Tolly. TEST REPORT

#220101

February 2020

Commissioned by
Dell Technologies

iDRAC9 Telemetry Streaming

Evaluation of The Performance and Efficiency of Telemetry Streaming in the New iDRAC9 v4.0 Release

Executive Summary

The datacenter of the future will rely on artificial intelligence (AI) to drive predictive & proactive management of the server infrastructure. In turn, AI systems rely on ingesting massive amounts of granular data. iDRAC9's Telemetry Streaming, available with the new iDRAC9 Datacenter license, is architected to provide highly granular, high-value data using resource-efficient data streaming.

Dell Technologies commissioned Tolly to benchmark the number of data points collected from a representative Dell EMC PowerEdge server via iDRAC9 v4.00.00.00 in a 24-hour period and compare that to a prior, less data-rich monitoring implementation. Additionally, data collection via polling and streaming were compared for network efficiency.

Tolly found that iDRAC9 Telemetry Streaming running on a representative test system produced over 2.9M data points in a 24-hour period compared with 6,240 data points for a typical polled server using the existing Redfish interfaces and a reasonable real-world polling interval of 30 minutes. With Telemetry Streaming each 24-hour report can be collected via a ***single HTTP request*** compared with a per-report best case of 17,280 HTTP requests to collect a single report daily. See Figure 1 & Table 1.

The Bottom Line

iDRAC9 v4.0 Telemetry Streaming provides:

1. High-value data including GPU monitoring, Advanced CPU metrics, Storage Drive SMART Logs and Advanced Monitoring
2. Over 2.9 million data points streamed to the analytics system in a 24-hour period per server
3. Data streaming architecture that is massively more efficient than device polling
4. Time series data that assures no missed data points which can occur in large polled environments
5. Agent-free, side-band management of server components

Performance of Telemetry Streaming vs Polling Over 24 Hours

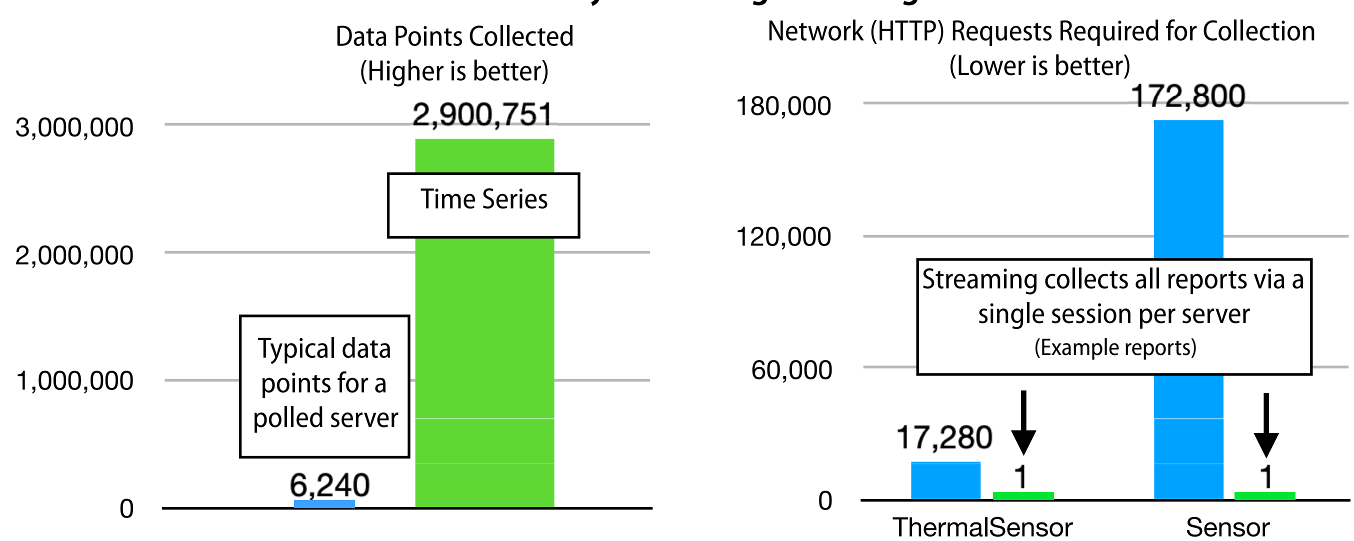

Notes: Test data gathered from Dell EMC PowerEdge R740xd with GPU. ThermalSensor report is a best case for polling as all data is from a single source. Single HTTP request needed for data collection telemetry streaming is too low to be visible in bar chart. All data is for a single server.

