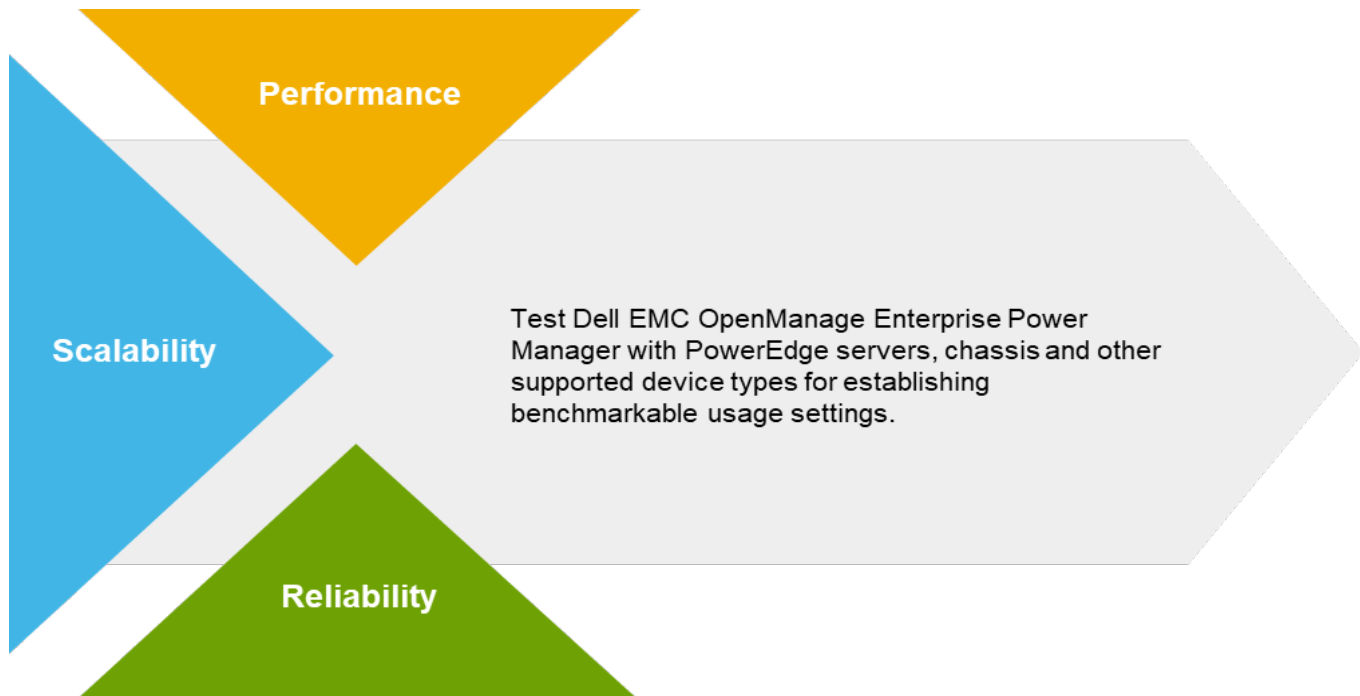


## Benchmark the Performance, Reliability and Scalability of Dell EMC OpenManage Enterprise Power Manager 2.0 in Data Center Environment



### Abstract

This technical white paper provides information about the best practices that you can follow while using Dell EMC OpenManage Enterprise Power Manager 2.0 to optimize the usage of supported devices in a data center environment.

August 2021

## Revisions

Date	Description
August 2021	Scalability whitepaper for Power Manager 2.0 release with OpenManage Enterprise version 3.6.1.

## Acknowledgements

This technical white paper is documented by the following members of the Dell EMC Server Engineering Team.

Authors: Ashish Singh, Aarsi Kumar, Jimmy Muraleedharan

Support: Shruthi Ravor

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2020-2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [8/4/2021] [Technical Document]

# Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents .....	3
Executive Summary .....	5
1 Introduction.....	6
2 Generic Recommendations.....	7
3 Test Environment and Methodology.....	9
4 Scalability and Performance tests for Power Manager 2.0.....	10
5 Addition of devices and groups in Power Manager.....	11
5.1 Configuration Details and Generic Recommendations .....	11
5.1.1 Configuration Details .....	11
5.1.2 Recommendations.....	11
5.2 Test Results in Scaled infrastructure (Devices and Groups) .....	12
5.2.1 Addition of devices and groups in Power Manager.....	12
5.2.2 Removal of devices and groups in Power Manager.....	12
6 Metric Collection.....	13
6.1 Configuration Details and Generic Recommendations .....	13
6.1.1 Configuration Details .....	13
6.1.2 Recommendations.....	13
6.2 Test Results.....	14
6.2.1 Small Deployment - CONFIG 1 .....	14
6.2.2 Large Deployments - CONFIG 2 and CONFIG 3.....	15
6.2.3 Large Deployments - CONFIG 4 and CONFIG 5.....	16
7 Power Policy (Power Capping).....	18
7.1 Configuration Details and Generic Recommendations .....	18
7.1.1 Configuration Details .....	18
7.1.2 Recommendations.....	18
7.2 Test Results in Scaled Infrastructure.....	18
7.3 Test Results for Power Policy Actions.....	19
8 Emergency Power Reduction.....	21
8.1 Configuration Details and Generic Recommendations .....	21
8.1.1 Configuration Details .....	21
8.1.2 Recommendations.....	21
8.2 Test Results in Scaled infrastructure (Devices and Groups) .....	21

8.3	Test Results for EPR actions.....	22
<b>9</b>	<b>Alert Thresholds .....</b>	<b>23</b>
9.1	Configuration Details and Generic Recommendations .....	23
9.1.1	Configuration Details .....	23
9.1.2	Recommendations.....	23
9.2	Test Results in Scaled Infrastructure.....	24
<b>10</b>	<b>Importing Physical Groups and Device Association from a CSV file .....</b>	<b>25</b>
10.1	Configuration Details and Generic Recommendations .....	25
10.1.1	Configuration Details .....	25
10.1.2	Recommendations .....	26
10.2	Test Results in Scaled Infrastructure.....	26
<b>11</b>	<b>Plugin Actions (Install, Disable, Enable, Uninstall).....</b>	<b>28</b>
11.1	Configuration Details and Generic Recommendations .....	28
11.1.1	Configuration Details .....	28
11.1.2	Recommendations .....	28
11.2	Test Results in Scaled Infrastructure.....	28
<b>12</b>	<b>Longevity Test Results .....</b>	<b>30</b>
12.1	Metric Collection .....	30
12.1.1	Configuration Details .....	30
12.2	Test Results.....	30
12.3	Recommendations.....	30
12.3.1	Disk Usage .....	31
12.3.2	Average CPU Utilization and Memory Utilization for Small Deployment .....	31
12.3.3	Average CPU Utilization and Memory Utilization for Large Deployments .....	32
<b>13</b>	<b>Troubleshooting.....</b>	<b>33</b>
13.1	Metric Collection happening at non-regular intervals in a scaled datacenter environment.....	33
13.2	Association of devices to Physical Groups may fail when trying to create more than 100 groups using csv import method.....	34
<b>14</b>	<b>Conclusion.....</b>	<b>35</b>
<b>15</b>	<b>Technical Support and Resources .....</b>	<b>36</b>
15.1	Related Resources .....	36

## Executive Summary

This technical whitepaper provides an overview about the performance, reliability, and scalability testing of Power Manager version 2.0 deployed on OpenManage Enterprise version 3.6.1. This whitepaper contains reports that are run to assess the product quality and stability of the following Power Manager features:

- Metric collection for Power and Thermal History, CPU Utilization, I/O Utilization, Memory Utilization, System Air Flow History, VM Power and PDU Instant Power.
- Power Cap configuration through Power Policy deployment.
- Emergency Power Reduction (EPR) configuration on Devices and Groups.
- Alert Thresholds configuration and monitoring on Devices and Groups.
- Import Physical Groups and Device Association from a CSV file.
- Addition or removal of Devices and Groups in Power Manager.
- Plugin actions such as install, enable, disable, and uninstall.

This document provides recommendations to achieve maximum performance using Power Manager in a scaled infrastructure.

# 1 Introduction

Dell EMC OpenManage Enterprise Power Manager is a plugin to Dell EMC OpenManage Enterprise (OME) console which uses fine-grained instrumentation to provide increased visibility over power consumption, anomalies, and utilization. Also, Power Manager alerts and reports on power and thermal events with servers, chassis, and groups consisting of servers and chassis. This enables increased control, faster response time, greater accuracy, and broader decision-making intelligence than would otherwise be possible.

This technical white paper summarizes Power Manager version 2.0 system testing environment, methodology, resultant data, and recommendations to achieve performance similar to the test results.

## 2 Generic Recommendations

Power Manager features are similar to OpenManage Enterprise that are dependent on the underlying hardware configurations and the resources allocated to the hosted Virtual Machine. To optimize the setup and task scheduling in Power Manager, prioritize the hardware allocation using the following guidelines:

### **Memory:**

Memory is a gating factor to a number of tasks which can be run simultaneously. It is recommended to manually increase the RAM from default value while running a maximum scaled configuration to allow more concurrent processes. 16GB RAM is the minimum requirement for Small scale deployments whereas 32GB is the minimum requirement for Enterprise scale deployments.

