

Collectd DCPMM Plugin HLD

Requirement

	Name
1.0	Uses APIs available with libpmwapi in Intel® PMWatch v3.0 or above
2.0	Uses libipmctl v01.00.00.3262 or above
3.0	Uses DCPMM DIMM firmware v01.00.00.5127 or above
4.0	Should have configurable metrics group selection
5.0	Should have configurable collection interval

Overview

Collectd DCPMM plugin monitors Intel® Optane™ DC Persistent Memory and provides memory performance and health information metrics.

Design

dcpmm plugin

The dcpmm plugin collects memory performance, health information metrics and timestamp information listed below.

Name	Type	Type Instance	Description	Comment
Memory performance metrics				
total_bytes_read	media	total_bytes_read	Number of bytes transacted by the read operations	
total_bytes_written	media	total_bytes_written	Number of bytes transacted by the write operations	
read_64B_ops_rcvd	media	read_64B_ops_rcvd	Number of read operations performed to the physical media in 64 bytes granularity	
write_64B_ops_rcvd	media	write_64B_ops_rcvd	Number of write operations performed to the physical media in 64 bytes granularity	
media_read_ops	media	media_read_ops	Number of read operations performed to the physical media	
media_write_ops	media	media_write_ops	Number of write operations performed to the physical media	
host_reads	controller	host_reads	Number of read operations received from the CPU (memory controller)	
host_writes	controller	host_writes	Number of write operations received from the CPU (memory controller)	
read_hit_ratio	buffer	read_hit_ratio	Measures the efficiency of the buffer in the read path. Range of 0.0 - 1.0	
write_hit_ratio	buffer	write_hit_ratio	Measures the efficiency of the buffer in the write path. Range of 0.0 - 1.0	
Health information metrics				
health_status	health	health_status	Overall health summary (0: normal 1: non-critical 2: critical 3: fatal)	
lifespan_remaining	health	lifespan_remaining	The module's remaining life as a percentage value of factory expected life span	
lifespan_used	health	lifespan_used	The module's used life as a percentage value of factory expected life span	
power_on_time	health	power_on_time	The lifetime the DIMM has been powered on in seconds	
uptime	health	uptime	The current uptime of the DIMM for the current power cycle in seconds	
last_shutdown_time	health	last_shutdown_time	The time the system was last shutdown. The time is represented in epoch (seconds)	
media_temperature	health	media_temperature	The media's current temperature in degrees Celsius	
controller_temperature	health	controller_temperature	The controller's current temperature in degrees Celsius	
max_media_temperature	health	max_media_temperature	The media's the highest temperature reported in degrees Celsius	

max_controller_temperature	health	max_controller_temperature	The controller's highest temperature reported in degrees Celsius	
Timestamp				
tsc_cycles	timestamp	tsc_cycles	The number of tsc cycles during each interval	
epoch	timestamp	epoch	The timestamp in seconds at which the metrics are collected from DCPMM DIMMs	

Plugin configuration

The following configuration options should be supported by dcpmm collectd plugin:

Name	Description	Comment
Interval	The collection interval in seconds at which the metric counts are collected	Defaults to global Interval value. This will override the global Interval value for dcpmm plugin. None of the other plugins will be affected.
CollectHealth	Health information metrics will be collected if set to true	Default value is false.
CollectPerfMetrics	Memory performance metrics will be collected if set to true	Default value is true.
EnableDispatchAll	This parameter helps to seamlessly enable simultaneous health and memory performance metrics collection in future.	This is unused at the moment and must always be false.

