

Collectd 101

- [What is collectd?](#)
- [Why use collectd?](#)
- [collectd architecture](#)
- [collectd data flow](#)
- [collectd statistics:](#)
- [collectd Values](#)
- [collectd Data Sets](#)
- [Examples of types in collectd](#)
- [collectd notifications](#)
- [Notifications and thresholds](#)
- [Collectd cool features](#)
 - [Exec Plugin](#)
 - [Collectd Plain Text Protocol](#)
 - [Filter Configuration](#)
 - [Network plugin](#)
 - [Writing a simple plugin](#)
- [References](#)

What is collectd?

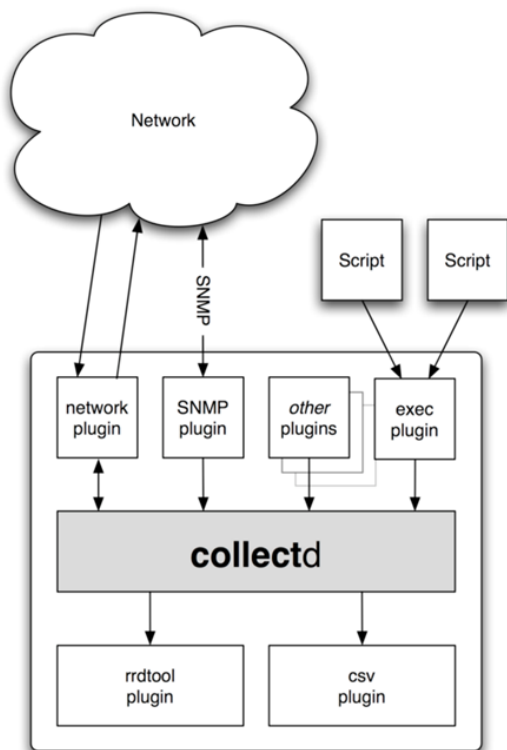
- collectd is a system statistics collection daemon; it uses plugins to collect system statistics and can publish those statistics in a number of ways.
- Free open source project, GNU General Public License, version 2 (GPLv2), the remaining files are licensed under other open source licenses
- collectd also provides some simple thresholding and event notification capabilities

Why use collectd?

- The collectd API, the selection of available plugins (90+) coupled with the modular nature of collectd architecture, provides a generic path to expose platform statistics from all existing plugins and newly developed plugins to OpenStack or other fault management applications.
- collectd also has bindings for several programming languages, which allows a developer to implement a plug-in in their language of choice, without having to modify the receiver of the collected metrics, for example Ceilometer itself.


collectd architecture

- plugins in collectd fall under one of the following categories:
- Input plugins: read system statistics at a regular interval, and dispatch the values to the collectd daemon.
 - Output plugins: receive dispatched values from the daemon and output/write/dispatch those values in various output formats (e.g. JSON, SNMP, AMQP, MySQL, HTTP, etc).
 - Binding plugins: which provide language bindings for collectd to allow plugins to be written in languages other than C.
 - Logging plugins: which write information to log files or dispatch messages to syslog.
 - Notification plugins: enable limited monitoring support in collectd.
 - Other plugins which can both read and dispatch values.



Source: <https://collectd.org/images/architecture-schematic.png>

collectd data flow



collectd statistics:

Statistics in collectd consist of a value list. A value list includes:

- Values
- Value length: the number of values in the data set.
- Time: timestamp at which the value was collected.
- Interval: interval at which to expect a new value.
- Host: used to identify the host.
- Plugin: used to identify the plugin.
- Plugin instance (optional): used to group a set of values together. For e.g. values belonging to a DPDK interface.
- Type: unit used to measure a value. In other words used to refer to a data set.
- Type instance (optional): used to distinguish between values that have an identical type.
- meta data: an opaque data structure that enables the passing of additional information about a value list. "Meta data in the global cache can be used to store arbitrary information about an identifier"

Host, plugin, plugin instance, type and type instance uniquely identify a collectd value

collectd Values

Values, can be one of:

- Derive: used for values where a change in the value since it's last been read is of interest. Can be used to calculate and store a rate.
- Counter: similar to derive values, but take the possibility of a counter wrap around into consideration.
- Gauge: used for values that are stored as is.
- Absolute: used for counters that are reset after reading collectd consist of a value list.

collectd Data Sets

Values lists are often accompanied by data sets that describe the values in more detail. Data sets consist of:

- A type: a name which uniquely identifies a data set.
- One or more data sources (entries in a data set) which include:
 - The name of the data source. If there is only a single data source this is set to "value".
 - The type of the data source, one of: counter, gauge, absolute or derive.
 - A min and a max value.

Examples of types in collectd

Examples of types in types.db:

```
bitrate value:GAUGE:0:4294967295
```

```
counter value:COUNTER:U:U
```

```
if_octets rx:COUNTER:0:4294967295,tx:COUNTER:0:4294967295
```

In the example above if_octets has two data sources: tx and rx.

collectd notifications

Notifications in collectd are generic messages containing:

- An associated severity, which can be one of OKAY, WARNING, and FAILURE.
- A time.
- A Message
- A host.
- A plugin.
- A plugin instance (optional).
- A type.
- A types instance (optional).
- Meta-data.