### **CPU:**

Power Manager leverages the services from OpenManage Enterprise for running multiple tasks by using different threads concurrently. It is recommended to add more cores for better performance since few of the database operations in Power Manager are CPU intensive. Increased hardware specifications, processor cores, and RAM yield to a better performance than what is detailed in this document. For Small Deployments, the minimum processor requirement is 4 cores whereas for Enterprise Deployments, the minimum requirement is 8 cores.

### **Network Bandwidth:**

Majority of the features in Power Manager are based on network interactions between the appliance and target devices. This information is then transmitted from the appliance to a client which is accessing it over the network. The time taken to execute a job depends on the available bandwidth and delay that exists in a network.

### **Timeout or Non-Responsive Targets:**

There are multiple cases where the communication from Power Manager to the target device may fail. There can be a performance impact on few of the Power Manager features; especially if this happens for many devices which are part of the task, for example, Metric Collection, Policy and EPR configurations. Possible scenarios include discovery with incorrect credentials (or non-administrator privileges), device malfunction or the management protocol is disabled in target devices. A WAN environment with minimum 10Mbps bandwidth and not more than 500ms delay is recommended for seamless functionality of OpenManage Enterprise and Power Manager tasks.

### **Hard Drive:**

The default drive space requirements for the appliance with Power Manager 2.0 remains as mentioned in OpenManage Enterprise version 3.6.1 User Guide. The minimum drive space requirement is 400 GB for Small and Enterprise Deployments.

Minimum system configuration required for OpenManage Enterprise Power Manager in scale environment







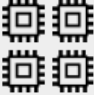

	Small Business	Enterprise Business
<b>Number of Devices</b>	 1000	 8000
<b>RAM Used</b>	 16 GB	 32 GB
<b>Hard Drive</b>	 400 GB	 400 GB
<b>Processor Cores</b>	 4	 8

Figure 1 Minimum system configuration requirements in scale environment



### 3 Test Environment and Methodology

The scalability, performance, reliability, and longevity tests for Power Manager Plugin 2.0 were performed on the following environments:

#### Small Business Customer's Environment:

This environment consists of 1000 managed systems in OpenManage Enterprise and the same set of devices monitored in Power Manager. CONFIG-1 has 1000 devices monitored in Power Manager plugin as part of Individual Devices, Static Groups and Physical Groups.

#### Enterprise Business Customer's Environment:

This environment consists of 8000 managed systems in OpenManage Enterprise and the same set of devices monitored in Power Manager. CONFIG - 2, 3, 4, and 5 are deployed as the test environment for Enterprise Business with 8000 devices monitored in Power Manager plugin as part of Individual Devices, Static Groups and Physical Groups.

Following is the list of configurations covered as part of this exercise.

Configuration details mentioned in all subsequent sections of this document are a subset of the same.

Table 1 Configuration Matrix for scalability test of Power Manager

Deployment Type	Configuration Theme	Configuration Name	Devices in Power Manager	Servers	Chassis	PDU's	VM	Monitored Groups
Small Business	Heterogenous Small Business Configuration	CONFIG-1	1000	500	80	20	400	200
Enterprise Business	Servers and Virtual Machines	CONFIG-2	8000	1000	-	-	7000	-
	Servers and PDU Devices	CONFIG-3	8000	7000	-	1000	-	300
	Servers, Chassis and PDU's	CONFIG-4	8000	7500	200	300	-	500
	Heterogenous Large Business Configuration	CONFIG-5	8000	4500	200	300	3000	400

**Note:** Tests are run at different Metric Gathering Intervals supported in Power Manager plugin such as 15 mins, 30 mins and 1-hour durations for 10 days.

## 4 Scalability and Performance tests for Power Manager 2.0

The primary focus of scalability and performance test was to collect and analyze various performance parameters of Power Manager version 2.0 for the following features:

- **Metric Collection:** Metric gathering for Power and Thermal History, CPU Utilization, I/O Utilization, Memory Utilization and System Air Flow History.
- **Power Policy:** Power Cap configuration through Power Policy deployment on Devices and Groups.
- **Emergency Power Reduction (EPR):** EPR Configuration on Devices and Groups.
- **Alert Thresholds:** Monitoring of Alert Thresholds (Power and Temperature) configured in Power Manager for Devices and Groups.

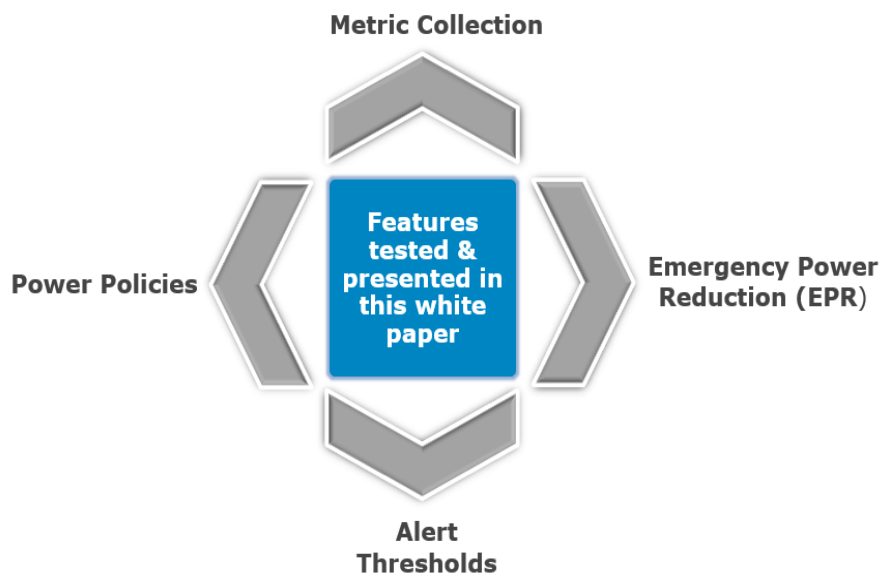


Figure 2 Power Manager features tested in scale environment

In addition to the core features mentioned above, a brief analysis is provided on different plugin actions (such as Install, Disable, Enable and Uninstall), modification of devices and groups added to Power Manager (addition and removal of Individual Devices and Groups) and on Import of Physical Groups and Device Association from a CSV file.

The tests were performed at various loads as mentioned in the configuration matrix at each section of the validated features. In addition to this, the longevity and reliability tests were performed, where the high priority features of Power Manager were executed at regular intervals over a period of 10 days. The performance data is collected and updated for these tests as well.

## 5 Addition of devices and groups in Power Manager

Add servers and chassis to Power Manager through **Add Device(s)** or **Add Group(s)** options in **Plugins > Power Management > Power Manager Devices** tab. Following test results include addition and removal of devices and groups in Power Manager for the test environments listed in Table 1 and recording the performance data.

### 5.1 Configuration Details and Generic Recommendations

#### 5.1.1 Configuration Details

Power Manager version 2.0 supports monitoring and management up to 8000 devices. You can add devices for monitoring and management either as Individual Devices, Static Groups or Physical Groups.

Tests for Individual Devices and Physical or Static Groups is performed on two configurations each. In **CONFIG-1**, 1000 devices were added to Power Manager and **CONFIG-4**, 8000 devices were added to Power Manager.

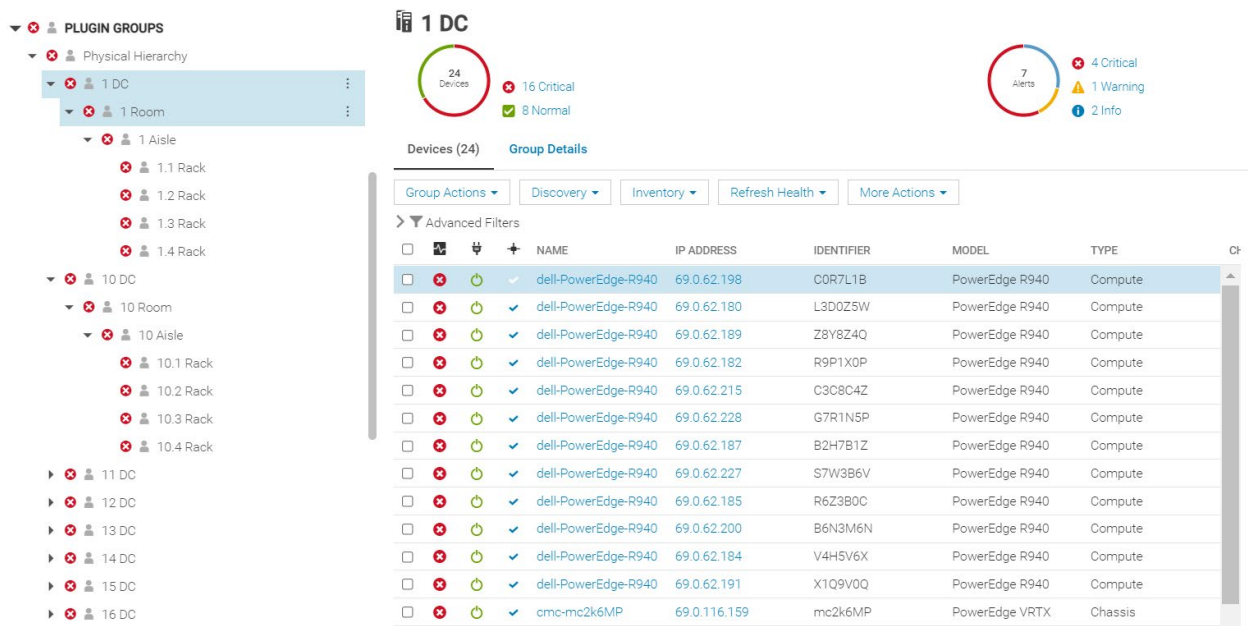


Figure 3 Nested Physical Groups added to Power Manager

#### 5.1.2 Recommendations

After analyzing the test results, following are the recommendations for adding Devices and Groups to Power Manager:

- It is recommended to create Static and Physical Groups with not more than 40 devices for addition to Power Manager for smooth and better management.
- Internet Explorer may take longer time to load the list of devices when compared to other browsers. Hence, it is recommended to use Google Chrome or Firefox browsers.
- Clearing browser cache can yield better results.

