Collectd Abstraction Layer usecases

This is a WIP

The suggested Abstraction Layer requires an extension to the collectd internal protocol to allow for plugins to be dynamically added/removed.

Overview

The idea behind this wiki is to look at the use cases for an Agent REST API, in this case collectd, to help define and implement a REST API that will allow us to dynamically do the following with collectd:

- * Enable/disable/or restart resource monitoring
- * Get values/notifications
- * Get capabilities
- * Get the list of metrics being collected
- * Flush the list of metrics
- * Set thresholds for resources
- * Blacklist resources
- * Support some sort of buffering mechanism, and should be able to configure
- * Get the timing information for the agent and do aTiming sync if required.

Phase 1 in the diagram below represents the current collectd implementation. Phase 2 is what we want to enable.



Usecases

Enable Plugin





Reconfigure Plugin



Get Values

TODO

See Cache contents

Get the list of metrics being collected

Flush Cache

Get Notifications



Heartbeat





Get Capabilities

TODO

Blacklist resources/meters

TODO

Buffer

TODO

Time Sync

get the timing information for the agent and do aTiming sync if required.

TODO

Authentication

For OpenStack Integration - keystone will be used for Authentication

Upgrade/Reconfiguration today without changes to collectd

Upgrade or reconfiguration changes to the collectd daemon on the fly can be done today through the use of containers - whereby you would always spin up collectd in a container context, and when you want to reconfigure the agent, you spin up a new instance of collectd with the new configuration and you terminate the older instance of collectd (in the other container). some consideration needs to be given for configuration of the collectd socket between instances and other configurable parameters that might be shared between multiple instances