Source: Tolly, December 2019 | Polling | Telemetry Streaming | Figure 1

Dell iDRAC9 v4.0

This new version of iDRAC9 represents a significant step forward that will aid customers in advancing the strategic automation of their datacenters.

Most notably, iDRAC9 v4.0 adds the industry-standard DMTF Redfish[1] v1 telemetry service. Available only with the Datacenter license, this update provides not only additional, high-value data reports but also combines with new data streaming technology to provide significantly more raw data with dramatically improved network efficiency. Furthermore, iDRAC9 v4.0 leverages an agent-free architecture to gather data without any operating system involvement.

High-Value Data

iDRAC9 data is delivered in industry-standard JSON format which is both machine- and human-readable. Figure 2 illustrates the data flow from iDRAC9 into ingress collectors, analytics and visualization systems.

The data is well-suited for ingestion into any number of multi-vendor analytics systems.

The iDRAC9 telemetry not only provides more data, it provides data of particular interest to data center managers:

- Thermal - Direct temperature reporting from components which is more far more accurate than using arrays of sensors. Time series data that can be used at a data center level.
- Power - System voltage information and power (W) consumed.
- Compute Usage Per Second (CUPS) — Provided via the Intel management engine (ME), this is a real-time performance index of CPU, memory and I/O information. iDRAC9 provides averages every five seconds which is more frequent than some competitors.
- NIC - Detailed statistics for network interface cards. Each NIC resident in the system provides separate statistics. Over 40 metrics are provided for each instance. These cover data transmitted/received as well as error counts, etc.
- Storage Disk SMART Logs - This report provides 19 different data points including power cycle, error counts and other vital data storage-related metrics.

Network Efficiency of Streaming

Previously, typically most data was collected via polling. That is, a monitoring server would repeatedly poll each server it is collecting metrics on at the desired interval to collect the most recent "bucket" of information.

This approach presented challenges for large environments. As the number of servers increased, the workload and time required to poll all servers also increased. Furthermore, should the polling interval exceed the update interval for a given data point, old data would be overwritten with new data. The result being that old data would be lost and never delivered into the analytics system.

Additionally, the time required (i.e. latency) for the outbound polling can become significant over the course of a day.

iDRAC9 v4.0 with the Datacenter license does away with the need for polling and replaces it with telemetry streaming which is incomparably more efficient than polling.

iDRAC Telemetry Streaming Test Details
Dell EMC PowerEdge R740xd with GPU

Report	Data Points Per Report
StorageDiskSMARTData*	5,472
NICStatistics*	368,940
GPUMetrics*	566,280
CPUMemMetrics*	289,872
PowerStatistics	28,720
StorageSensor	386,496
NICSensor	32,208
MemorySensor	394,814
ThermalMetrics	184,305
CPUSensor	34,296
FanSensor	3,320
ThermalSensor	32,208
CUPS	68,880
PSUMetrics	5,716
Sensor	354,288
PowerMetrics	144,936
Total	2,900,751

Notes: Detailed breakdown by report name of data presented in Figure 1. * Reports only available with Datacenter license.

Source: Tolly, December 2019 Table 1

[1] DMTF's Redfish® is a standard designed to deliver simple and secure management for converged, hybrid IT and the Software Defined Data Center (SDDC). See dmtf.org for more information.

With streaming, the collecting server issues a single HTTP GET request which remains active and automatically receives all data for a single statistics report. Figure 1 provides two examples that illustrate the network efficiency benefits of telemetry streaming.

The first example, ThermalSensor, provides a best-case situation where all of the data be retrieved from a single report. Over the course of a single day, the polling approach would require 17, 280 HTTP GET requests to be issued. (One HTTP GET every 5 seconds, resulting in 720 per hour and 17,280 per day.) Each request measured in our testing required .303 seconds to complete resulting in 86.094 minutes of collection time for that single set of data on a single server. With telemetry streaming all that network overhead is eliminated.

The second example, Thermal, shows what happens when multiple URIs need to be queried for data. As the number of URIs to be polled increases, the network overhead increases - in this case to 172,800 calls over 24 hours. With telemetry streaming, again, all this data is collected via a single HTTP GET that remains open to the server.

Consistent Data Format

Telemetry Streaming delivers content in a consistent format using a standard schema format. All reports are generated using JavaScript Object Notation (JSON). Thus, it is a simple task for systems, or people, to use and process the data. See Figure 4 for a partial example of a report generated by Telemetry Streaming.

Previously, different data reports used different schema formats. This necessitated custom scripting to parse each report according to its particular schema.

