

Telemetry Streaming with iDRAC9 – Storage Reports

Abstract

Dell EMC PowerEdge Servers with iDRAC9 4.x Datacenter stream data to help IT administrators better understand the inner workings of their server environment. This white paper explains how to use the Telemetry Storage reports for getting the details of the drives and cards on the PowerEdge servers.

July 2020

Revisions

Date	Description
July 2020	Initial release

Acknowledgments

Authors: Sailaja Mahendrakar, Sankara Gara, Cyril Jose, Texas Roemer, Sankunny Jayaprasad.

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 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [7/27/2020] [White Paper] [ID 430]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	3
Executive summary.....	4
1 Introduction.....	5
1.1 Prerequisites.....	5
2 Telemetry Storage Metric Reports	6
2.1.1 StorageDiskSMARTData Metric Report.....	7
2.1.2 NVMeSMARTData Metric Report.....	8
2.1.3 StorageSensor Metric Report.....	9
2.2 Supported Devices for Storage Metric Reports.....	9
2.3 Special Cases of Hot plug Storage Devices.....	9
2.4 Workflow Examples	10
A Technical support and resources	14

Executive summary

Beginning with iDRAC9 v4.00.00.00 firmware and a Datacenter license, IT managers can integrate advanced server hardware operation telemetry into their existing analytics solutions. Telemetry is provided as granular, timeseries data that is streamed, or pushed, compared to inefficient, legacy polling, or pulled, methods. The advanced agent-free architecture in iDRAC9 provides over 180 data metrics that are related to server and peripherals operations. Metrics are precisely timestamped and internally buffered to allow highly efficient data stream collection and processing with minimal network loading. This comprehensive telemetry can be fed into analytics tools to predict failure events, optimize server operation, and enhance cyber resiliency.

This paper details the telemetry data from SSD SMART drives and NVMe SMART drives, which could be used to assess the health and state of the drives on the server.

1 Introduction

Telemetry Storage reports are helpful to assess the health and state of the various drives on the server and becomes necessary as the number of drives increase on the server. One key set of metric reports are the Storage reports. Storage reports provide data for both SSD Smart drives as well as NVMe SMART drives. This data can be used to assess the health and state of the drives on the server. Three telemetry storage reports are available to get the details of the drives.

- Storage SMART disk data (StorageDiskSMARTData)
- NVMe SMART data (NVMeSMARTData)
- Storage Sensor (StorageSensor)

1.1 Prerequisites

- The Telemetry feature is available on iDRAC9 firmware version 4.00.00.00 or above and requires a Datacenter license.

2 Telemetry Storage Metric Reports

The health state of the SSD drives behind PERC storage controller can be assessed with the telemetry metric report named "StorageDiskSMARTData". And similarly, the health state of the PCIeSSD devices which are not controlled by SWRAID controller, can be derived with the telemetry metric report named "NVMeSMARTData". The temperature information of storage devices can be found from the telemetry metric report named "StorageSensor".

Below are the Telemetry storage metric reports supported in iDRAC 4.00 release:

- StorageDiskSMARTData: PERC SSD disk smart information
- NVMeSMARTData: PCIeSSD/NVMe device smart information
- StorageSensor: Temperature reading for internal server storage devices

Table 1 Telemetry storage metric reports

Telemetry Storage Reports	Metrics	Description
StorageDiskSMARTData	CRCErrrorCount,CommandTimeout,CurrentPendingSectorCount,DriveTemperature,ECCERate,EraseFailCount,ExceptionModeStatus,MediaWriteCount,PercentDriveLifeRemaining,PowerCycleCount,PowerOnHours,ProgramFailCount,ReadErrorRate,ReallocatedBlockCount,UncorrectableErrorCount,UncorrectableLBACount,UnusedReservedBlockCount,UsedReservedBlockCount,VolatileMemoryBackupSourceFailures	SSD SMART information.
NVMeSMARTData	AvailableSpare,AvailableSpareThreshold,CompositeTemperature,ControllerBusyTimeLower,ControllerBusyTimeUpper,CriticalWarning,DataUnitsReadLower,DataUnitsReadUpper,DataUnitsWrittenLower,DataUnitsWrittenUpper,HostReadCommandsLower,HostReadCommandsUpper,HostWriteCommandsLower,HostWriteCommandsUpper,MediaDataIntegrityErrorsLower,MediaDataIntegrityErrorsUpper,NumOfErrorInfoLogEntriesLower,NumOfErrorInfoLogEntriesUpper,PercentageUsed,PowerCyclesLower,PowerCyclesUpper,PowerOnHoursLower,PowerOnHoursUpper,UnsafeShutdownsLower,UnsafeShutdownsUpper	PCIeSSD/NVMe device smart information.
StorageSensor	TemperatureReading	Temperature information for server storage internal devices.

