a hierarchial structure: Environment/Cluster/Host/Metric to be able to display overview graphs and host-specific metrics. It communicates directly to the Graphite API to determine which Environments, Clusters, Hosts and Metrics are currently stored in Graphite.

14.22 Grockets

Grockets is a node.js application which provides streaming JSON data over HTTP from Graphite.

14.23 HoardD

HoardD is a Node.js app written in CoffeeScript to send data from servers to Graphite, much like collectd does, but aimed at being easier to expand and with less footprint. It comes by default with basic collectors plus Redis and MySQL metrics, and can be expanded with Javascript or CoffeeScript.

14.24 Host sFlow

Host sFlow is an open source implementation of the sFlow protocol (http://www.sflow.org), exporting a standard set of host cpu, memory, disk and network I/O metrics. The sflow2graphite utility converts sFlow to Graphite's plaintext protocol, allowing Graphite to receive sFlow metrics.

14.25 hubot-scripts

Hubot is a Campfire bot written in Node.js and CoffeeScript. The related hubot-scripts project includes a Graphite script which supports searching and displaying saved graphs from the Composer directory in your Campfire rooms.

14.26 jmxtrans

jmxtrans is a powerful tool that performs JMX queries to collect metrics from Java applications. It is requires very little configuration and is capable of sending metric data to several backend applications, including Graphite.

14.27 Ledbetter

Ledbetter is a simple script for gathering Nagios problem statistics and submitting them to Graphite. It focuses on summary (overall, servicegroup and hostgroup) statistics and writes them to the nagios.problems metrics namespace within Graphite.

14.28 Logster

Logster is a utility for reading log files and generating metrics in Graphite or Ganglia. It is ideal for visualizing trends of events that are occurring in your application/system/error logs. For example, you might use logster to graph the number of occurrences of HTTP response code that appears in your web server logs.

14.29 Pencil

Pencil is a monitoring frontend for graphite. It runs a webserver that dishes out pretty Graphite URLs in interesting and intuitive layouts.

14.30 Rocksteady

Rocksteady is a system that ties together Graphite, RabbitMQ, and Esper. Developed by AdMob (who was then bought by Google), this was released by Google as open source (http://google-opensource.blogspot.com/2010/09/get-ready-to-rocksteady.html).

14.31 Scales

Scales is a Python server state and statistics library that can output its data to Graphite.

14.32 Seyren

Seyren is an alerting dashboard for Graphite.

14.33 Shinken

Shinken is a system monitoring solution compatible with Nagios which emphasizes scalability, flexibility, and ease of setup. Shinken provides complete integration with Graphite for processing and display of performance data.

14.34 statsd

statsd is a simple daemon for easy stats aggregation, developed by the folks at Etsy. A list of forks and alternative implementations can be found at <http://joemiller.me/2011/09/21/list-of-statsd-server-implementations/>

14.35 Structured Metrics

structured_metrics is a lightweight python library that uses plugins to read in Graphite's list of metric names and convert it into a multi-dimensional tag space of clear, sanitized targets.

14.36 Tasseo

Tasseo is a lightweight, easily configurable, real-time dashboard for Graphite metrics.

14.37 Therry

Therry ia s simple web service that caches Graphite metrics and exposes an endpoint for dumping or searching against them by substring.

14.38 TimeseriesWidget

TimeseriesWidget adds timeseries graphs to your webpages/dashboards using a simple api, focuses on high interactivity and modern features (realtime zooming, datapoint inspection, annotated events, etc). Supports Graphite, flot, rickshaw and anthracite.

CHAPTER

FIFTEEN

WHO IS USING GRAPHITE?

Here are some organizations that use Graphite:

- Orbitz
- Sears Holdings
- Etsy (see http://codeascraft.etsy.com/2010/12/08/track-every-release/)
- Google (opensource Rocksteady project)
- Media Temple
- Canonical
- Brightcove (see http://opensource.brightcove.com/project/Diamond/)
- Vimeo

And many more

CHAPTER

SIXTEEN

RELEASE NOTES

16.1 0.9.10

5/31/12

Graphite 0.9.10 has been released and is now available. The packages for Whisper, Carbon, and Graphite-web are available via several sources:

- Pypi (and by extension, pip)
- http://pypi.python.org/pypi/graphite-web/
- http://pypi.python.org/pypi/carbon/
- http://pypi.python.org/pypi/whisper/
- Github
- https://github.com/graphite-project/graphite-web/downloads
- https://github.com/graphite-project/carbon/downloads
- https://github.com/graphite-project/whisper/downloads
- Launchpad
- https://launchpad.net/graphite/0.9/0.9.10

This release contains a fabulous amount of incremental improvement over 0.9.9. Some highlights include: * Fixes to several annoying Composer and Dashboard UI bugs * Import of Saved Graphs into Dashboards * Fixes to cache-full behavior for carbon-cache and carbon senders (relay and aggregator) * Many new useful render functions and graph options * Improvements to the rendering engine and fixes to many rendering bugs * Support for rendering graphs as annotated SVG * Better organized and more flexible Graphite-web config layout (local_settings.py)

Upgrading from 0.9.9 should be as simple as updating the packages. It is recommended but not necessary that local_settings.py be recreated based on the newly shipped local_settings.py.example as it includes many newly exposed settings and an improved organization and comments. Carbon's config files also have a few new settings as well to check out.

The Graphite project is also in the midst of some project changes. For those who have not yet noticed, the Graphite codebase has been moved to Github (http://github.com/graphite-project) and split into individual components (Graphite-web, Carbon, Whisper, and soon Ceres). The Launchpad project remains active in supporting the project with its Answers (http://answers.launchpad.net/graphite/) and Bugs (http://bugs.launchpad.net/graphite/) functionality.

Development going forward will focus on preparing what will become Graphite 0.10.0 which will include support for the Ceres database format as well as a major refactor of the Carbon daemon (nicknamed "Megacarbon"). The master

branches of the project should be considered to be in an 'alpha' state for the time being and subject to backwards-incompatible changes. Fixes to the current version will be maintained in the 0.9.x project branches but no 0.9.11 version is planned for the time being.

A big thanks goes out to all those who have helped the project in contributions of time and energy in the form of code contributions, testing, discussion, and helping each other out with support questions. Additional thanks are due to Aman Gupta (tmm1) for all of his great work on the rendering engine and other fixes, Sidnei Da Silva for his work migrating the project to Github and his fixes, and everyone who's taken the time to answer questions on the Answers site and on IRC.

As always, if you need any assistance please ask a question or join us on IRC in #graphite on Freenode.

The following is a summary of changes since the last release:

16.1.1 New features

Whisper

- Allocate Whisper files in chunks by default (jordansissel)
- Allow Whisper files to be allocated sparsely (jordansissel)
- Add whisper-merge command to copy data from one file to another (sidnei)
- Add whisper-dump utility (amosshapira)

Graphite Dashboard

- New button to retrieve Graph URL (octplane)
- Add button to send email of rendered graph as attachment (bkjones)
- Allow relative 'until' time to be set in dashboard (daniellawrence)
- Add ability to import Graphs into dashboards from URL or Saved Graphs

Rendering Engine

- New minorY option to configure minor gridlines (whd)
- New alpha() function to set individual color alpha values (tmm1)
- Allow areaAlpha to set alpha values for all styles of stacked graphs (tmm1)
- New minimumAbove() function: draw only series whose min is above n (tmm1)
- New areaBetween() function: draw the area between two graph lines (tmm1)
- New holtWintersConfidenceArea() function: display area between Holt-Winters confidence bands (tmm1)
- New SVG output format with embedded graph metadata (tmm1)
- New metric whitelist/blacklist functionality using pattern files
- New filterBelowPercentile() function: remove data below n percentile from a series (tmm1)
- New removeAbovePercentile() and removeAboveValue() functions to remove outliers (tmm1)
- New removeBelowPercentile() and removeBelowValue() functions to match above counterparts
- New aliasSub() function: perform a regex search/replace on metric names (tmm1)

- New rangeOfSeries() function: reduces multiple series into the value range of each point (saysjonathan)
- New movingMedian() function: moving median, similar to movingAverage (recursify)
- New multiplySeries() function: combine series by multiplying them
- New hideYAxis option (mdeeks)
- New percentileOfSeries() function: Combines series into the value at n percentile for each point
- New tranformNull() function: transforms None values to specified (cbrinley)
- New scaleToSeconds() function: scales values based on series step (redbaron)
- New aliasByMetric() function: trims all but the last element of metric name in legend (obfuscurity)
- New uniqueLegend option to filter duplicate metric names in legend (mdeeks)
- New vtitleRight option to label 2nd Y-axis