Notifications and thresholds

<https://collectd.org/documentation/manpages/collectd-threshold.5.shtml>

Collectd cool features

Exec Plugin

- Executes scripts / applications and reads values back that are printed to STDOUT by that program.
- Extends the daemon in an easy and flexible way.
- Can also be used to call a bash script that does something with the notification from the threshold plugin.

Example of a script used by the exec plugin to write notifications to a file /tmp/notifications :

```
<Plugin exec>

#       Exec "user:group" "/path/to/exec"

       NotificationExec "stack" "write_notification.sh"

</Plugin>
```

write_notification.sh just writes the notification passed from exec through STDIN to a file (/tmp/notifications).

```
#!/bin/bash
```



```
rm /tmp/notifications

while read x y
do
    echo $x$y >> /tmp/notifications
done
```

an example generated file is shown below:

```
cat /tmp/notifications contents
Severity:FAILURE

Time:1472552207.385

Host:pod3-node1

Plugin:dpdkevents

PluginInstance:dpdk0

Type:gauge

TypeInstance:link_status

DataSource:value

CurrentValue:1.000000e+00

WarningMin:nan

WarningMax:nan

FailureMin:2.000000e+00

FailureMax:nan
```

Hostpod3-node1, plugin dpdkevents (instance dpdk0) type gauge (instance link_status): Data source "value" is currently 1.000000. That is below the failure threshold of 2.000000.

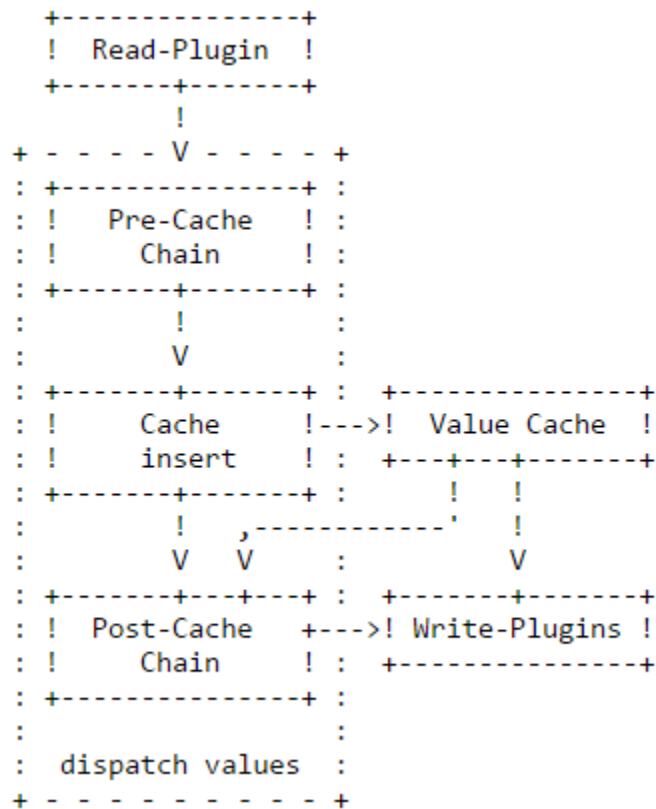
Collectd Plain Text Protocol

Submit statistics and notifications to the daemon as well as query the current value of collected statistics.

Plugins currently using this protocol are [Exec](#) (partially) and [UnixSock](#)

Filter Configuration

- Starting with collectd 4.6 there is a powerful filtering infrastructure implemented in the daemon.
- The concept has mostly been copied from ip_tables.
- Terminology:
 - Match: criteria to select specific values.
 - Target: action that is to be performed with data.
 - Rule: The combination of any number of matches and at least one target
 - Chain: a list of rules and possibly default targets. Rules tried in order and if one matches, the associated target will be called.



- When "read" plugins call the dispatch functions to dispatch values, the pre-cache chain is run.
- The values are then added to the internal cache.
- The post-cache chain is run after the values have been added to the cache.
- Allows you to remap value names.
- The cache is also used to convert counter values to rates.
- More info @ https://collectd.org/documentation/manpages/collectd.conf.5.shtml#filter_configuration

Network plugin

- Uses a [binary protocol](#)
- UDP transport
- Sends data to a remote instance of collectd, receives data from a remote instance, or both at the same time.
- Data which has been received from the network can be Forwarded again.
- It's possible to sign or encrypt the network traffic.

Writing a simple plugin

Please see <https://wiki.opnfv.org/display/fastpath/Collectd+how+to+implement+a+simple+plugin>

Coding Style: <https://wiki.ith.intel.com/display/HA/CollectD+plugin+code+style>

References

<https://collectd.org/>

<https://collectd.org/faq.shtml>

<https://collectd.org/documentation.shtml>

https://collectd.org/wiki/index.php/Notifications_and_thresholds

https://collectd.org/wiki/index.php/First_steps

https://collectd.org/wiki/index.php/Plugin_architecture

https://www.netways.de/fileadmin/images/Events_Trainings/Events/OSMC/2015/Slides_2015/collectd_Thresholds_Plugin_and_Icinga_-_Florian_Forster.pdf

<https://collectd.org/documentation/manpages/collectd.conf.5.shtml>

https://collectd.org/wiki/index.php/Plain_text_protocol

https://collectd.org/wiki/index.php/Plugin_architecture