2.1.1 StorageDiskSMARTData Metric Report

Table 2 Metrics of the StorageDiskSmartData telemetry report and the corresponding description of the metrics:

Telemetry Metric Report	Metrics	Description of the Metric
StorageDiskSMARTData	CRCErrCount	CRC error count
	CommandTimeout	The number of aborted operations due to disk timeout
	CurrentPendingSectorCount	
	DriveTemperature	Drive temperature in celsius
	ECCERate	Uncorrected read errors reported
	EraseFailCount	Erase fail count
	ExceptionModeStatus	Exception mode status
	MediaWriteCount	Media (SSD) write count
	PercentDriveLifeRemaining	Drive Life Remaining (DLR) in percentage
	PowerCycleCount	This attribute indicates the count of full hard disk power on/off cycles.
	PowerOnHours	The raw value of this attribute shows total count of hours (or minutes, or seconds, depending on manufacturer) in power-on state
	ProgramFailCount	Program fail count since drive was deployed
	ReadErrorRate	Read error rate
	ReallocatedBlockCount	Reallocated block count
	UncorrectableErrorCount	Uncorrectable error count
	UncorrectableLBACount	Uncorrectable LBA(Logical block addressing) count
	UnusedReservedBlockCount	Unused reserved block count
	UsedReservedBlockCount	Used reserved block count
VolatileMemoryBackupSourceFailures	Volatile memory backup source failures	

2.1.2 NVMeSMARTData Metric Report

Metrics of the NVMeSMARTData telemetry report and the corresponding description of the metrics:

Table 3 NVMeSMARTData Report

NVMeSMARTData	AvailableSpare	Specifies the remaining spare capacity available as a normalized percentage (0 to 100%)
	AvailableSpareThreshold	The available spare value below which an asynchronous event completion may occur. The value is indicated as a normalized percentage (0 to 100%)
	CompositeTemperature	Indicates the current composite temperature (in Kelvin) of the controller and namespace(s) associated with that controller
	ControllerBusyTimeLower	Contains the lower part of the amount of time in minutes the controller is busy with I/O command
	ControllerBusyTimeUpper	Contains the upper part of the amount of time in minutes the controller is busy with I/O commands
	CriticalWarning	Indicates critical warnings about the controller state
	DataUnitsReadLower	Specifies the lower part of the count of 512 byte data units the host has read from the controller. This value is reported in thousands and is rounded up
	DataUnitsReadUpper	Specifies the upper part of the count of 512 byte data units the host has read from the controller. This value is reported in thousands and is rounded up
	DataUnitsWrittenLower	Specifies the lower part of the count of 512 byte data units the host has written to the controller. This value is reported in thousands and is rounded up
	DataUnitsWrittenUpper	Specifies the upper part of the count of 512 byte data units the host has written to the controller. This value is reported in thousands and is rounded up
	HostReadCommandsLower	Specifies the lower part of the number of read commands completed by the controller
	HostReadCommandsUpper	Specifies the upper part of the number of read commands completed by the controller
	HostWriteCommandsLower	Specifies the lower part of the number of write commands completed by the controller
	HostWriteCommandsUpper	Specifies the upper part of the number of write commands completed by the controller
	MediaDataIntegrityErrorsLower	Contains the lower part of the count of detected unrecovered data integrity errors. Includes errors such as uncorrectable ECC, CRC checksum failure, or LBA tag mismatch
	MediaDataIntegrityErrorsUpper	Contains the upper part of the count of detected unrecovered data integrity errors. Includes errors such as uncorrectable ECC, CRC checksum failure, or LBA tag mismatch
NumOfErrorInfoLogEntriesLower	Contains the lower part of the count of error information log entries over the life of the controller	