Dell Technologies

Telemetry Streaming - iDRAC9 v4.0

Performance and Feature Evaluation

Tolly. Certified

Tested December 2019

Test Configuration

To illustrate the differences in data provided with telemetry streaming, engineers built out a Dell EMC PowerEdge R740xd with GPU capabilities. The particular configuration was chosen to show the richness of data available. Each particular system, naturally, will have its own specific set of components

Integrating iDRAC9 Telemetry into Analytics Solutions

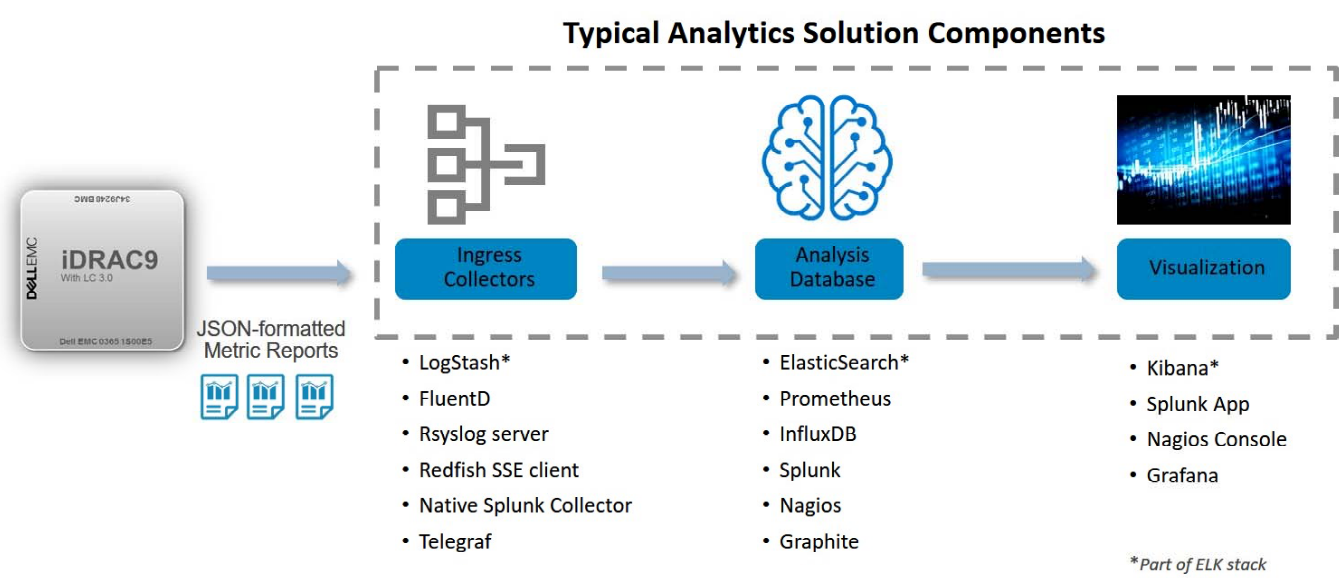

Source: Dell Technologies, December 2019

Figure 2

and, thus, specific telemetry. Key elements of the system used for evaluation data can be found in Figure 3.

Typical Data Points of Polled Server

When using polling, your analytics gathering is literally racing against the clock. Given the time overhead of polling, there is a finite amount of data that can be pulled from even a single server during each 24-hour period.

Engineers calculated the data points that could be collected from a single polled server by first counting the unique data points available for the reports available from a Redfish-based polled server as configured for this test which was 130 distinct data points per polling cycle. See Table 2.

Benefits of iDRAC's Agent-free Architecture

In the past, access to information such as network interface card statistics required software agents in the operating system to serve as the information channel. This required OS-specific agents and could place overhead on the operating system when providing data collection services.

Being able to retrieve data directly from devices without the need for any operating system agent both simplifies the communication between iDRAC9 and server components and removes the data collection burden from the operating system.

Dell Technologies has leveraged industry standard protocols such as MCTP, NC-SI and NVMe-MI to communicate to each peripheral device such as PERC RAID controllers, Ethernet NICs, Fibre Channel HBAs, SAS HBAs, and NVMe drives. This architecture is the result of lengthy, multi-year partnerships with industry-leading vendors to provide agent-free device management in our PowerEdge servers. This architecture is the key enabler for how comprehensive telemetry data is captured by iDRAC9.