## 5.2 Test Results in Scaled infrastructure (Devices and Groups)

Below mentioned test results are on updating the list of devices and groups in Power Manager after clearing browser cache in Google Chrome version 91.0.4472.164.

### 5.2.1 Addition of devices and groups in Power Manager

Table 2 Test results on adding Devices and Groups into Power Manager

Configuration Name	Devices in Power Manager	Groups in Power Manager	Added As	Time Taken
CONFIG - 1	1000	-	Individual Devices	~ 3s
CONFIG - 1	1000	200	Static Groups	~ 3s
CONFIG - 4	8000	-	Individual Devices	~ 5s
CONFIG - 4	8000	500	Static Groups	~ 6s

### 5.2.2 Removal of devices and groups in Power Manager

Table 3 Test results on removing Devices and Groups from Power Manager

Configuration Name	Devices in Power Manager	Groups in Power Manager	Added As	Time Taken
CONFIG - 1	1000	-	Individual Devices	~ 3 s
CONFIG - 1	1000	100	Static Groups	~ 3 s
CONFIG - 4	8000	-	Individual Devices	~ 4s
CONFIG - 4	8000	500	Static Groups	~ 6s

**Note:** Tests are run for removal of devices and groups from Power Manager to verify when the setting for **Delete Power Manager Data** is configured as **No**.

## 6 Metric Collection

Metric details are captured from monitored servers, chassis, PDU's and virtual machines at regular intervals of 15, 30, or 60 minutes as per the configured Power Manager Settings. This metrics is represented through the Metrics and Monitoring History graphs:

- Power History
- Thermal History
- CPU Utilization
- I/O Utilization
- Memory Utilization
- System Air Flow History

View this historical data in **Power Management and Monitoring** tab in **Device Details** section for **Individual Devices** and in **Group Details** tab for Static and Physical Groups.

### 6.1 Configuration Details and Generic Recommendations

#### 6.1.1 Configuration Details

Metric Collection is supported on Individual Devices, Static Groups and Physical Groups in Power Manager. Metric gathering starts when the devices and groups are added to Power Manager or when the supported device types are Associated to a Rack Physical Group. The default setting of Metric gathering interval automatically changes to 30 minutes when more than 1000 devices are added to Power Manager.

#### 6.1.2 Recommendations

After analyzing the test results, following are the recommendations for Metric Collection:

- It is recommended to create Static and Rack Physical Groups with not more than 40 devices for Metric Collection to occur seamlessly.
- Make sure that the target devices are reachable to Power Manager and is responsive with respect to the management protocols. Higher number of unresponsive devices can have impact in metric collection due to the number of attempted retries.
- In case there are high number of unresponsive devices, it is recommended to have Metric gathering interval set to 60 minutes which will help avoid performance deterioration.

---

**Note:** The test results captured in all the following sections is for four hours of Metric Collection and the graphs are plotted for the CPU and Memory utilization of metric collection service.

---

## 6.2 Test Results

### 6.2.1 Small Deployment - CONFIG 1

Table 4 Metric Collection results for Small Deployment – Config1

Configuration Name	Devices in Power Manager	Servers	Chassis	PDUs	VMs	Metric Interval	Average Time Taken
CONFIG-1	1000	500	80	20	400	15 min	3 min

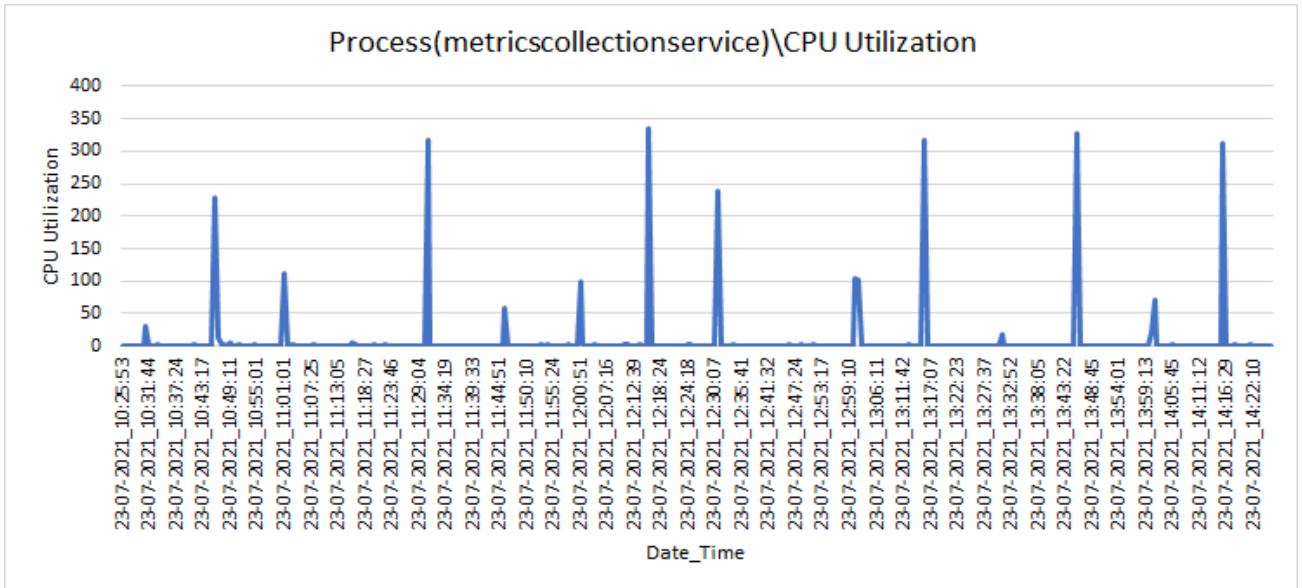


Figure 4 CPU Utilization for Small Deployment for 4 hours

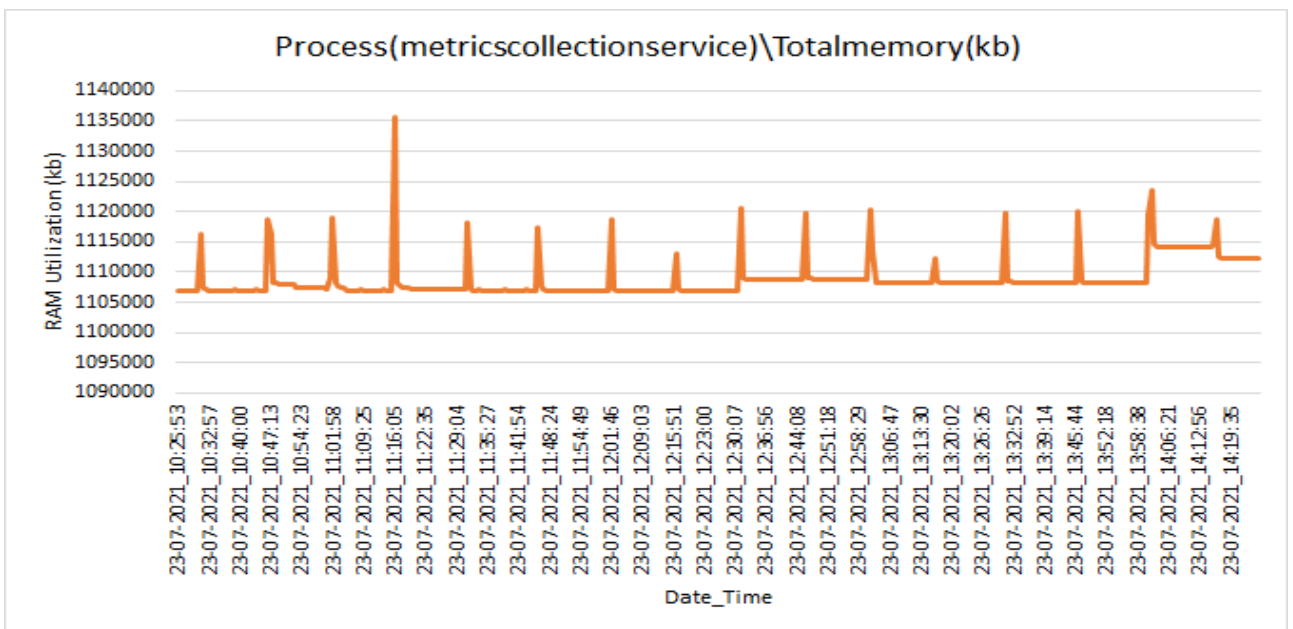


Figure 5 Memory Utilization for Small Deployment for 4 hours

## 6.2.2 Large Deployments - CONFIG 2 and CONFIG 3

Table 5 Metric Collection results for Config 2 and Config 3

Configuration Name	Devices in Power Manager	Servers	Chassis	PDUs	VMs	Metric Interval	Average Time Taken
CONFIG - 2	8000	1000	-	-	7000	30 min	6 min
CONFIG - 3	8000	7000	-	1000	-	30 min	6 min