	NumOfErrorInfoLogEntriesUpper	Contains the upper part of the count of error information log entries over the life of the controller
	PercentageUsed	Specifies a vendor specific estimate of the percentage of NVM (Non Volatile Memory) subsystem life used based on the actual usage and the manufacturer's prediction of NVM life
	PowerCyclesLower	Specifies the lower part of the number of power cycles
	PowerCyclesUpper	Specifies the upper part of the number of power cycles
	PowerOnHoursLower	Specifies the lower part of the number of hours powered on. Hours may not include time that the controller was powered and in a non-operational power state
	PowerOnHoursUpper	Specifies the upper part of the number of hours powered on. Hours may not include time that the controller was powered and in a non-operational power state
	UnsafeShutdownsLower	Contains the lower part of the number of unsafe shutdowns. This count is incremented when a shutdown notification (CC.SHN) is not received prior to loss of power
	UnsafeShutdownsUpper	Contains the upper part of the number of unsafe shutdowns. This count is incremented when a shutdown notification (CC.SHN) is not received prior to loss of power

2.1.3 StorageSensor Metric Report

Metrics of the StorageSensor telemetry report and the corresponding description of the metrics:

StorageSensor	TemperatureReading	Temperature information for server storage internal devices
---------------	--------------------	---

2.2 Supported Devices for Storage Metric Reports

- StorageDiskSMARTData: SAS/SATA SSD drives behind PERC storage controller
- NVMeSMARTData: PCIeSSD/NVMe devices (drives and HHHL cards) not behind SWRAID controller
- StorageSensor: PCIeSSD/NVMe devices (drives and HHHL cards), SAS/SATA HDD/SSD drives behind PERC, BOSS M.2 drives and SATA HDD / SSD behind SWRAID controller

2.3 Special Cases of Hot plug Storage Devices

When the drives get hot swap, all the three Storage reports will reflect the change on the fly within the report Interval.

2.4 Workflow Examples

Configuration of the Storage reports could be done with either of interface racadm/redfish/SCP.

Example:


```
racadm>>set idrac.telemetry.EnableTelemetry Enabled
[Key=idrac.Embedded.1#Telemetry.1]
Object value modified successfully
```

```
racadm>>set idrac.TelemetryNVMeSMARTData.1.EnableTelemetry Enabled
[Key=idrac.Embedded.1#TelemetryNVMeSMARTData.1]
Object value modified successfully
```


