

Transform Datacenter Analytics with iDRAC9 Telemetry Streaming

Tech Note by

Kim Kinahan
Michael E. Brown
Rick Hall
Doug Iler

Summary

Telemetry Streaming, a new feature in iDRAC9 v4.0 enabled by the new Datacenter License, can produce more high-value (comprehensive and accurate) data faster than with previous versions.

There is a huge amount of untapped machine data in your IT infrastructure: use iDRAC9 Telemetry Streaming and analytics to leverage that data to optimize your server management and operations.

Introduction

With the advent of the new iDRAC9 v4.00.00.00 firmware release and the Datacenter license, IT managers can now integrate advanced telemetry about the server hardware operation into their existing analytics solutions. This telemetry is provided as granular, time-series data that can be streamed versus using inefficient, legacy polling methods. The advanced agent-free architecture in iDRAC9 provides over 180 data metrics (with more coming) related to server and peripherals operations that are precisely time-stamped and internally buffered to allow highly efficient data stream collection and processing with minimal network loading. This comprehensive telemetry can be fed to popular analytics tools to predict failure events, optimize server operation, and enhance cyber-resiliency.

Telemetry and Analytics

Telemetry has been around for decades and has been used in various business applications, from hospitals monitoring patients to oil and gas drilling systems to weather balloons transmitting meteorological data. The definition of *Telemetry* is an “automated communications process by which measurements are made, and other data collected at remote or inaccessible points are transmitted to receiving equipment for monitoring.”

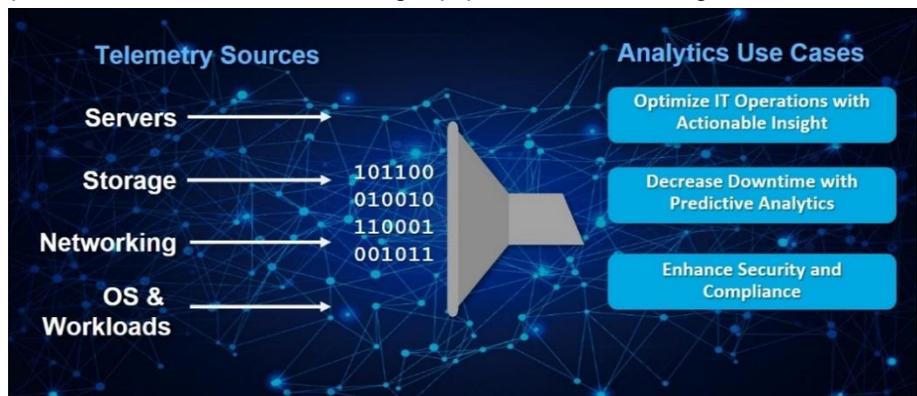


Figure 1. Telemetry Monitoring in a Typical Data Center

In the era of “Big Data,” IT managers leverage a wide range of telemetry from their infrastructure in their monitoring tools, as shown in Figure 1. However, increasingly that telemetry is also used in AI-based analytics to gain operational insight into their datacenter operations. This is far more powerful than using simple alerting and monitoring techniques that typically only report health and status via SNMP alerts or IMPI traps.

Using analytics tools, IT managers can more proactively manage by analyzing trends and discovering insightful relationships between seemingly unrelated events and operations. A recent survey found that 61% of IT decision-makers considered data and analytics very important to their business growth strategy/digital transformation efforts.¹

Some of the use cases for data center analytics are:

- **Predictive analytics:** Customers can perform an in-depth analysis of server telemetry, including device parametric data to proactively replace failing devices. In one case, an IT team used analytics on telemetry from memory devices to develop an algorithm that predicted eventual failure. This allows proactive replacement of suspect devices during scheduled maintenance windows, significantly improving uptime and SLA quality.
- **Optimized IT operations:** You can perform time-series analysis of vital server metrics to gain insights into optimizing server operation, including tracking of power, temperature, CPU, and I/O performance, etc. One industry that makes extensive use of analytics is High-Frequency Trading, where every millisecond of compute counts in accelerating automated trades. Detailed telemetry is commonly used to discover ways to squeeze out more performance from servers, which becomes a key competitive advantage in this industry.
- **Security:** AI-based analytics can respond far faster to security events. You can enhance security AI and forensics by monitoring the say of unusual user login activity or physical intrusion events on your servers.