Figure 6 CPU Utilization for Large Deployments (Configurations 2 and 3) for 4 hours

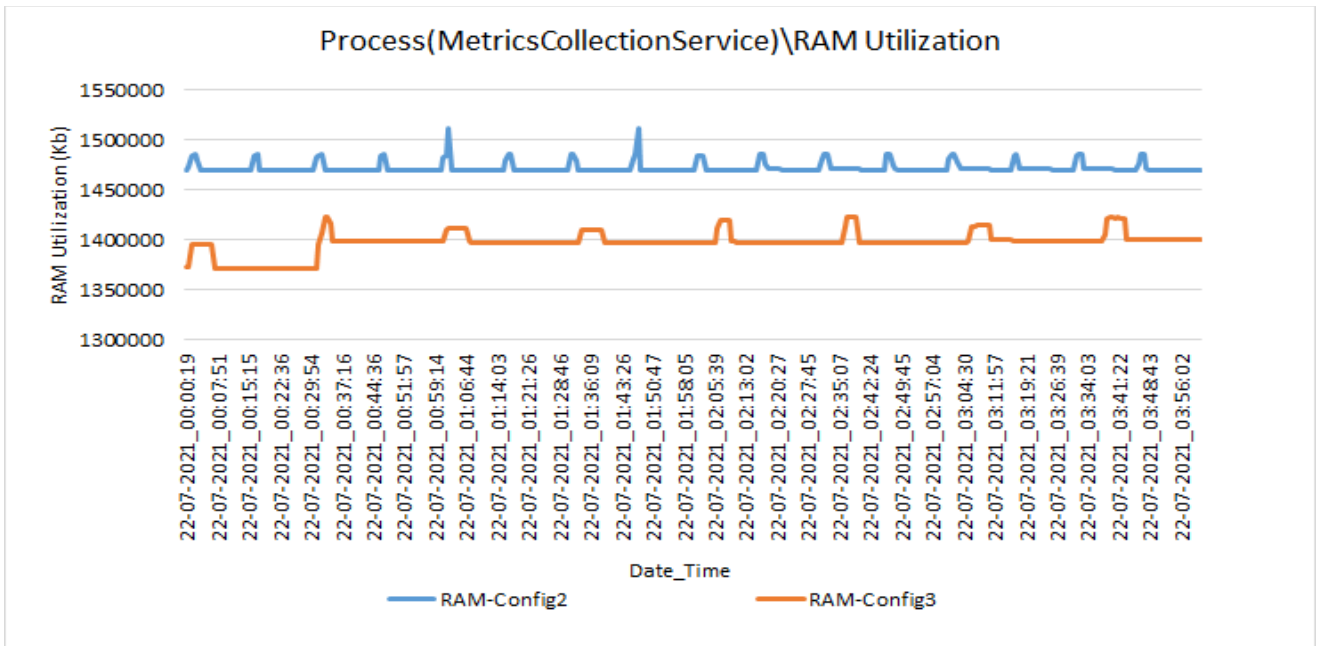


Figure 7 Memory Utilization for Large Deployments (Configurations 2 and 3) for 4 hours

### 6.2.3 Large Deployments - CONFIG 4 and CONFIG 5

Table 6 Metric Collection results for Config 4 and Config 5

Configuration Name	Devices in Power Manager	Servers	Chassis	PDUs	VMs	Metric Interval	Average Time Taken
CONFIG - 4	8000	7500	200	300	-	30 min	3 min
CONFIG - 5	8000	4500	200	300	3000	30 min	6 min



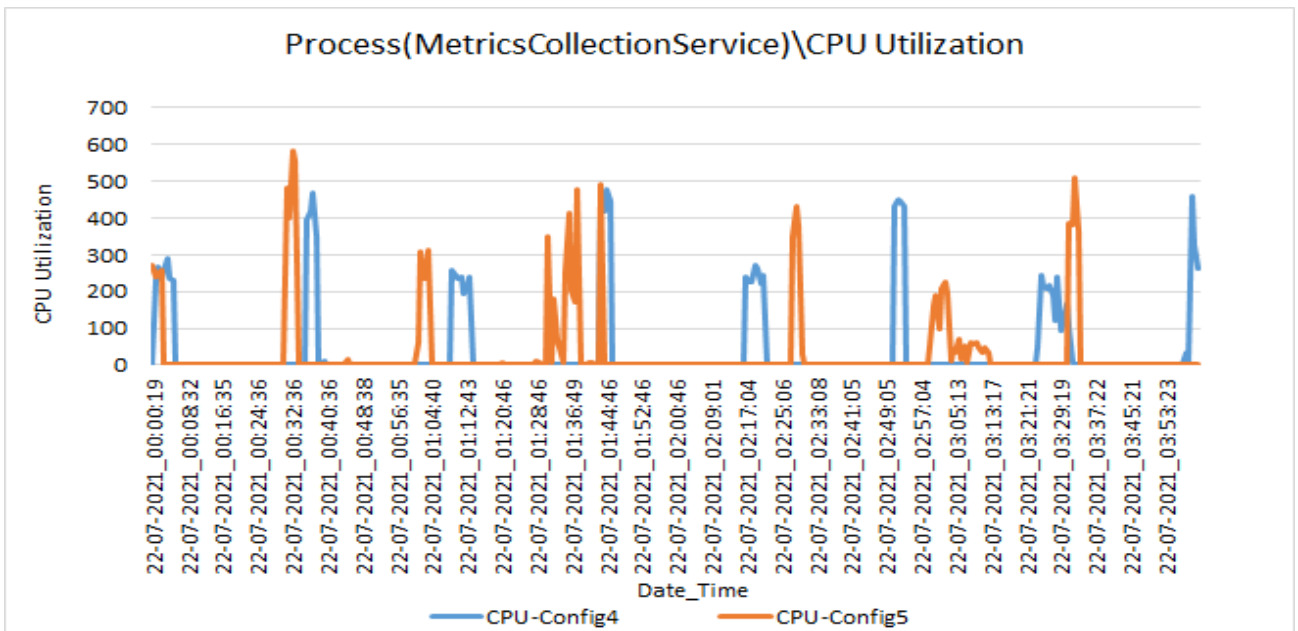


Figure 8 CPU Utilization for Large Deployments (Configurations 4 and 5) for 4 hours

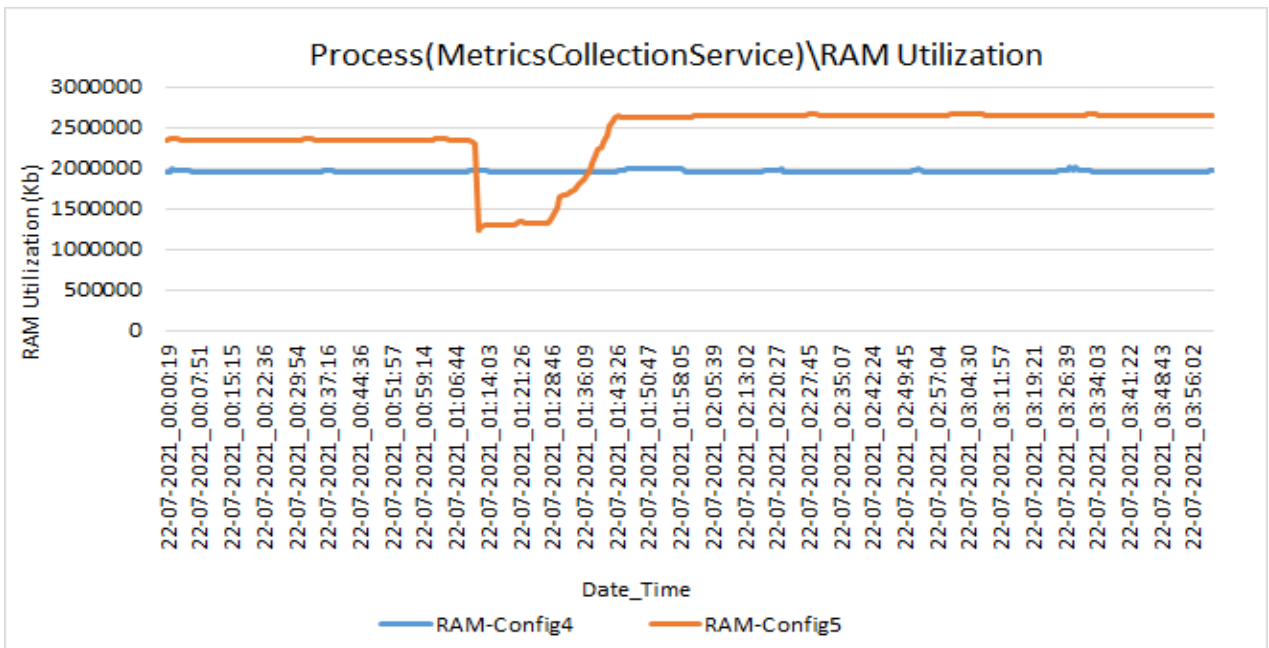


Figure 9 Memory Utilization for Large Deployments (Configurations 4 and 5) for 4 hours

## 7 Power Policy (Power Capping)

Create and apply power policies for regulating the power consumption of a device or a group of devices. Power Manager supports applying a Power Policy on a single target device or single static or physical group at a time from the policy creation wizard.