Carbon

- Allow flock() mode to be configured for Whisper
- Allow flushing of rrdcached before rrd data fetches (shufgy)
- Add ability to configure carbon metric prefix (jblaine)

16.1.2 Bug fixes

Whisper

- Record only the last value when duplicate timestamps are sent (knyar)
- Fix rrd2whisper.py script to work with newer python-rrdtool api

Carbon

- Fix full drain of queue after cache-full event when flow-control is enabled in both client and carbon-cache
- Fix unnecessary drop of a single metric point when cache is full
- Fix instrumentation of carbon-relay (darrellb)

Webapp

- Fix reading of Gzip'd whisper files and remote reading of RRDs
- Fix registration of Event model in admin site
- Fix events() to work with timezone aware dates
- Fix Event model to use tagging properly and fix compatibility with MySQL (hellvinz)
- Fix compatibility of built-in json module in events and graphlot
- Fix loading of saved graphs where a target has a '%' in the name

Rendering Engine

- Fix removal of whitespace above stacked graphs with yMax setting (tmm1)
- Use powers of 2 when calculating yStep and yUnitSystem=binary (tmm1)
- Force 100% usage of vertical space when yMax=max
- Compact memcached keys to keep size under 250 after Django processing (Kevin Clark)
- Fix alignFromTrue functionality in summarize() (tmm1)
- Fix cases of mismatched units in holt-winters bootstraps (lapsu,tmm1)
- Force integer in moving average window parameter (lapsu)
- Fix incorrect cache fetch when storage dir is symlinked (mk-fraggod)
- Fix infinite loop in Y-axis render when series range is very-very small
- Fix "Undo Function" button when braces expressions are present in the target
- Fix legend column calculation (darrellb)
- Fix broken aliasByNode() (darrellb)
- Fix rendering failures when infinite values are present in series
- Fix legend text overlap with dual Y-axis mode (nleskiw)
- Fix missing hunk of graph on right side with Dual Y-axis
- Fix cactiStyle() handling of None values
- Fix rendering breakage during DST time switch
- Allow multiple named stacks of metrics (aleh)
- Fix incorrect/misaligned graphs when series with unaligned steps are mixed in a graph
- Properly shift over series that have a later start time than the graph start

Composer

- Fix JS error on IE due to tailing list commas (reed-r-lance)
- Fix usage of + instead of %20 for spaces in URL encoding in composer view
- Fix display of a broken image rather than "No Data" when last target is removed
- Fix the loss of multiple targets when loading a saved graph with new params (vilkaspilkas)
- Fix unremovable duplicate metrics

Dashboard

- Fix automatic edit field selection on click (octplane)
- Fix usage of browser cache-busting uniq parameter to be filtered from memcache key (sidnei)
- Fix inability to remove Graphs with duplicate target lists

16.1.3 Other improvements

Carbon

- Match time units used in storage-schemas.conf with those in the webapp (ohlol)
- Only log Carbon queue fullness once (sidnei)
- Only log Carbon queue space free if it was once full (sidnei)
- Log a message with the affected filename when a Whisper update fails (bmhatfield)
- Move carbon instance logs to their own own directory to prevent clobbering
- Prevent carbon-aggregator from clobbering aggregated values when aggregating to same-name
- Add SSL option to amqp publisher (sidnei)
- Remove duplicate dot metric path filtering for performance (drawks)
- Refactor of schema validation to give more informative errors
- Add reloading of rewrite-rules and aggregation-schemas for consistency

Webapp

- Refactor settings.py to allow more complete configuration in local_settings.py
- Make Graphite compatible with Django 1.4
- Add jsonp support for /browser endpoint
- Make it harder to break metric browsing with a bad DATA_DIRS entry

Rendering Engine

- Make asPercent() much more flexible and useful
- stddev() function made more robust
- Allow metrics to begin with a braces-wildcard
- Prevent drawAsInfinite() lines from affecting Y axis height (bmhatfield)
- Pass through time with secondly rather than minutely resolution to rrdfetch (tmm1)
- Tree branches should display above all leaves (mdeeks)
- Add alignToInterval to hitcount() function similar to summarize() (jwoschitz)
- Fix PieGraph missing function
- Allow timeShift() to shift forward as well as backward