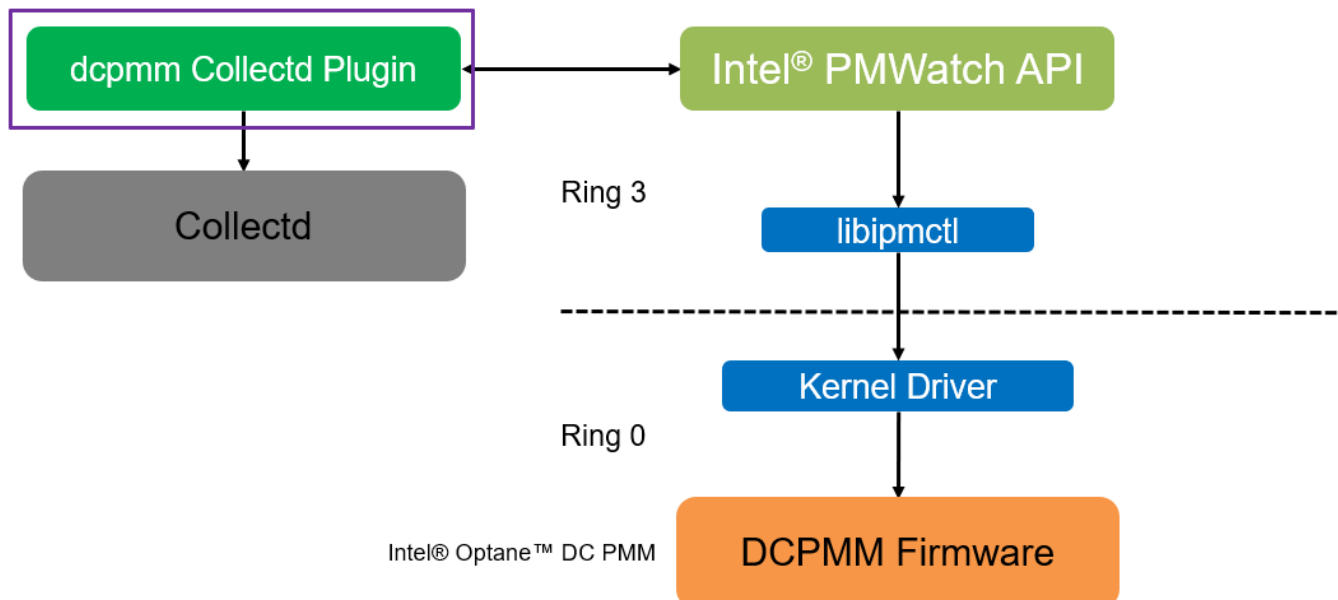
Here is an example of the plugin configuration section of collectd.conf file:

```
<Plugin dcpmm>
  Interval 10.0
  CollectHealth false
  CollectPerfMetrics true
  EnableDispatchAll false
</Plugin>
```

Implementation details

To enable monitoring of Intel® Optane™ DC Persistent Memory, APIs provided by Intel® PMWatch is used. This is an open source tool and is available at <https://github.com/intel/intel-pmwatch>.

The following diagram shows the high level architecture of dcpmm plugin.



plugin API	libpmwapi API	Description
dcpmm_config		Parse configuration provided in collectd.conf and register read callback if the correct configuration is provided
dcpmm_init	PMWAPIGetDIMMCount	Obtain the number of DCPMM DIMMs
	PMWAPIStart	Passes the configuration and starts the collection
dcpmm_read	PMWAPIRead	Reads the metric values
dcpmm_shutdown	PMWAPIStop	Stops the collection

Considerations

Configuration Considerations

The recommended collection interval is 1 second or above.

Deployment Considerations

If Intel® Optane™ DC Persistent Memory (DCPMM) is not available in the system, the plugin will be unloaded during initialization callback.

API/GUI/CLI Considerations

Equivalence Considerations

Security Considerations

Alarms, events, statistics considerations

Redundancy Considerations

Performance Considerations

The recommended collection interval is 1 second or above.

Testing Consideration

The timing interval requirement needs to be taken into consideration when conducting tests.

The Tests should be carried out on a system underload as well as a relatively idle system.

Other Considerations

Impact

The following table outlines possible impact(s) the deployment of this deliverable may have on the current system.

Ref	System Impact Description	Recommendation / Comments
1		

Key Assumptions

The following assumptions apply to the scope specified in this document.

Ref	Assumption	Status
1		

Key Exclusions

The following exclusions apply to the scope discussed in this document.

Ref	Exclusion	Status
1		