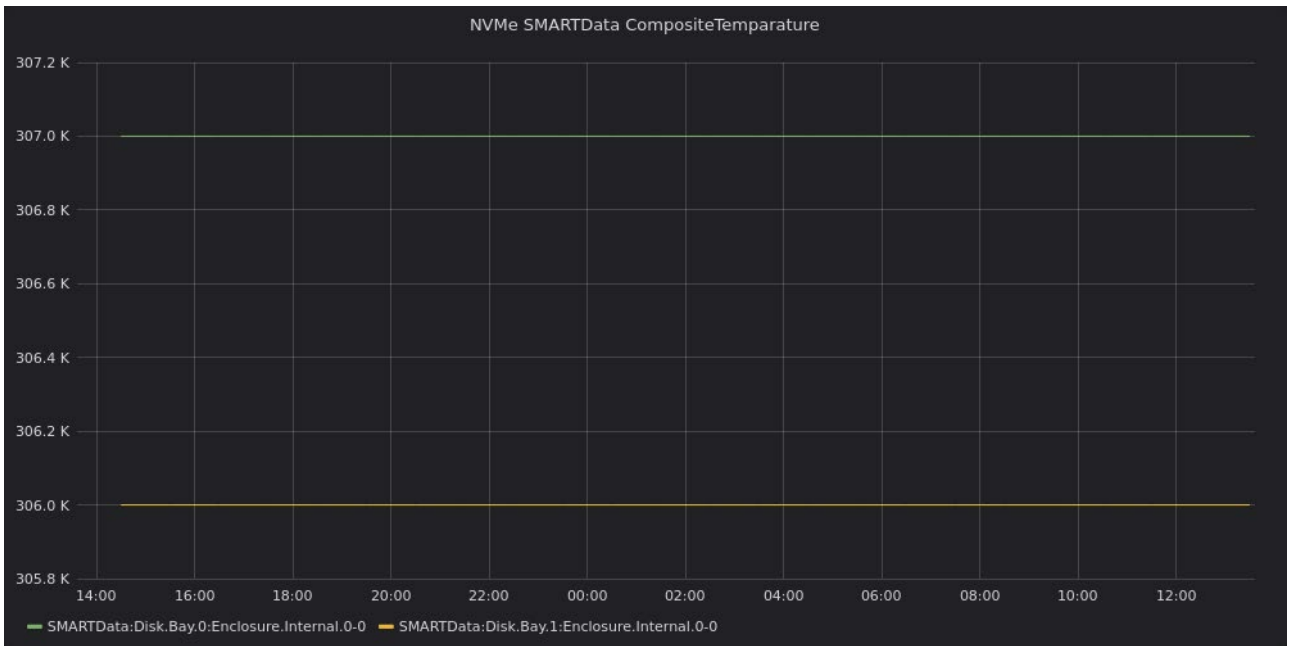
After the above configuration, NVMeSMARTData Report can be pulled, or streamed through http subscription, remote syslog or server-side events (SSE) based methods, Further details of the configuration could be found in the Telemetry Getting Started whitepaper.