Composer

- Don't reorder targets when applying functions
- Refactor of Graph Options menu

Dashboard

- Explicitly size img tags to keep scroll position intact during reloads
- Default the navBar as collapsed when loading an existing dashboard view
- Show wildcards in top nav browsing view
- Allow dashboards to have any character in title (octplane)
- Make "Remove All Graphs" and "Change Size" dialogs modal (dannyla)
- Make the new "north" navbar the default

16.2 0.9.11

8/20/2013

Graphite 0.9.11 is now available for usage. Source bundles are available from GitHub:

- https://github.com/graphite-project/graphite-web/archive/0.9.11.tar.gz
- https://github.com/graphite-project/carbon/archive/0.9.11.tar.gz
- https://github.com/graphite-project/whisper/archive/0.9.11.tar.gz

Graphite can also be installed from Pypi via pip. Pypi bundles are here:

- http://pypi.python.org/pypi/graphite-web/
- http://pypi.python.org/pypi/carbon/
- http://pypi.python.org/pypi/whisper/

16.2.1 Upgrading

It's recommended to install all three 0.9.11 packages together for the most success, however in this case *graphite-web* can be installed separately from carbon if necessary. *Carbon* and *Whisper* must be updated together due to the coupling of certain changes.

Graphite 0.9.11 now requires a Django version of at least 1.3. Ensure this dependency is satisfied before updating *graphite-web*

As always, comparing the example config files with existing ones is recommended to ensure awareness of any new features.

16.2.2 Security Notes

This release contains several security fixes for cross-site scripting (XSS) as well as a fix for a remote-execution exploit in graphite-web (CVE-2013-5903). Patches for the past three prior releases are available in these gists:

- 0.9.10
- 0.9.9
- 0.9.8

In a pinch, the following url mapping can be removed by hand if the remote-rendering feature is not being used:

```
diff --git a/webapp/graphite/render/urls.py b/webapp/graphite/render/urls.py
index a94a5d1..f934b43 100644
--- a/webapp/graphite/render/urls.py
+++ b/webapp/graphite/render/urls.py
@@ -15,7 +15,6 @@ limitations under the License."""
 from django.conf.urls.defaults import *

 urlpatterns = patterns('graphite.render.views',
-  ('local/?$','renderLocalView'),
   ('~(?P<username>[^/]+)/(?P<graphName>[^/]+)/?','renderMyGraphView'),
   ('', 'renderView'),
 )
```

Finally, The setting of Django's SECRET_KEY setting is now encouraged and exposed in local_settings.py as well.

16.2.3 New Features

Graphite-web

- Properly return an HTTP 400 on missing query parameter in metrics/search endpoint (dieterbe)
- cumulative() is now superceded by consolidateBy() which supports min/max/avg/sum (nleskiw)
- Make graphlot target host configurable for easier embedding (dieterbe)
- Allow graphlot graphs to be embedded for use in dashboard apps (dieterbe)
- When wildcarding, prefer matching metric files to directories with the same name (tmm1)
- New header design and css cleanups (obfuscurity)
- New composer button to open the target in graphlot (magec)
- timeshift() can now shift beyond current time, allowing better current-over-week charts (mgb)
- Unit scaling added to cactiStyle (drawks)
- Support RRD files in index.json view (obfuscurity)
- Support for alternate target[] url syntax (luxflux)
- New countSeries() function which returns the cardinality of a wildcard (obfuscurity)
- Bootstrap data for movingAverage and movingMedian (seveas)
- movingAverage and movingMedian now optionally take time periods to specify window size (danielbeardsley)
- jsonp support in events/get_data (gingerlime)
- Ace editor for manually editing dashboard json (jordanlewis)
- New stddevSeries(), timeStack() functions (windbender)
- Remove ugly graph image background in dashboard (frejsoya)
- y-axis divisors for determining y-axis scale are now configurable (wfarr)
- Allow any characters in axis labels
- Target grammar now supports scientific notation for numbers
- New identity() function (dieterbe)
- Update default color scheme (obfuscurity)

- Dont blow up on permissions errors while walking directories (log instead)
- Encourage users to set SECRET_KEY uniquely with a warning