### 7.1 Configuration Details and Generic Recommendations

#### 7.1.1 Configuration Details

Applying a Power Policy is supported on individual devices, static groups, and physical groups. Tests were performed by creating 400 Power Policies and applying them on 400 devices sequentially.

#### 7.1.2 Recommendations

After analyzing the test results, following are the recommendations for using Power Policy feature:

- It is recommended to create Static and Rack Physical Groups with not more than 40 devices for all Power Policy actions.
- Ensure that the target devices are reachable to Power Manager and is responsive with respect to the management protocols. Higher number of non-responsive devices and network latency can cause failures during Power Policy deployments.
- If policies are applied at individual device level, ensure that you select a maximum of up to 200 policies to perform any simultaneous policy related actions on the devices for maximum performance optimization.
- If multiple Power Policies are applied on an Individual Device or a Group, then there is an increase in the computation time for Power Manager that results in slightly higher time duration for applying policies in a scaled infrastructure.

### 7.2 Test Results in Scaled Infrastructure

Test results were very similar when Power Policies related actions were performed on individual devices, Static Groups, and Physical groups. Hence the data captured for all the following actions are on Physical Groups. Following actions were tested in scaled infrastructure for Enterprise business Large scale deployment.

- Apply Power Policy
- Disable Power Policy
- Enable Power Policy
- Delete Power Policy

Power Policy action is applied on 400 devices that are in a single Physical Group (Datacenter level) across 10 racks, to ensure that the policy action is triggered simultaneously for all devices without any lag or obstruction.

The following graphs represent the CPU and Memory utilization of Task Execution Service during different power policy actions.

### 7.3 Test Results for Power Policy Actions

Table 7 Test Results for Power Policy related actions for Heterogeneous Large Business Configuration

Devices in Power Manager	Tested Devices	Action	Time Taken
8000	400 Servers + 30 Chassis	Create policies	1 m 52 s
8000	400 Servers + 30 Chassis	Disable policies	1 m 9 s
8000	400 Servers + 30 Chassis	Enable policies	1 m 14 s
8000	400 Servers + 30 Chassis	Delete policies	1 m 7 s




Figure 10 CPU Utilization for Power Policy related actions




Figure 11 Memory Utilization for Power Policy related actions

## 8 Emergency Power Reduction

Emergency Power Reduction (EPR) is a feature in Power Manager which helps you to reduce power consumption of devices immediately during a power emergency. Apply EPR on Individual Devices using the **Apply EPR** option that is available in **Power Management and Monitoring > Policies and EPR** tab. And, apply EPR on Static and Physical Groups using the **Apply EPR** option that is available in **Group Details > Policies and EPR** tab. Throttle and Shut Down are two modes of applying EPR on target devices and groups.

**EPR Throttle:** On applying EPR (Throttle), all the selected devices consume extremely low power and hence impacts the performance of the target device.

**EPR Shut Down:** On applying EPR (Shut Down), all the selected devices monitored as Individual device or as part of the groups are gracefully shut down.

### 8.1 Configuration Details and Generic Recommendations

#### 8.1.1 Configuration Details

Applying the EPR is supported on individual devices, static groups and physical groups. Test for individual devices has been performed by applying 400 EPR's (sequentially) for 400 servers in Power Manager, from REST interface with a delay of 5 seconds set between subsequent EPR actions. In case of groups, the devices were distributed across 5 Static and Physical Groups (40 servers each) and the EPR was applied sequentially when the preceding task was completed. Data is collected for the EPR Throttle and EPR Shut Down actions from Power Manager.

#### 8.1.2 Recommendations

After analyzing the impact and test results, following are the recommendations for using the EPR feature:

- EPR Throttle action throttles down the device power down to an extremely low level, which impacts the performance of target device. Hence the recommendation is to use this feature only in case of an emergency.
- It is recommended to create Static and Rack Physical Groups with a maximum 40 devices for applying an EPR.
- Make sure that the target devices are reachable to Power Manager and is responsive with respect to the management protocols. Higher number of non-responsive devices or network latency can cause performance impact while applying an EPR. For more information about the protocols used in Power Manager, see the OpenManage Enterprise Power Manager User's Guide available on the support site.
- To disable the EPR's, select a maximum of 200 EPR's at a time for better performance and results.

### 8.2 Test Results in Scaled infrastructure (Devices and Groups)

Test results were very similar when the EPR actions were performed on Individual Devices, Static Groups and Physical Groups.

The following graphs represent the CPU and Memory utilization of task execution service during different EPR actions.

### 8.3 Test Results for EPR actions

Table 8 Test Results for EPR actions for Heterogeneous Large Business Configuration

Devices in Power Manager	Tested Devices	EPR Action	Time Taken
8000	400 Servers	Enable EPR Throttle	8 m 30 s
8000	400 Servers	Disable EPR Throttle	5 m 42 s
8000	400 Servers	Enable EPR Shutdown	8 m 18 s
8000	400 Servers	Disable EPR Shutdown	8 m 14 s




Figure 12 CPU Utilization for EPR Actions

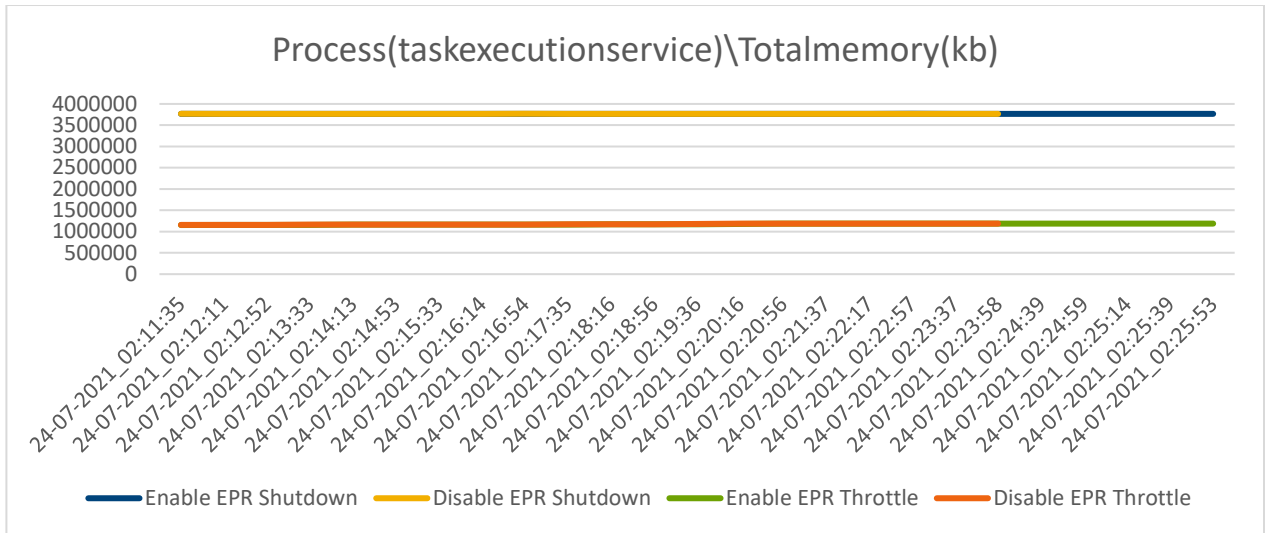


Figure 13 Memory Utilization for EPR Actions

## 9 Alert Thresholds

Power Manager leverages the Events and Alerts feature in OpenManage Enterprise and supports different alerts such as Threshold violation alerts, Group membership change alerts, Policy Violation alerts and so on. Power and Temperature Threshold violation alerts are used to notify the administrator when the power consumption and temperature has exceeded certain configured values so that the administrator can take immediate actions.

### 9.1 Configuration Details and Generic Recommendations

#### 9.1.1 Configuration Details

Violation of Power and Temperature thresholds' and corresponding event generation has been tested simultaneously on 250 devices in a configuration where 8000 devices are monitored in Power Manager with no alert actions configured.

#### 9.1.2 Recommendations

After analyzing the test results, following are the recommendations for Threshold Alerting functionality of Power Manager.

---

**Note:** The benchmarks mentioned in this document for Alert Thresholds are common for OpenManage Enterprise console.

---

- Reception of alerts is supported with 5,000 alerts per minute for 60 minutes without any drop when no alert actions are configured. This implies that the appliance receives 85 alerts per second continuously for 1 hour without any alert drop.
- With one alert action (email or trap forwarding) configured and mapped with 8,000 discovered devices, reception of alerts is supported up to 85 alerts per second for a duration of 60 minutes with a possibility of 20 percent alert drop. Restarting OME services help to receive a greater number of alerts.
- With all the alert actions configured; reception of alerts is supported till 9 alerts per second with a possibility of 25% alert drop, if continued for long duration. Restarting OME services help to receive a greater number of alerts.
- Key factors determining alert reception in case of configured alert actions are:
  - Number of devices selected.
  - Time interval between alert reception.
  - Number of alert actions configured.

The graphs in the following test results represent the CPU and Memory utilization of event processing service captured during alert reception.