Source: Dell Technologies, December 2019

Test System Configuration (Partial)
Dell EMC PowerEdge R740xd with GPU

Inventory	
CPUs 1 & 2	Xeon Gold 6130 (16 cores each)
DIMMs A1-B6	Hynix HMA84GR7AFR4N-VK 32 GB Dual-Rank DDR4 2666 MHz
DIMMs A7-B12	Micron 18ASF2G72XF12G6V21AB 16 GB Single-Rank DDR4 2666 MHz
Power Supplies 1 & 2	Delta 2000 W AC Supply
Video Controllers in Slots 1, 4, 8	TU104GL [Tesla T4]
Integrated NIC 1, Port 1	Intel Corporation Ethernet 10G 4P X710 SFP+ rNDC
Integrated NIC 1, Ports 2-4	Intel Corporation Ethernet 10G X710 rNDC
NIC in Slot 6	Intel Corporation Ethernet 10G 2P X550-t Adapter
RAID Controller in Slot 5	PERC H730P Adapter
Disks 0-5 in Backplane 1 on RAID Controller in Slot 5	INTEL 119 GB SSDSC2BB120G7R
Disks 12-17 in Backplane 1 on RAID Controller in Slot 5	SAMSUNG 399 GB MZILS400HEGR0D3
Disks 18-23 in Backplane 1 on RAID Controller in Slot 5	SEAGATE 999 GB ST1000NX0443
Disks 6-11 in Backplane 1 on RAID Controller in Slot 5	TOSHIBA 299 GB AL13SXB30ENY

Source: Tolly, December 2019

Figure 3

Assuming a polling cycle of 30 minutes, this would result in 48 collections of data over a 24-hour period providing a total of 6,240 data points.

Additionally, taking a single telemetry report as an example, and the best case scenario for the older data collection methods, Telemetry Streaming delivered 32,208 observed data points in the ThermalSensor report over the 24 hour period using a single, long-lived HTTP connection. Attempting to replicate this exact level of data collection using the older Redfish polling methods would require 12 HTTP requests every minute, totaling 720 per hour, and 17,280 requests over 24 hours. In a real-world datacenter environment, this level of traffic is not scalable across a typical deployment of thousands of servers. Thus, most real-world deployments monitor at a much more reduced rate, with typical systems being polled at 30 minute or hourly intervals.

Telemetry Streaming: CPUMemMetrics Report (Partial) Standard Schema, JSON Format


```
{
  "@odata.type": "#MetricReport.v1_2_0.MetricReport",
  "@odata.context": "/redfish/v1/$metadata#MetricReport.MetricReport",
  "@odata.id": "/redfish/v1/TelemetryService/MetricReports/
CPUMemMetrics",
  "Id": "CPUMemMetrics",
  "Name": "CPU Mem Metrics Metric Report Definition",
  "ReportSequence": "1340",
  "MetricReportDefinition": {
    "@odata.id": "/redfish/v1/TelemetryService/
MetricReportDefinitions/CPUMemMetrics"
  },
  "Timestamp": "2019-12-10T11:15:15-06:00",
  "MetricValues": [
    {
      "MetricId": "CPUC0ResidencyHigh",
      "Timestamp": "2019-12-10T11:15:11-06:00",
      "MetricValue": "1",
      "Oem": {
        "Dell": {
          "ContextID": "CPU.Socket.1",
          "Label": "CPU.Socket.1 CPUC0ResidencyHigh"
        }
      }
    },
    {
      "MetricId": "CPUC0ResidencyLow",
      "Timestamp": "2019-12-10T11:15:11-06:00",
      "MetricValue": "831805080",
      "Oem": {
        "Dell": {
          "ContextID": "CPU.Socket.1",
          "Label": "CPU.Socket.1 CPUC0ResidencyLow"
        }
      }
    },

    {
      "MetricId": "NonC0ResidencyHigh",
      "Timestamp": "2019-12-10T11:15:11-06:00",
      "MetricValue": "7966",
      "Oem": {
        "Dell": {
          "ContextID": "CPU.Socket.1",
          "Label": "CPU.Socket.1 NonC0ResidencyHigh"
        }
      }
    },
```


Note: This report continues on and provides over 200 metrics across a 10 minute reporting period.

Source: Tolly, December 2019

Figure 4

Polled Server Unique Data Points Per Cycle

Report	Data Points Per Report
PowerStatistics	20
StorageSensor	24
NICSensor	2
MemorySensor	24
ThermalMetrics	11
CPUSensor	2
FanSensor	6
ThermalSensor	2
CUPS	4
PSUMetrics	4
Sensor	22
PowerMetrics	9
Total	130

Source: Tolly, December 2019

Table 2

About Tolly

The Tolly Group companies have been delivering world-class information technology services for over 30 years. Tolly is a leading global provider of third-party validation services for vendors of information technology products, components and services.

You can reach the company by E-mail at sales@tolly.com, or by telephone at +1 561.391.5610.

Visit Tolly on the Internet at:
http://www.tolly.com

Terms of Usage

220101 ef-O wt-2020-02-12—VerO