Carbon

- Improvements to setup.py rpm generation and basic init scripts (bmhatfield)
- Allow alternate update rate at shutdown (Daniel314)
- Add support for new fallocate() allocation method in Whisper (slackhappy)
- Improvements to noisy logging (nleskiw, drawks)
- Protect against writes outside the storage tree
- Performance fixes to rate limiting, removal of unnecessary locks (drawks)
- Alternate write strategies for carbon-cache (max size, random) (drawks)
- carbon-aggregator aware consistent-hashing for carbon-relay (slackhappy)
- Allow custom umask to be passsed to twisted at startup (egnyte)
- New options WRITE_BACK_FREQUENCY to control frequency of partially-aggregated output (jdanbrown)
- Improve consistent-hashing performance when replication factor is 1 (slackhappy)
- Various code cleanups (sejeff)
- Allow a timestamp of -1 to be sent to aggregator to set to current time (gwillem)
- Allow log rotation to be handled by an external process (justinvenus)
- min/max aggregation methods are now supported (ishiro)

Whisper

- Better commandline sanity checking and messaging (sejeff)
- Handle SIGPIPE correctly in commandline utils (sejeff)
- Option to intelligently aggregate values on whisper-resize (jens-rantil)
- Use more efficient max() instead of sorted()[-1] (ryepup)
- Add fallocate() support (slackhappy)
- Improve handling of exceptional fetch cases (dieterbe)
- Improve rrd2whisper's handling of rrd files
- Improve error messaging on retention errors at create time (lambdafu)

16.2.4 Bug fixes

Graphite-web

- broken nPercentile() and related functions
- Python 2.4 compatibility in browser endpoint (dcarley)
- Missing URL parameters in composer load

- Fix to multiplySeries to return the expected type (nleskiw)
- Don't blow up when empty series passed to cactiStyle (mattus)
- Trailing commas in js breaking ie (nleskiw, davecoutts)
- Remove extra and unnecessary rendering while loading saved graphs (hostedgraphite)
- Broken entry of timezone in composer menu (hcchu)
- constantLine() not drawing across the entire graph (mattsn0w)
- SVG rendering broken when using secondYAxis (obfuscurity)
- Expect url-encoded octothorpes in colorList (magec)
- Display relative times properly in dashboard (daveconcannon)
- cactiStyle() blows up with empty series (eranrund)
- Remove problemmatic and unnecessary url encoding
- Several pathExpressions missing which caused trouble in certain function combinations (dieterbe,colby,kovyrin)
- Use non-linux-specific datetime formatter %I instead of %l (richg)
- Use os.sep properly for path separation (justinc)
- Negative numbers not allowed in yAxis input box
- scale() misreports itself in legend when using small decimals
- colorList incorrectly cast to an int in some cases (rckclmbr)
- removeBelow* menu items adding the wrong functions to target list (harveyzh)
- nPercentile renders it's name incorrectly (TimZehta)
- CSV rendering does not respect tz parameter
- Missing max interval in xAxisConfigs causes long-term graphs with few points to render with a 12hr axis config
- Stacked graphs not filling completely in staircase mode
- Stacked graphs and many drawAsInfinite() lines do not draw cleanly
- Graphlot does not handle event timestamps properly (matthew keller)
- sin() time() and randomWalk() incorrectly using float times (jbrucenet)
- legend height is incorrect when secondYAxis used (obfuscurity)
- Expanded wildcards in legends are misordered (dieterbe)
- Regression in formatPathExpression (jeblair)
- index.json returns leading periods when WHISPER_DIR does not endin a trailing slash (bitprophet)
- Regression in areaMode=all causes only the last series to be filled (piotr1212)
- Default to settings.TIMEZONE if timezone unknown (gingerlime)
- Negative filled graphs render from bottom rather than 0 (piotr1212)
- Composer and Dashboard XSS fixes (jwheare, sejeff)
- Fix persistence of tz aware datetime in non-postgres databases
- Fix insecure deserialization of pickled objects (CVE-2013-5093)
- Lots of documentation improvement (jeblair,bclermont,lensen,cbliard,hvnsweeting)

Carbon

- Empty lines match everything in whitelist (gographs)
- Storage-schemas dont auto reload when they should
- Carbon-relay per-destination metrics are broken
- Regression in MAX_CREATES_PER_MINUTE where values >60 were set to 0 (jeblair)
- Memory leak in carbon-aggregator in certain cases (lbosson)
- Python2.4 compatibility in AMQP send/receive (justinvenus)
- Cache/queue sizes are misreported (bitprophet)
- NaN values shouldn't be passed through from amqp (llaurent)

Whisper

- Python2.4 compatibility for whisper-dump.py (snore)
- Correct filtering of duplicate values to ensure last-write-wins

16.3 0.9.2

I've received a bunch of great bug reports and feedback on the 0.9 release and have resolved lots of minor issues this week as a result. So please try out the new 0.9.2 release on the downloads page and keep the bug reports coming!