However, to perform effective analytics, you need data: lots and lots of it to feed Machine or Deep Learning techniques effectively. The larger the data set, the more accurate the analysis becomes as evidenced by the petabytes of data that social media uses in analytics of user attributes and buying behaviors

The Streaming Advantage in iDRAC9

Telemetry streaming's big performance advantage is in reducing the overhead needed to get the complete data stream from a remote device. Retrieving telemetry using polling can result in an enormous number of discrete commands being issued, which is very challenging in scaling across a large datacenter. With iDRAC9 Telemetry Streaming, you get time-series and detailed statistics reports delivered directly to a variety of analytics collection tools with higher efficiency by removing the need for issuing individual commands for each piece of data. The streaming configuration is flexible so users can modify the number of metrics they require, the report interval (30 seconds for example), and enable reports to be sent immediately upon detection of critical events in the server (like a PSU failure say).

In summary, the advantages of Streaming over Polling are:

- **Better Scalability:** Polling requires a lot of scripting work and CPU cycles to aggregate data and suffers from scaling issues when we are talking about 100's or 1000's of servers. Streaming data, in contrast, can be pushed directly into popular analytics tools like Prometheus, ELK stack, InfluxDB, Splunk without the overhead and network loading associated polling.

- **More Accuracy:** Polling can also lead to data loss or “gaps” in sampling for time series analysis; it is usually only a snapshot of current states, not the complete picture over time. You might miss critical peaks or excursions in data.
- **Less Delay:** Data can be severely delayed in time due to needing multiple commands to get a complete set of data and the inability to poll simultaneously from a central management host. Streaming more accurately preserves the time-series context of data samples.

Consequently, streaming is a far more efficient and accurate way to gather telemetry.

Telemetry Excellence with the iDRAC9 Datacenter License

iDRAC9 v4.0, with the Datacenter license, offers over 180 telemetry metrics on various server devices and sensors. These metrics also form the basis of our SupportAssist Collection Report, an incredibly useful tool that captures over 5,000 pieces of diagnostic data and log files for troubleshooting server issues. iDRAC9 Telemetry Streaming does all the heavy lifting for you by internally sampling and storing all the data points and then streaming them out in reports at a frequency that fits your needs. iDRAC9 can deliver almost 3 million metrics a day to transform the accuracy of analytics processing for your data center!

Telemetry can be delivered via the following methods:

- Redfish Server-Sent Events (SSE), a DMTF standard for streaming data²
- Redfish subscription for pushing events, another DMTF standard
- Remote Syslog, a protocol for pushing logs for centralized monitoring
- Non-streaming, scripted polling via the iDRAC9 RESTful API (though not as efficient as streaming as discussed earlier)

The data is formatted using JSON (JavaScript Object Notation) and can be easily adapted to connect many analytics solutions on the market, as shown in Figure 2.

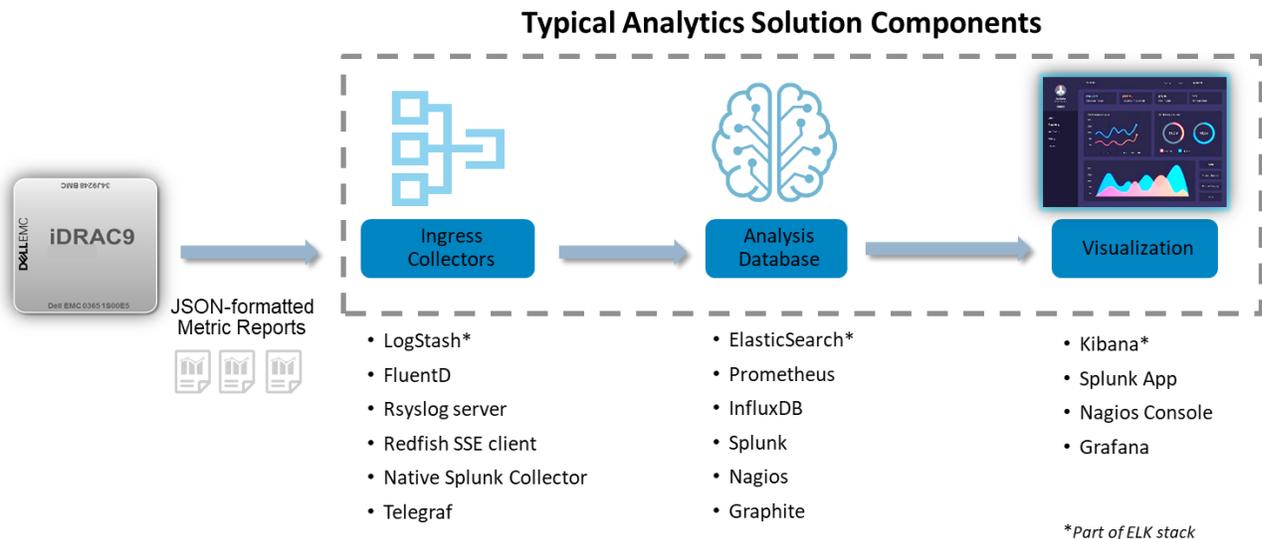


Figure 2. Integrating iDRAC9 Telemetry Streaming with Popular Analytics Solutions