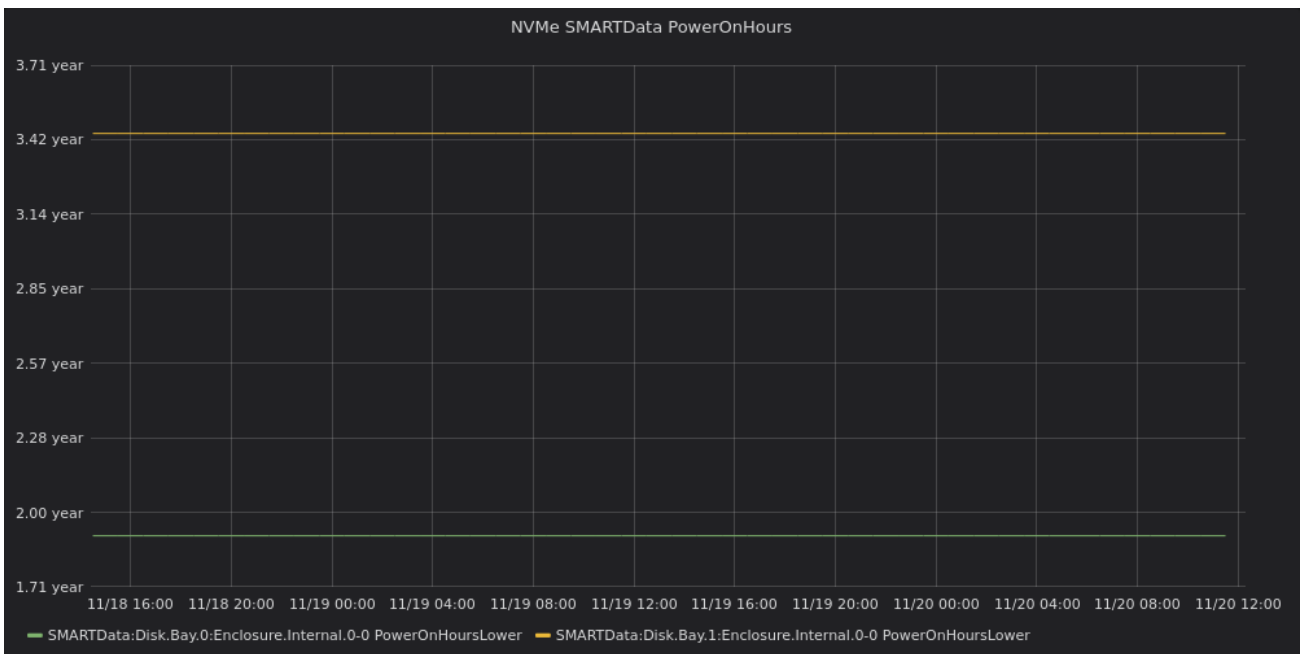
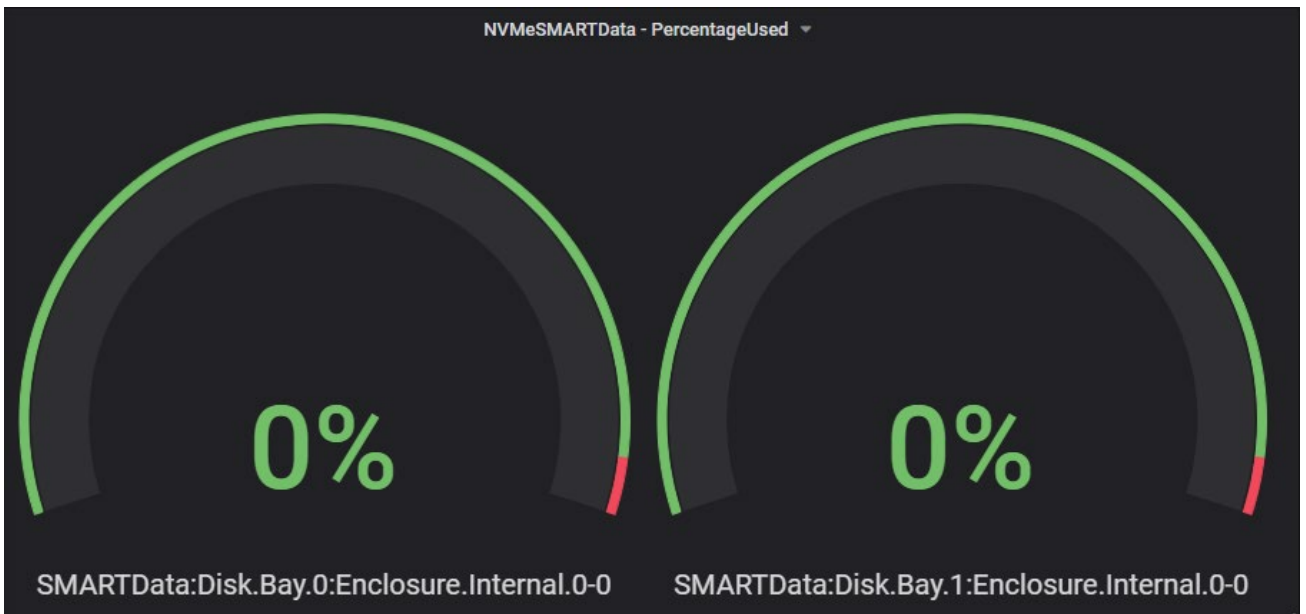
Similarly, other storage reports can be configured and streamed to assess the health and state of the drives on the server.

Below examples are from servers having multiple SSD's and Smart drives. The metrics of the Storage reports visualized below for example, could be used get the insight of the drive temperatures, drive states, drive lifetime, and so on.









A Technical support and resources

iDRAC Telemetry Workflow Examples

<https://github.com/dell/iDRAC-Telemetry-Scripting/>

Open-source iDRAC REST API with Redfish Python and PowerShell examples.

<https://github.com/dell/iDRAC-Redfish-Scripting>

The iDRAC support home page provides access to product documents, technical white papers, how-to videos, and more.

www.dell.com/support/idrac

iDRAC User Guides and other manuals

www.dell.com/idracmanuals

Dell Technical Support

www.Dell.com/support