- ChrisMD

16.4 0.9.3

7/16/08

This release is an incremental improvement over 0.9.2, including lots of bug fixes, major enhancements to the installer, and several new handy scripts. Thanks to everyone who submitted bug reports and questions. The next few Graphite releases will continue to focus on quality rather than new features. In particular, 0.9.4 will include a re-write of the carbon backend, which will be much simpler and easier to administer and troubleshoot. I am also working on porting lots of internal documentation to this wiki. My goal is to have a 1.0 release by the end of the year, which must be well-documented, easy to deploy, easy to troubleshoot, and of course as bug-free as possible. If there is time a new feature or two might make it in, but this is not the primary focus.

- ChrisMD

16.5 0.9.4

1/30/09

It's been a good 6 months since the last release. Not much has changed aside from a few minor enhancements and some good bug fixes, unfortunately I've not had nearly as much time as I'd like to dedicate to working on Graphite. Regardless, it is getting more mature slowly but surely. In the next few months I may be in a better position to get more real work done on it, but we shall see. For now I'd just like to thank everyone who has given me great questions and bug reports, your feedback is what keeps this project moving. Thanks.

- ChrisMD

16.6 0.9.5

1/4/10

It's hard to believe it's been an entire year since the last release of Graphite. This just goes to show how good I am at procrastination. After taking a look at the old 0.9.4 release I can safely say that the new 0.9.5 release is very significant. Here are the biggest changes:

Graphite now supports [[[federated storage]]] (for better scalability) # Carbon was completely rewritten using Twisted (much cleaner, more configurable, and more fault tolerant) # The installation process now uses distutils (finally!) # Graphite, Carbon, and Whisper are now three separate packages (for more flexible deployment) # Graphite's browser UI fully migrated to pure ExtJS 3.0 (cleaner code, less bugs) # Many many bug fixes as always

I'd like to thank everyone in the community who has been using graphite and contributing to the project. I don't usually do new year's resolutions, but I've got a good idea for one this year. I really want to get away from infrequent huge releases like this one and get back to very frequent small releases in the true spirit of open source. So my resolution is to release *something* once a month. It will probably usually just be bug fixes, or perhaps some much needed documentation. So look forward to something new in February!

- *ChrisMD*

16.7 0.9.6

2/26/10

This has probably been the most active month of Graphite development since the project was open sourced. Lots of community members have contributed code and ideas to help move Graphite forward. I'm really excited about this, the project is gaining momentum and I hope we can keep that up by continuing with the new monthly release cycle. To give credit where it is due, here is a list of this month's most active users and what they've been working on (in no particular order):

- Lucio Torre - AMQP support
- jdugan - beautification of the Y-axis labels via the yUnitSystem option
- Nick Leskiw - the YAxis=right rendering option
- Kraig Amador - tons of rendering options/functions such as yLimit, timeShift(), log(), sumSeriesWithWildcard(), new filtering functions, and much more! (Kraig you're the man!)
- Arthur Gautier - debian packaging
- Elliot Murphy - packaging, inclusion in Ubuntu
- fp - RHEL / CentOS RPM packaging
- and many more...

Thanks to everyone who has gotten involved with Graphite, your support helps motivate others (especially me).

Many of these new features are really great but unfortunately undocumented, but the good news is that my focus for March is going to be 100% on *documentation*. There may not be an actual code release in March but I hope to get a substantial amount of documentation written right here on this wiki. Stay tuned.

