

Getting Started with Telemetry Monitoring

Understand your infrastructure's performance in detail
using Intel telemetry data.

How-to Guide

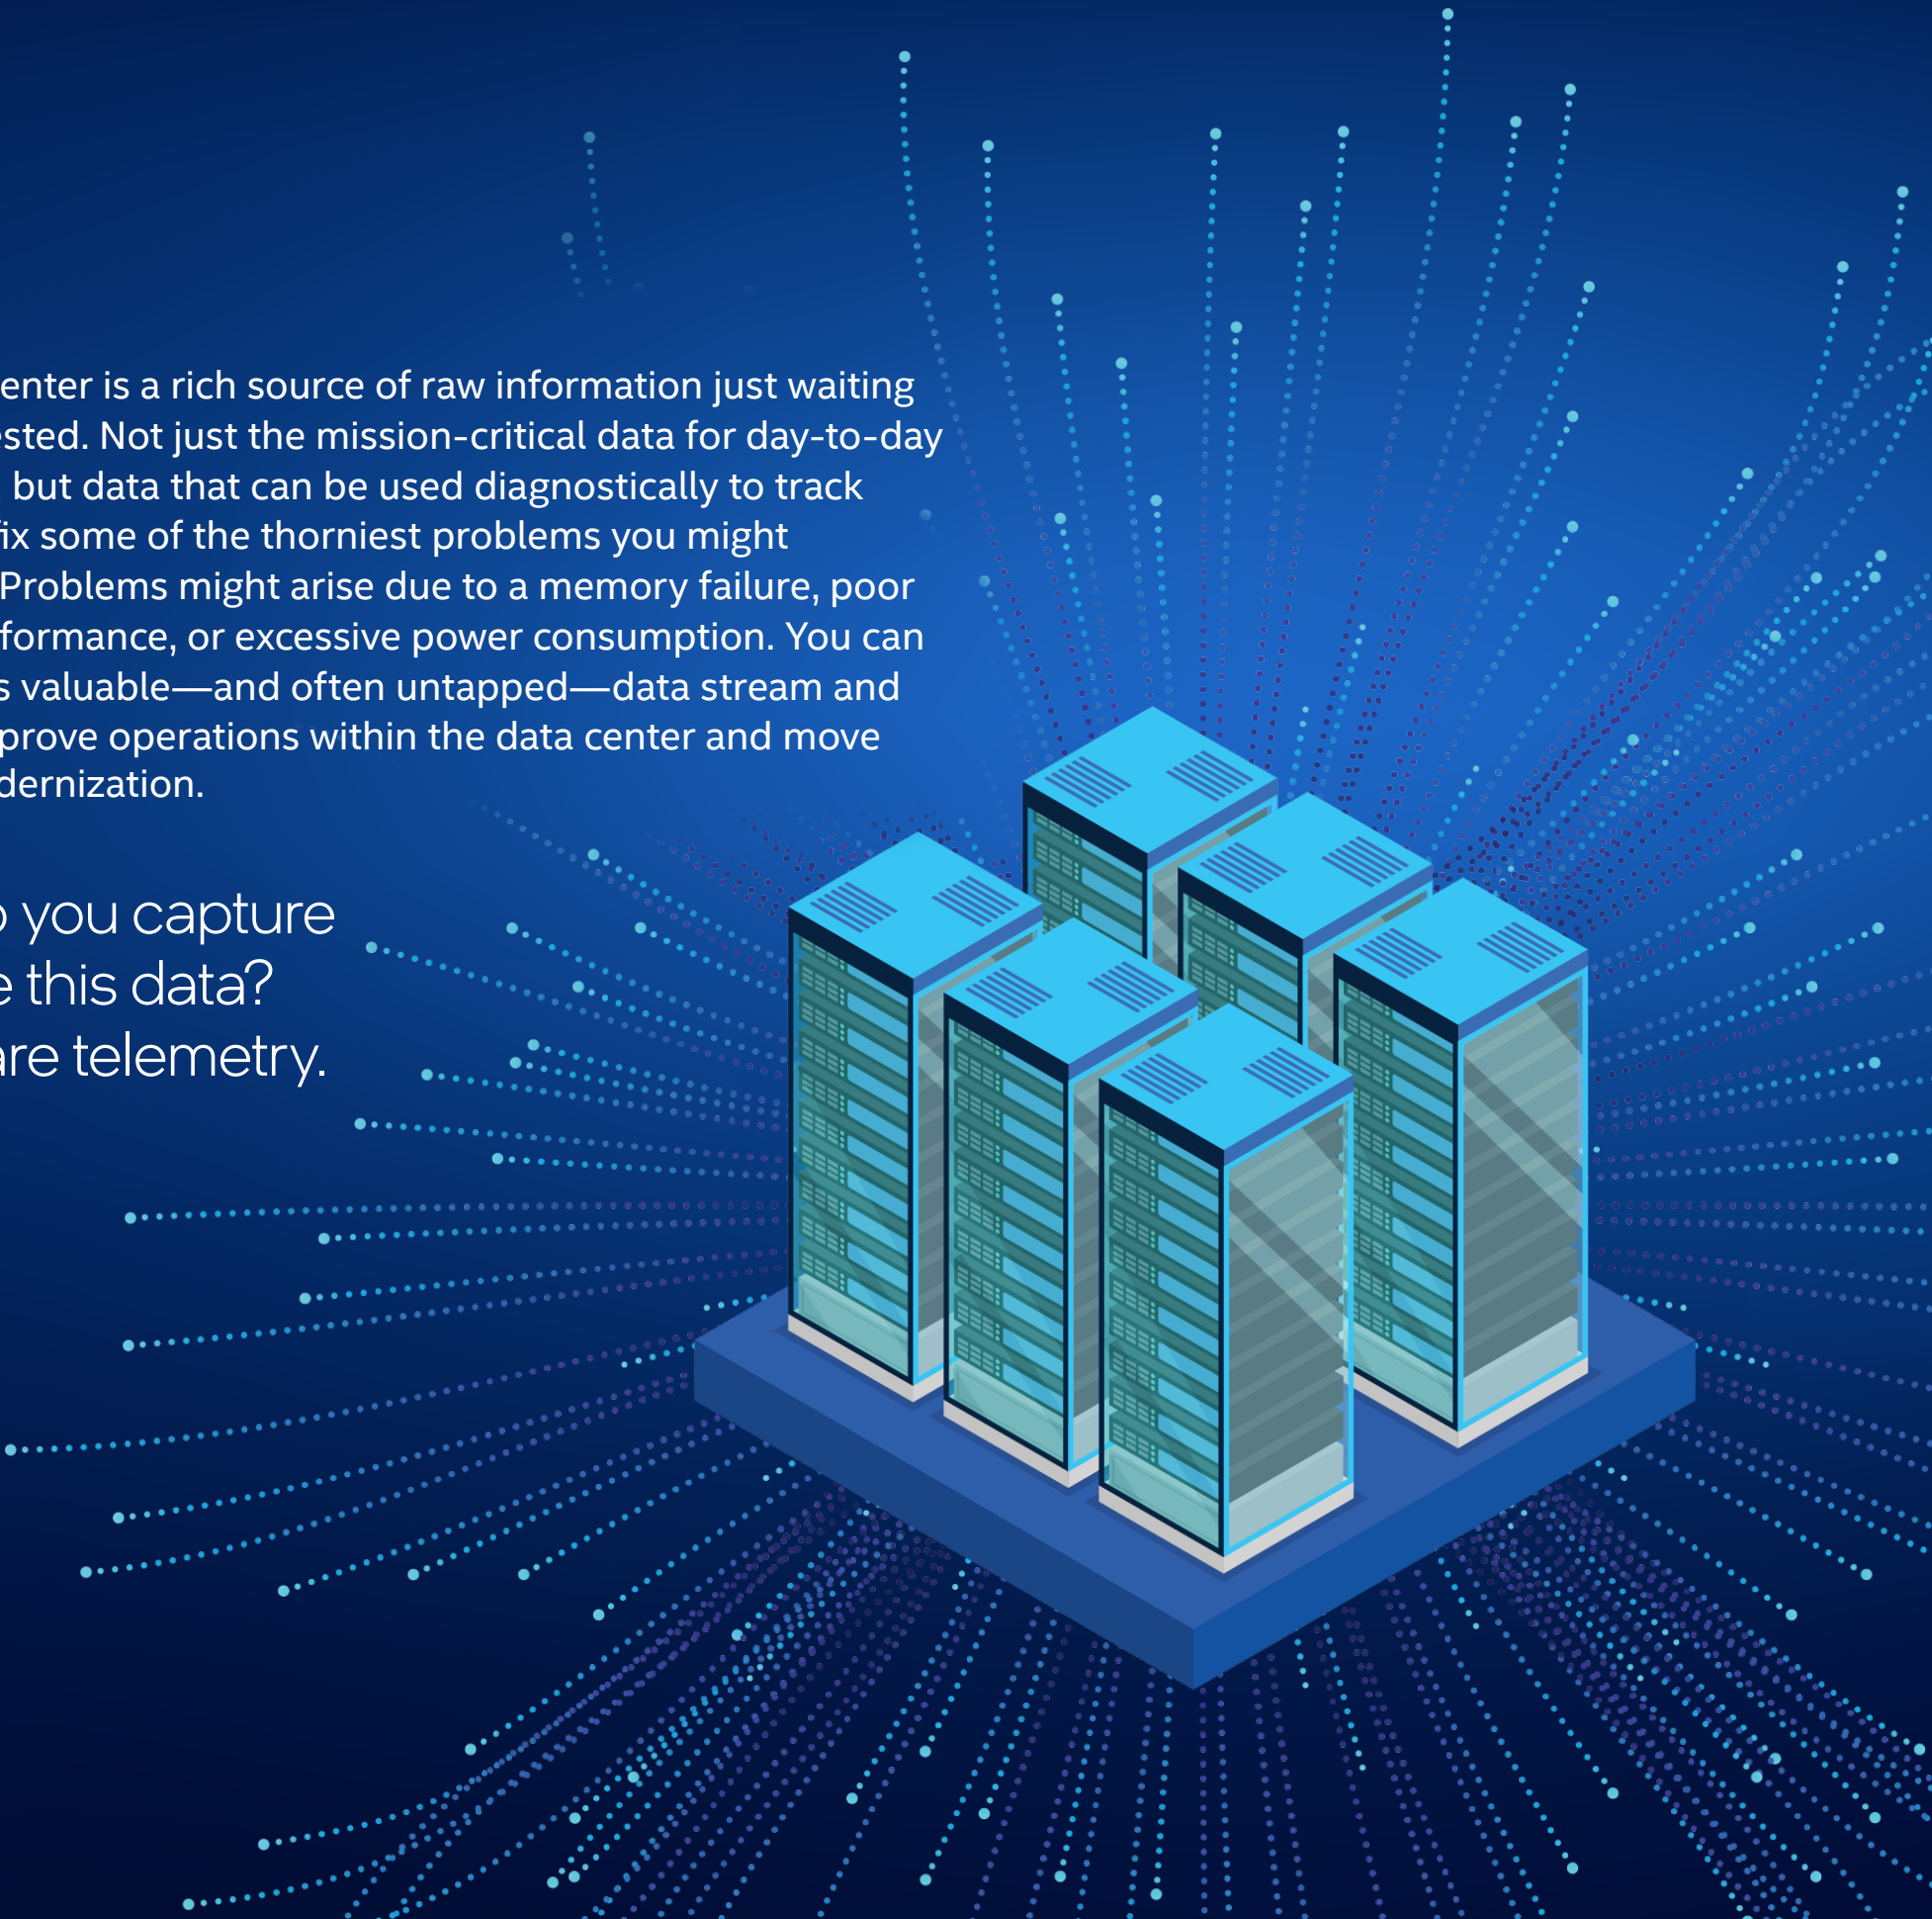


Contents:

What is hardware telemetry?	2
Why use telemetry?	3
What type of data should I collect with telemetry?	4
What's the best way to build a telemetry software stack?	5
What are other benefits of telemetry monitoring?	5
Where do I go for more information?	6

Your data center is a rich source of raw information just waiting to be harvested. Not just the mission-critical data for day-to-day operations, but data that can be used diagnostically to track down and fix some of the thorniest problems you might encounter. Problems might arise due to a memory failure, poor system performance, or excessive power consumption. You can capture this valuable—and often untapped—data stream and use it to improve operations within the data center and move toward modernization.

How do you capture and use this data?