## 9.2 Test Results in Scaled Infrastructure

Table 9 Results for Alert Thresholds in scaled environment

Configuration Name	Devices in Power Manager	Tested Devices	Alerts Generated	Time Taken
CONFIG - 4	8000	250	500	2 m 3 s

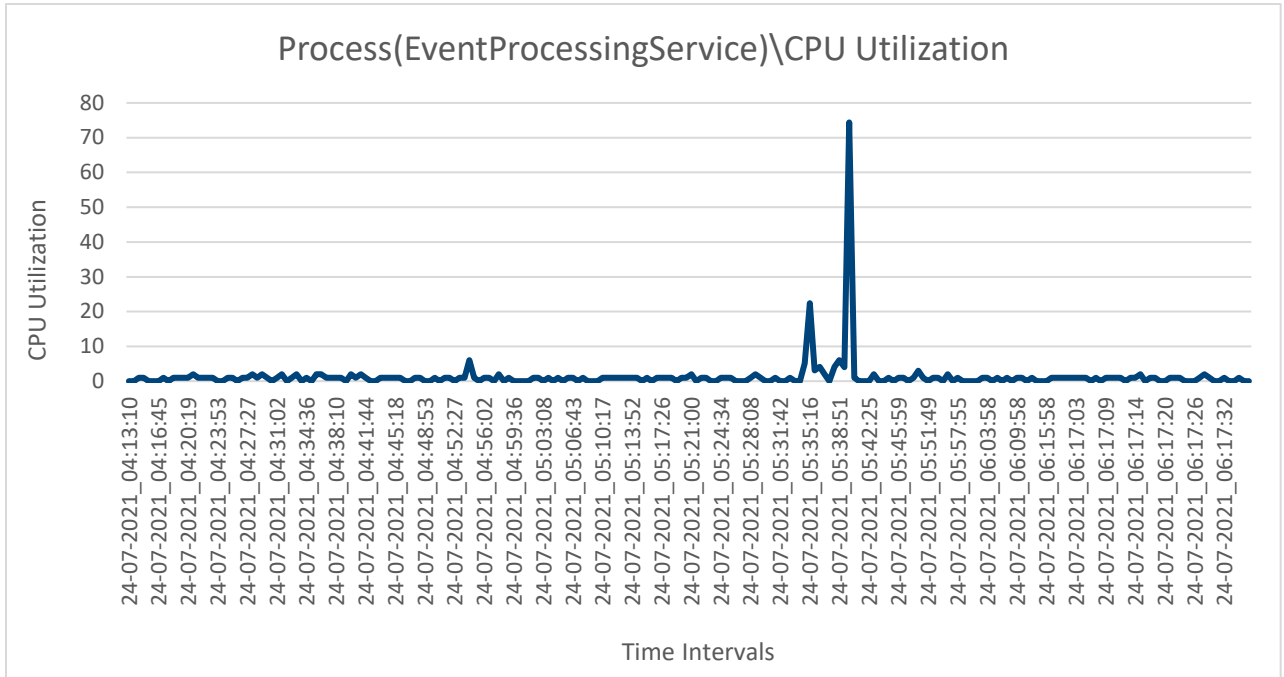


Figure 14 CPU Utilization during Alert Reception in scaled infrastructure




Figure 15 Memory Utilization during Alert Reception in scaled infrastructure



# 10 Importing Physical Groups and Device Association from a CSV file

Creation of physical group hierarchies for existing devices in a data center can be a time-consuming activity and hence, Power Manager facilitates importing the devices as physical groups through a csv file.

## 10.1 Configuration Details and Generic Recommendations

### 10.1.1 Configuration Details

Tests are conducted to verify the appliance resource utilization and task execution time after creating 180 nested physical groups (18 Datacenters, 18 Rooms, 18 Aisles, and 126 Racks) and associating the physical groups to 3150 devices through **Import** option available in **Plugins > Power Management > Power Manager Devices > Physical Groups** tab. The progress of importing physical groups and device association from CSV file is referenced from **Physical Group CSV Import Task** (System Generated) in **Monitor > Jobs** page.

Following is a sample of hierarchy followed for physical group creation with each Rack associated to 25 devices each and 175 devices per datacenter.

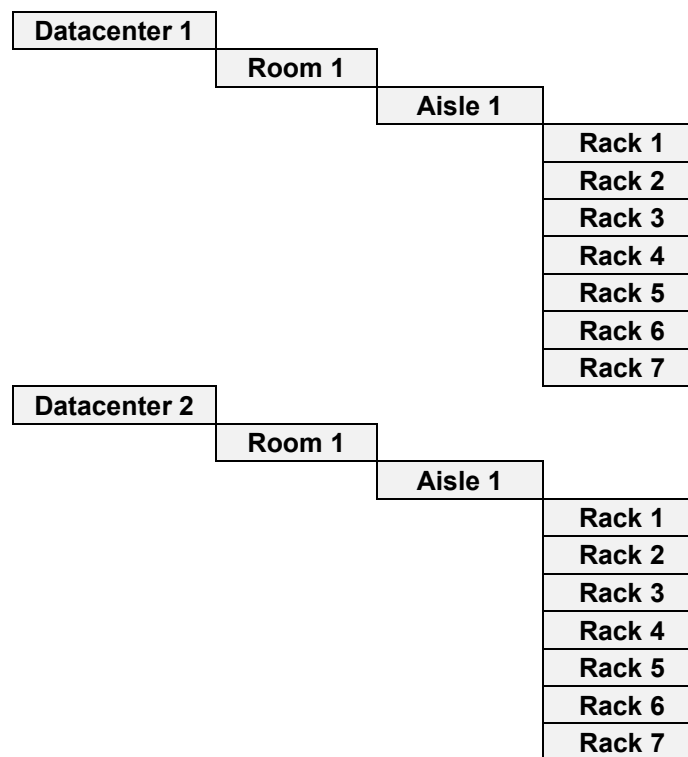


Figure 16 Physical group hierarchy

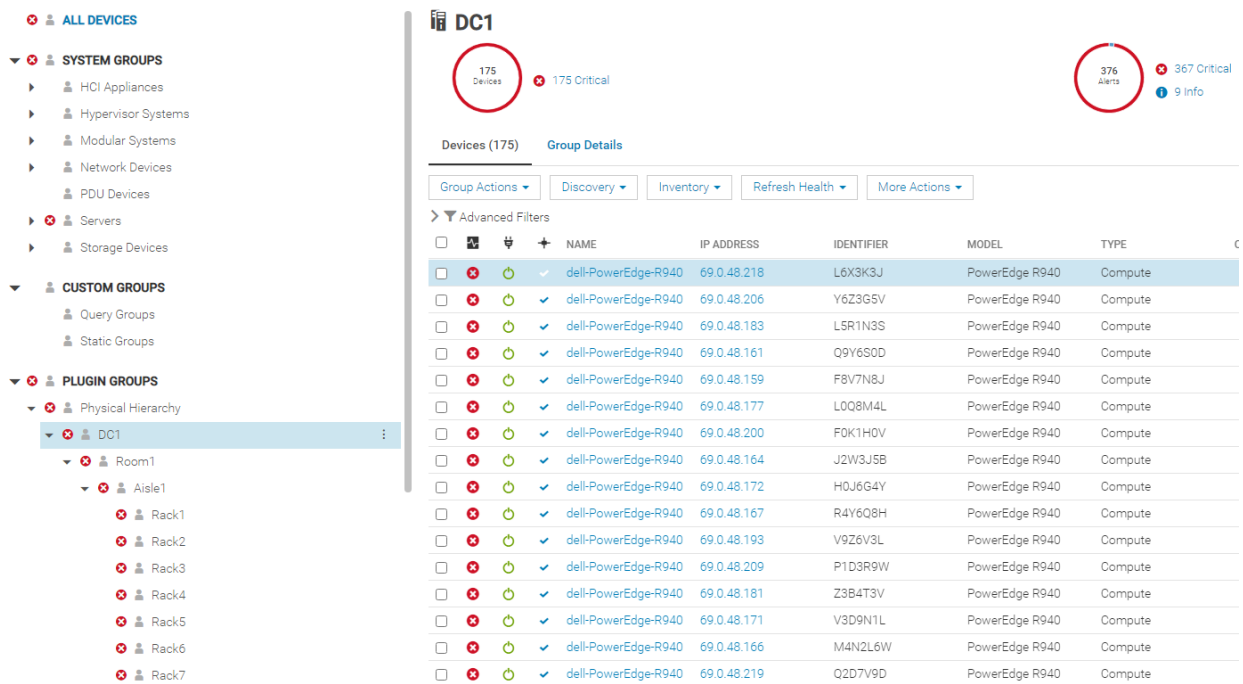


Figure 17 Hierarchy created after Importing Physical Groups and Device Association from CSV file

## 10.1.2 Recommendations

After analyzing the test results, following are the recommendations for utilizing the feature to Import Physical Groups from CSV file in Power Manager.

- For better performance, the physical hierarchies and device associations can be split into multiple files and imported as different tasks one after the other.
- Time taken for physical group creation is comparatively higher for nested groups, when compared to the rack groups created at root level.
- Time taken for device association is comparatively higher when there are larger number of devices discovered in OpenManage Enterprise.
- For better results in a scaled infrastructure, it is recommended not to execute other jobs (such as Policies and EPR application) which leverage Task Execution Service in the appliance, while import of physical groups and device association is in progress from a CSV file.

**Note:** Import of multiple files are not supported at the same time. Import option is disabled in **Plugins > Power Management > Power Manager Devices > Physical Groups** tab until the **Physical Group CSV Import Task** is completed.