- ChrisMD

16.8 0.9.7

1/8/11

A little late but better than never, Graphite 0.9.7 is now out and available for download. It available through PyPI (http://pypi.python.org/pypi) and the Launchpad project page (https://launchpad.net/graphite). Here is a quick-rundown of the new features and some nice bug fixes:

16.8.1 Features

- Composer UI menus have been updated to reflect all currently available functions and options
- New threshold() function allows you to draw a horizontal line with a custom color and legend name (though color and legend name are not available through composer UI yet)
- New summarize() function allows you to draw data at a lower precision than it is stored at (ie. draw hourly datapoints for minutely data)
- New group() function allows you to specify a collection of metrics for passing to other functions that require a single arg, without using wildcards.
- Retention configurations support a new more convenient syntax (see [https://bugs.launchpad.net/graphite/+bug/697896 Bug #697896])
- Carbon's logging of every whisper update can be disabled now (set LOG_UPDATES = False in carbon.conf)
- Carbon-relay can now specify ports for remote carbon-caches
- Timezones can now be specified at render-time using Olson timezone names (see [http://pytz.sourceforge.net/ pytz])
- Saved MyGraphs now support a hierarchical structure when dots are used in the saved graph names
- By popular request, carbon now ignores improperly formatted datapoint lines rather than disconnecting the client ([https://bugs.launchpad.net/graphite/+bug/589476 Bug #589476])
- X-axis labeling has been revamped to avoid overlapping and confusing labels
- RPM and source RPM packages are available for download. Note that they currently **do not check dependencies** and **do not perform post-install tasks**. This means they are suitable for upgrades but the usual install doc will need to be followed for new installations. Please contribute feedback regarding these packages so we can make them work out of the box on Fedora and CentOS.

16.8.2 Bugs Fixed (woefully incomplete)

- [https://bugs.launchpad.net/graphite/+bug/528228 Bug #528228] - fixed 'tz' parameter for specifying custom timezone at render-time
- [https://bugs.launchpad.net/graphite/+bug/676395 Bug #676395] - fixed timeShift() function
- [https://bugs.launchpad.net/graphite/+bug/690586 Bug #690586] - fixed log() function to work with negative data
- [https://bugs.launchpad.net/graphite/+bug/684563 Bug #684563] - fixed carbon-cache.py –config parameter
- [https://bugs.launchpad.net/graphite/+bug/660861 Bug #660861] - fixed nonNegativeDerivative() math for wrapping counters
- [https://bugs.launchpad.net/graphite/+bug/591948 Bug #591948] - fixed X-axis labeling that was inaccurate in some situations

- [https://bugs.launchpad.net/graphite/+bug/595652 Bug #595652] - fixed bug preventing clustering from working with default settings.py
- [https://bugs.launchpad.net/graphite/+bug/542090 Bug #542090] - fixed Y-axis labeling issue for large values with small variance
- Dozens more...

The best part is, the work in this release has continued to be largely a community effort. Almost all bugs that got fixed were reported from users and the vast majority have been fixed because of contributed patches and highly detailed bug reports. In other words, this ain't a one-man show! Thanks to everyone who has contributed code, bug reports, documentation, questions, and answers.

In the interest of not being incredibly wrong again, I will refrain from putting a date on when the next Graphite release will be out. But it will not be another year, that's for sure... Several projects I have to do for work in the coming months are going to involve major enhancements to Graphite, unlike this past year during which I've really only worked on it in my spare time. Thanks again to everyone and happy new year!

- ChrisMD

16.9 0.9.8

4/3/11

Graphite 0.9.8 is now out and available for download. It available through PyPI (http://pypi.python.org/pypi) and the Launchpad project page (https://launchpad.net/graphite).

This release is a major step forward for Graphite, with a long list of substantive enhancements only 3 months after the last release. One of the highlights is the move of our documentation to readthedocs.org, the docs are now built using Sphinx and they live in trunk under the 'docs' folder. Just commit any changes and readthedocs.org will automatically update by pulling changes from launchpad nightly.

A special thanks goes out to AppNexus (http://appnexus.com/), who sponsored the development of two awesome new features. First is the new carbon-aggregator daemon. This new daemon lets you configure the calculation of aggregate metrics at storage time instead of using a heavy-weight sumSeries or averageSeries at rendering time. This daemon can also rewrite metric names. You manage it like the other two carbon daemons, via carbon.conf. Documentation on configuring carbon-aggregator will be coming soon.

AppNexus also sponsored the development of the new Dashboard UI. This new interface allows you to put together dashboards containing many graphs quickly and easily. You can save a dashboard and view it later. Note that this is a basic implementation for now

Beyond that, there are many other new features so please read through the changelog carefully.