Hardware telemetry.



What is hardware telemetry?

Generally speaking, telemetry refers to a stream of data coming from any number of sources, including servers, software, and environmental equipment. For hardware telemetry, this data is generated by the components on your hardware platform, such as the CPU, memory, and PCIe interface.

Registers on Intel® Xeon® Scalable processors monitor cache, CPU frequencies, memory bandwidth, and input/output (I/O) accesses. Hardware telemetry is generated from a robust set of model-specific registers (MSRs) and sensors on the Intel® platform.



Figure 1. Hardware telemetry data generated by Intel architecture reveals conditions and patterns that can help you understand and orchestrate what's going on in each server, CPU, and core



Why use telemetry?

Hardware telemetry can provide a detailed picture of exactly what is going on in all your servers and their components. This can help you find the root cause of a performance problem or degradation: whether the problem stems from the hardware (as in memory) or a workload imbalance on servers or components. These insights can speed and simplify problem resolution, freeing your time up for more interesting and more proactive work.

Telemetry monitoring can make data center management easier and more efficient. Because it shows you where the bottlenecks and overloads are occurring over time, you can fine-tune processes and orchestrate workload placement based on real-time workload profiling. This can help you protect your network performance and balance the load across your infrastructure.

Telemetry can be used to check configurations and gain insights into resource utilization, power efficiency, and general system health. In a well-orchestrated data center, it can help you identify anomalies that might be early indicators of possible points of failure. This can shift IT efforts from constantly reacting to problems to identifying—or even predicting—the first signs of a problem.