Types of Telemetry Data

A summary of the types of telemetry that iDRAC9 has are:

New Telemetry Data with iDRAC9 4.0:

- Serial Data Log messages
- GPU Accelerator Inventory & Monitoring
- Advanced CPU Metrics
- Storage Drive SMART logs
- Advanced Memory Monitoring
- SFP+ Optical Transceiver Inventory & Monitoring

Existing Telemetry Data:

- Configuration – comprehensive settings for all devices (BIOS, iDRAC, NICs, RAID, etc.)
- Inventory: comprehensive server hardware and firmware reporting
- Performance: CPU, memory bandwidth and I/O usage (Compute Usage Per Second or CUPS)
- Performance and diagnostic statistics: PERC, NICs, Fiber Channel
- Sensors: voltage, temperature, power, connectivity status, intrusion detection
- Logs: SEL log, iDRAC diagnostics, Lifecycle Controller Log

Figure 3 illustrates an external analytics solution capturing and visualizing iDRAC9 Telemetry Streaming. In this case, CUPS performance data was streamed to InfluxDB for the data analysis, and Grafana then used for the visualization.

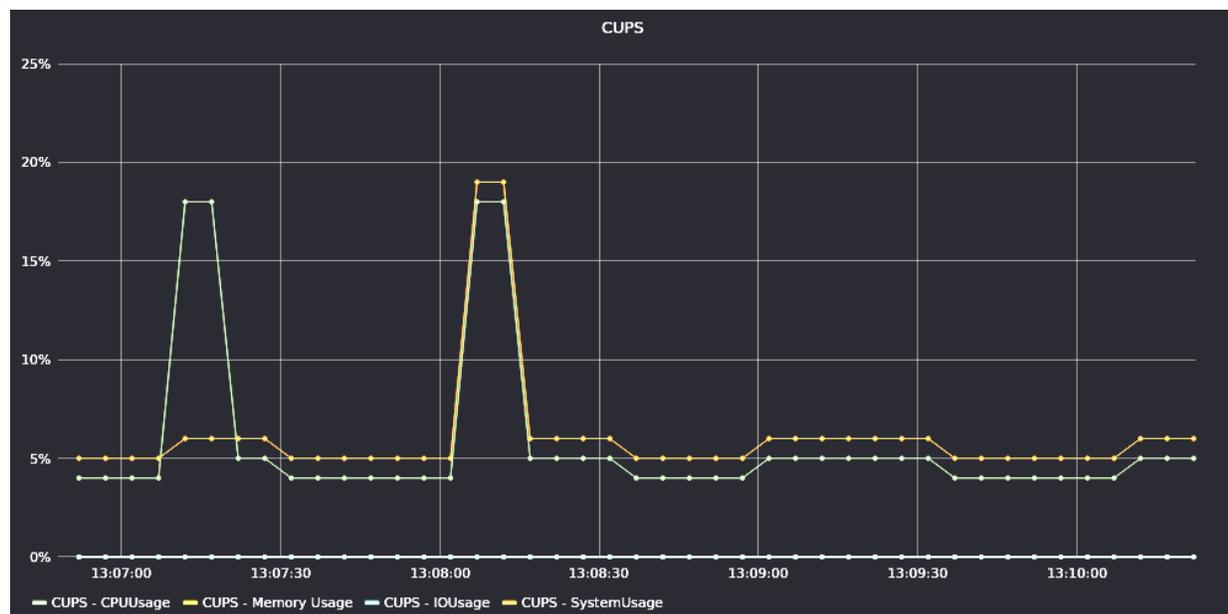


Figure 3. Example of iDRAC9 Telemetry for CUPS Performance Data

In Conclusion

Dell EMC continues to introduce innovations that help our customers automate the management of their IT infrastructure. iDRAC9 Telemetry Streaming represents a huge step forward in helping our customers leverage the extensive data available in their PowerEdge servers. Customers can easily stream this telemetry into their analytics tools and leverage advanced AI techniques to automate their IT systems management and operations further.

1. “2020 Global State of Enterprise Analytics”, published by MicroStrategy.
2. Server-Sent Events (SSE) is a server push technology (part of HTML5) enabling a client to receive automatic updates from a server via an HTTP/S internet connection.



PowerEdge DfD Repository
For more technical learning



Contact Us
For feedback and requests



Follow Us
For PowerEdge news