16.9.1 Changes

- New carbon-aggregator daemon can compute your aggregate metrics
- New Dashboard UI
- Upgraded to ExtJS 3.3
- All Documentation is moving to Sphinx in our bzr branch, HTML builds of it are hosted by readthedocs.org (http://graphite.readthedocs.org/)
- The recommended Apache setup is now officially mod_wsgi and not mod_python.
- New metric pattern syntax, eg. {{example.{foo,bar}.metric}}, matches both {{example.foo.metric}} and {{example.bar.metric}}

- Y-axis now draws much more useful labels for values much less 1
- The YAxis=left|right parameter has been renamed to yAxisSide=left|right
- Rewrote webapp/render/grammar.py to be much more readable
- Added new json api call /metrics/expand/?query=foo.* -> ["foo.bar", "foo.baz", ...]
- Added debugging manhole in carbon-cache.py (ssh-accessible python intepreter interface into carbon at runtime)
- Added new hitcount function (thanks to Shane Hathaway)
- The "User Graphs" tree now works properly for usernames that contain dots
- Fixed data roll-up bug in whisper
- Added AUTOFLUSH option in whisper/carbon for synchronous I/O
- and as always, many more smaller bug fixes

- ChrisMD

16.10 0.9.9

10/6/11

Graphite 0.9.9 is now out and available for download. It available through PyPI (http://pypi.python.org/pypi) and the Launchpad project page (https://launchpad.net/graphite).

This is a very substantial release. To give you an idea, the 0.9.8 release was cut from trunk around revision 380 while 0.9.9 was cut from revision 589, so that's almost as many commits as Graphite has ever had just since 0.9.8. The full changelog is too big for me to assemble nicely unfortunately, but I will try to cover all the important bits and if you're really curious you can see all the changes at http://bazaar.launchpad.net/~graphite-dev/graphite/main/changes

There are some really important things you need to know if you're upgrading from an earlier release (even trunk). Read all the change summaries below please!

16.10.1 API Changes

There have been API changes in whisper, carbon, and the webapp. If you are upgrading to 0.9.9 YOU MUST UPGRADE ALL 3 PACKAGES, if you mix 0.9.8 whisper with 0.9.9 carbon for example, it won't work. Upgrade all 3, and don't forget to use the –force.

The webapp has a new dependency on django.tagging (you should be able to simply 'pip install django-tagging')

16.10.2 New Default Behavior

We've addressed a security vulnerability with receiving pickled datapoints, see Bug #817247. This affects you in that the new default behavior is to use a more secure unpickler, which is slightly slower than the standard insecure unpickler. To revert to the less secure but faster approach previously used, you have to set USE_INSECURE_UNPICKLER=True in your carbon.conf.

16.10.3 Revamped Dashboard UI

- You can now use all the composer functionality by clicking on a dashboard graph
- You can drag and drop to move graphs around (and hover-drop to combine them!)

- There is an awesome new auto-completer interface available by going to the Dashboard menu, Configure UI, Completer. This may become the default in the future because its so awesome. (pro tip: try using the dashboard completer with * instead of * for some really powerful 'group by' functionality)

16.10.4 Other Stuff

- Tons of readthedocs.org improvements, also the example config files now have some great comment documentation
- Whisper now supports rollup aggregation methods other than averaging. The default is

still to average but there a new aggregation-schemas.conf (see Bug #853955) * To learn about the new metric metadata API that can be used to configure custom rollup aggregation methods read my answer to https://answers.launchpad.net/graphite/+question/173304 (you can skip the question part if you just care about the new API)

As for the current development focus, I can now finally work on the long-awaited merge of the 1.1 branch into trunk. The Ceres database will be in the next release, I'm going to try and merge it in (including the new refactored storage API) in the next week or so. I'll announce on graphite-dev when its available for testing. My aim is to get it fully documented for 1.0, which I'm targetting for end of this year. There might be an 0.9.10 first, depending on how many bugs are found in 0.9.9.

As always, thanks to everyone who has contributed, especially the following rockstar crew that made some major contributions in the past few months:

- Aman Gupta (tmm1)
- Nick Leskiw (nleskiw)
- Sidnei da Silva

- ChrisMD

CHAPTER

SEVENTEEN

INDICES AND TABLES

- *genindex*
- *modindex*
- *search*

PYTHON MODULE INDEX