## 10.2 Test Results in Scaled Infrastructure

Table 10 Results for Physical Group CSV Import Task in scaled in infrastructure

Configuration Name	Devices in OM Enterprise	Number of Nested Physical Groups	Number of Devices Associated	Time Taken
CONFIG - 1	1000	60	1000	35m
CONFIG - 4	8000	180	3150	2h 9m

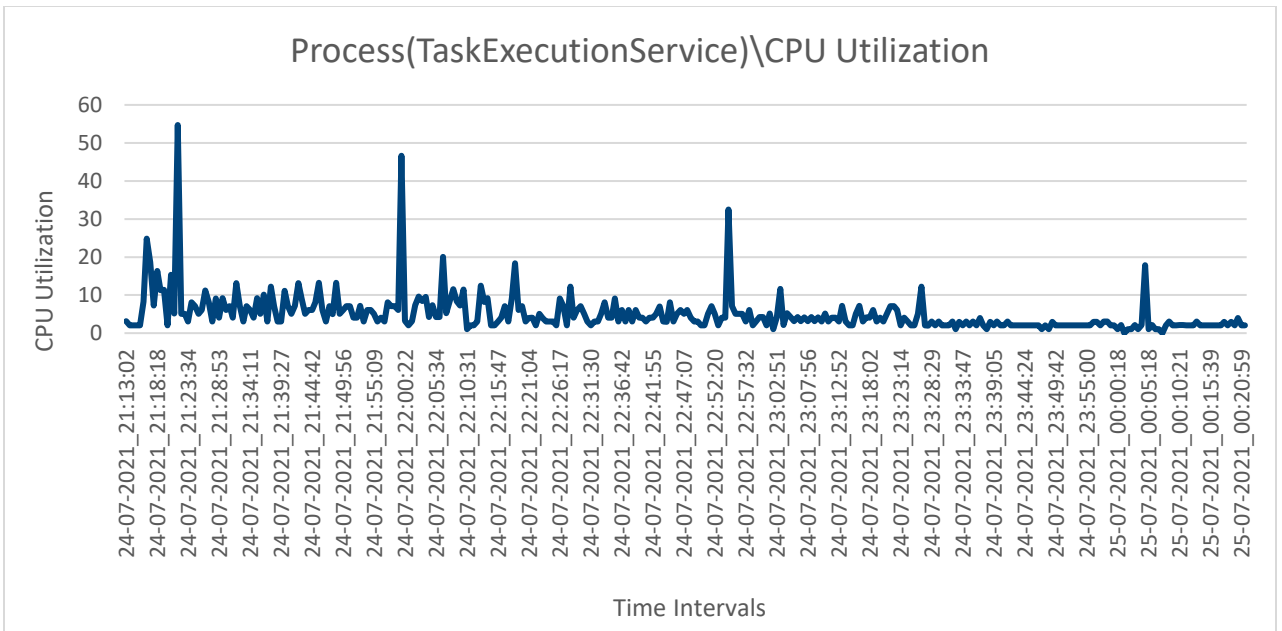


Figure 18 CPU Utilization during CSV Import (CONFIG-4) in scaled infrastructure

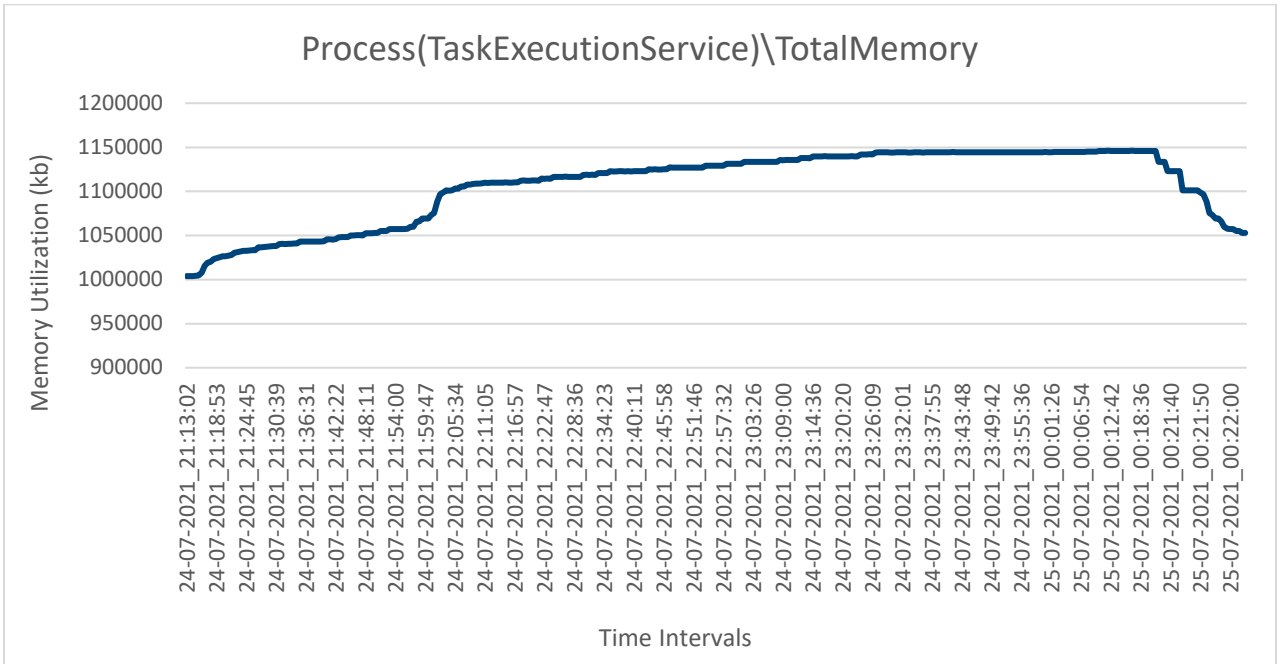


Figure 19 Memory Utilization during CSV Import (CONFIG-4) in scaled infrastructure

# 11 Plugin Actions (Install, Disable, Enable, Uninstall)

## 11.1 Configuration Details and Generic Recommendations

### 11.1.1 Configuration Details

Tests are run on a Small Business Environment with 1000 devices in Open Manage Enterprise and the same is monitored in Power Manager as part of 50 static and physical groups and on Enterprise Business Environment with 8000 devices monitored in Power Manager as part of 500 nested groups.

### 11.1.2 Recommendations

After analyzing the impact and test results, following are the recommendations for different plugin actions:

- Perform any plugin actions when none of the OpenManage Enterprise or Power Manager specific tasks are running.
- Ensure that all the Power Policies and EPR's are either disabled or deleted before performing a Disable or Uninstall action of Power Manager.

## 11.2 Test Results in Scaled Infrastructure

This section displays the time taken for completing different plugin actions on Power Manager 2.0 and corresponding OpenManage Enterprise services to respond in a scaled infrastructure. Tests are run in scaled OpenManage Enterprise appliance with Metric data available in individual device and group level for a week, Power Policies and Alert Thresholds configured for all monitored devices in Power Manager and when none of the Open Manage Enterprise tasks were in-progress.

**Note:** The time taken for download is dependent on the network bandwidth. Test result mentioned below is with a bandwidth of 1Gbps.

Table 11 Test Results for different Plugin Actions in a scaled infrastructure

Plugin Action	Devices in OpenManage Enterprise	Devices in Power Manager	Groups in Power Manager	Number of Power Policies	Number of Alert Thresholds	Time Taken
Download	8000	N/A	N/A	N/A	N/A	14s
Install	1000	N/A	N/A	N/A	N/A	1m 18s
	8000	N/A	N/A	N/A	N/A	1m 20s
Disable	1000	1000	25	1000	1000	59s
	8000	8000	500	4000	3000	1m 3s
Enable	1000	1000	25	1000	1000	1m 15s
	8000	8000	500	4000	3000	1m 18s

Plugin Actions (Install, Disable, Enable, Uninstall)

Uninstall	1000	1000	25	1000	1000	2m 12s
	8000	8000	500	4000	3000	2m 25s

## 12 Longevity Test Results

The primary goal for running these tests is to validate the reliability of Power Manager with OpenManage Enterprise when the appliance is utilized continuously for a defined period.

### 12.1 Metric Collection

#### 12.1.1 Configuration Details

Longevity test was run for a duration of 10 days. This test was run on individual devices and devices part of static and physical groups to monitor the performance in Small and Large deployment environments of OpenManage Enterprise. Following is the list of configurations where the test is executed.

Table 12 List of configurations deployed for Longevity Test

Deployment Type	Configuration Theme	Configuration Name	Devices in Power Manager	Servers	Chassis	PDU's	VMs	Monitored Groups
Small Business	Heterogenous Small Business Configuration	<b>CONFIG-1</b>	<b>1000</b>	500	80	20	400	200
Enterprise Business	Servers, Chassis and PDU's	<b>CONFIG-4</b>	<b>8000</b>	7500		300	-	500
	Heterogenous Large Business Configuration	<b>CONFIG-5</b>	<b>8000</b>	4500	200	300	3000	400

### 12.2 Test Results

- All monitored configurations have default OpenManage Enterprise tasks such as Default Inventory Task, and Global Health Task running on scheduled intervals.
- The time taken for Metric Collection remained constant for a period of 10 days for the supported number of devices.
- An uninterrupted service was observed while performing the Longevity test.
- All the devices which were monitored in longevity infrastructure was reachable to Power Manager and was responsive with respect to the management protocols. Higher number of non-responsive devices or network latency can impact in metric gathering intervals due to an increase in the number of retries attempted.
- CPU Utilization in the following graphs indicate 100 percent utilization of one processor core in the appliance.
- The time taken to open the OME console and to access different pages, slowly increases when more data is loaded to the database of the appliance. The test data was captured after every third day.