Greater efficiency and fewer overloads can reduce the frequency of failures and improve key performance indicators (KPIs), including total cost of ownership (TCO), reliability, security, performance, and power consumption.

What type of data should I collect with telemetry?

Keep in mind that when it comes to telemetry data, less can be more. The objective is not to collect the maximum amount of data, but to collect the right data in a useful amount. Collecting everything every few seconds can lead to enormous log files and, eventually, poor system performance.

A best practice is to start with just your highest priority data and gradually add more streams. In addition, it's a good idea to scan data less frequently until an issue is identified, and then to increase the frequency until the issue is resolved. Table 1 provides guidance on the types of data to collect and how often to do so.

Table 1. Guidelines for data-collection frequency

What Data to Collect	Why to Collect It	How Often to Collect It
System configuration, kernel version, and log messages	Troubleshooting and failure-trend detection	On reboot
Power (CPU, memory, and system) and thermal (inlet, outlet, DIMM, and CPU airflow)	Manage power and thermal efficiency; identify spikes as early indicators of failure	30 seconds to 1 minute
Performance of the CPU, cache, and memory	Target machines for workload colocation	1–10 seconds
Memory errors	Identify issues causing device degradation and early indicators of failure	1–5 minutes Trigger more frequent scans on first uncorrectable memory error (UME)
DIMM performance and health	Identify issues causing device degradation and early indicators of failure	15–30 minutes and on startup
Throughput	Failure-trend detection	30 seconds to 1 minute

Intel Telemetry Collector (ITC)

ITC is a reference collector that provides a quick introduction to the different metrics available, including power and thermal statistics, performance counters, process activities, threads, and operating-system-level disk, network, and memory statistics. With ITC, you can ingest and visualize data from various sources and multiple machines.¹

The ITC visualization highlights typical pain points such as memory bandwidth, NUMA imbalance, and interrupt request (IRQ) affinity issues. However, while ITC is a good first step into telemetry, it's not a fleet solution that can scale for hundreds or thousands of servers. For that level of telemetry, you'll want to advance to the next step—a scalable open source telemetry software stack—see the next section: [What's the best way to build a telemetry software stack?](#)

What's the best way to build a telemetry software stack?

You can use different stacks, including Elasticsearch, Logstash, and Kibana (ELK), for telemetry, but the Prometheus stack is the easiest and best supported. Intel supports the Prometheus tool chain by adding exporters for Intel hardware components, either through collectd or Telegraf exporters or through direct exporters. The stack includes data collection using exporters, integration with Kafka, a time-series Prometheus database, and a visualization tier for Grafana.

This Prometheus stack supports the open telemetry standard, is containerized, and is not difficult to set up. It can be supplemented by additional Intel hardware collectors to expose data from components such as Intel® Optane™ technology and the Data Plane Development Kit (DPDK) for networking data.



Figure 2. The open source Grafana platform, supported by Intel, visualizes and analyzes data on a dashboard for proactive monitoring and management of system trouble spots

What are other benefits of telemetry monitoring?

Telemetry monitoring can provide faster troubleshooting with high visibility into your data center's components and operations. It also provides many other benefits that extend into the overall modernization of your data center infrastructure:

- **Operational efficiency (operating expenses [OpEx]).** Increase density through orchestration based on telemetry to run more workloads on fewer servers.
- **Infrastructure efficiency (capital expenditures [CapEx]).** Identify the right-sized SKU for a workload in a fraction of the time with SKU emulation.
- **Higher utilization.** Improve your power-usage effectiveness (PUE). The average PUE for a data center has not improved since 2013.²
- **Better performance.** Optimize configuration code—the source of most performance problems—for an improved user experience.
- **Increased reliability.** Avoid those 3 a.m. crisis calls and keep your customers delighted. Thirty-eight percent of network outages take more than a day to find and resolve, and improving network resilience is the number-1 priority for 49 percent of global IT leaders.³

Where do I go for more information?

Learn more about telemetry monitoring by reading the paper, [“Building Intelligence into Your Cloud and Data Center.”](#)



¹ Talk to your Intel representative about how to get access to ITC.

² UptimeInstitute. “Data center PUEs flat since 2013.” April 2020.
<https://journal.uptimeinstitute.com/data-center-pues-flat-since-2013/>.

³ ComputerWeekly. “Multiple network outages costing enterprise businesses millions of pounds a year.” May 2020.
computerweekly.com/news/252483308/Multiple-network-outages-costing-enterprise-businesses-millions-of-pounds-a-year.

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Printed in USA

0421/SMR/PRW/PDF

Please Recycle 346868-001US