### 12.3 Recommendations

- To reduce the time taken in opening the User Interface (UI) console, restart the appliance.

- Clear the browser cache to yield better results.

The following graphs represent the Total CPU and Total Memory utilization of the appliance during the monitored period of longevity testing.

### 12.3.1 Disk Usage

Table 13 Disk Usage of all tested configurations (Longevity)

Deployment Type	Configuration Name	Devices in Power Manager	Servers	Chassis	PDU	VMs	Monitored Groups	Disk Usage
Small Business	CONFIG-1	1000	500	80	20	400	200	38.00 GB
Enterprise Business	CONFIG-4	8000	7500	200	300	-	-	44.98 GB
	CONFIG-5	8000	4500	200	300	3000	-	37.46 GB

### 12.3.2 Average CPU Utilization and Memory Utilization for Small Deployment




Figure 20 Total CPU and RAM Utilization for Metric Collection (Longevity) in CONFIG – 1

### 12.3.3 Average CPU Utilization and Memory Utilization for Large Deployments

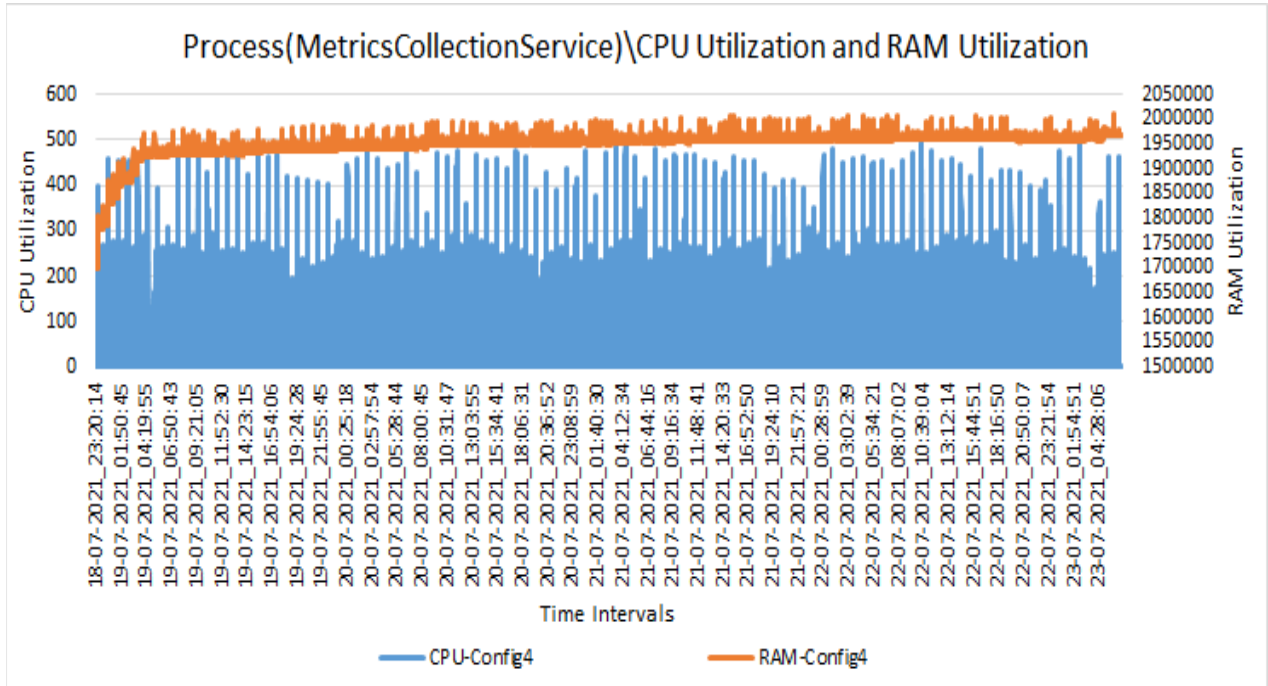


Figure 21 Total CPU and RAM Utilization for Metric Collection (Longevity) in CONFIG-4




Figure 22 Total CPU and RAM Utilization for Metric Collection (Longevity) in CONFIG-5



## 13 Troubleshooting

### 13.1 Metric Collection happening at non-regular intervals in a scaled datacenter environment

#### Issue

In the Metrics and Monitoring graphs for Individual Devices and Static Groups, the metric collection has happened at non-regular intervals when compared to the configuration set in **Power Manager Setting** for Metric Gathering Interval.

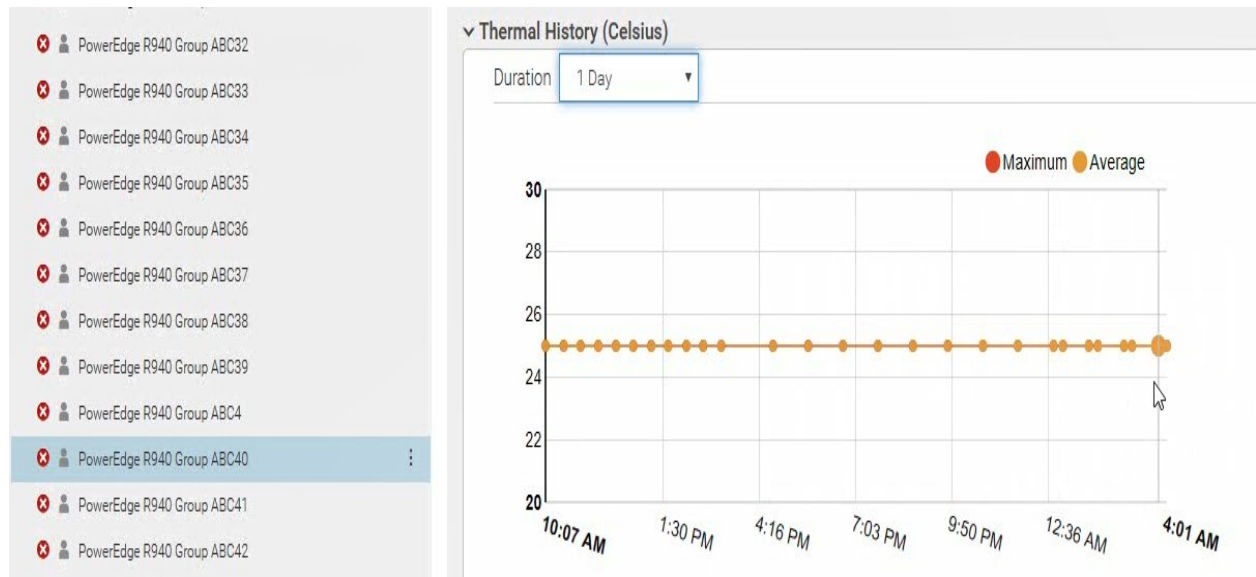


Figure 23 Metric Collection happening at non-regular intervals irrespective of the configuration set in Power Manager Setting for Metric Gathering Interval.

#### Resolution

Metric collection can happen at irregular intervals due to the following conditions in a datacenter environment.

- More than 10 percent of the Power Manager monitored devices are not reachable.
- Insufficient VM resources allocated or the resource consumption of guest VM is much higher that would cause the appliance to slow down.
- Network bandwidth lesser than 10Mbps or latency greater than 500ms.

To rectify the issue, ensure that the following conditions are met:

- All the monitored devices are reachable and responsive to the management protocols. For more information about the management protocols, see *OpenManage Enterprise Power Manager User's Guide*.
- Allocate sufficient resources to the VM.
- Ensure that there are no network latency issues.

## 13.2 Association of devices to Physical Groups may fail when trying to create more than 100 groups using csv import method

### Issue

When you import a CSV file with more than 100 groups and devices, the group creation is successful, but the association of devices to the groups may fail, with the following error message on the Jobs page under Physical Group CSV Import Task job: "Unable to complete the action because the value entered for GroupId is invalid."

### Resolution

To create all the groups successfully through CSV file, ensure that you perform one of the following actions:

- Re-import the CSV file

OR

- Create a CSV file specifying group details and import this file.
- Ensure that the import job is successful.
- Create another CSV file specifying device association to those groups, and then import the CSV file.

## 14 Conclusion

Dell EMC provides its customers with products that simplify and streamline their IT processes, and freeing administrator's time to focus on activities that help grow the business. This technical white paper provides comprehensive information about the test environment, recommendations, and test results in the form of neatly arranged infographics to make it easy for reading, scanning, and understanding. To maximize utilization, special notes and cautions are specified, wherever necessary. For more information about different Dell EMC PowerEdge servers, see the [Dell PowerEdge Servers Portfolio Guide](#).

## 15 Technical Support and Resources.

- [Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.
- To watch quick and short videos about handling the PowerEdge server components, visit the [QRL video website](#).
- [Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

### 15.1 Related Resources

- Knowledge Base for Dell EMC OpenManage Enterprise [HTML](#)
- Knowledge Base for Dell EMC OpenManage Enterprise Power Manager and Power Center [HTML](#)
- Dell EMC OpenManage Enterprise Power Manager Version 2.0 User's Guide [PDF HTML](#)
- Dell EMC OpenManage Enterprise Power Manager RESTful API Guide version 2.0 [PDF HTML](#)
- Dell EMC OpenManage Enterprise Power Manager 2.0 Release Notes [PDF HTML](#)