09-02-2026

Performance data

Performance data refers to the measurable metrics, statistics, and logs that provide insights into how the underlying hardware, software, or network resources supporting a software system are functioning. This data helps organizations monitor, analyze, and optimize the performance, reliability, scalability, and availability of the infrastructure that supports their applications and services.

This data enables teams to identify bottlenecks, diagnose issues, and make data-driven decisions for improvement. It can also help teams ensure that resources are being used efficiently to reduce costs and that systems are meeting their service level agreements.

Performance data typically includes:

- **CPU utilization**: Percentage of CPU capacity being used by applications or processes
- **Memory usage**: Amount of RAM being consumed by applications or processes
- **Disk I/O (input/output) operations**: Read and write speeds to disk storage
- **Network latency**: Time it takes for data to travel between two points in a network
- **Network bandwidth usage**: Amount of data being transmitted over a network connection
- **Error rates**: Frequency of errors occurring in infrastructure components
- **Database query performance**: Metrics like query execution time and the number of queries per second
- **Application response time**: Time taken by the infrastructure to process a request and return a response
- **Server uptime/downtime**: Duration for which servers or services remain operational or unavailable
- **Virtual machine (VM) resource utilization**: Metrics specific to VMs, such as vCPU usage, memory allocation, and disk usage
- **Container metrics**: Performance of containerized services, such as CPU, memory usage, and restart counts
- **Temperature and power usage**: Physical hardware metrics, such as server temperature and power consumption

The Splunk Common Information Model (CIM) add-on contains a Performance data model with fields and tags that describe performance tracking data, similar to what is described in this article. This includes CPU, memory, storage, uptime and more. You might also be interested in inventory data.

Before looking at documentation for specific data sources, review the Splunk Help information on general data ingestion for Splunk Enterprise, Splunk Cloud Platform or Splunk Observability Cloud.

Add-ons and apps

- Linux/Unix
  - Splunk Add-on for Unix and Linux
  - Splunk Add-on for Linux
  - Monitoring Linux - Metrics and Logs Forwarding
  - Linux Splunk App
- Microsoft
  - Splunk Add-on for Microsoft Windows
  - Splunk Add-on for Microsoft Hyper-V
  - Splunk Add-on for Microsoft Cloud Services
  - Monitoring Windows Containers - Metrics and Log Forwarding
- Oracle
  - Splunk Cloud Addon for OCI (Oracle Cloud Infrastructure)
  - Splunk App for Oracle Cloud Infrastructure (OCI)
  - Splunk Add-On for Oracle Cloud Infrastructure (OCI)
- Splunk Add-on for Sysmon
- Splunk Add-on for Google Cloud Platform
- Qualys Technology Add-on (TA) for Splunk
- Docker
  - Docker Connector
  - Docker App for Splunk
  - Splunk Connect for Docker
  - Docker Monitoring Add-on for Splunk
- Kubernetes Connector

Splunk Lantern articles for the Splunk platform

- Benchmarking filesystem performance on Linux-based indexers
- Comparing Intel and AMD hardware performance for the indexing tier
- Comparing search head hardware changes and their impact on Splunk platform search performance
- Improving hardware utilization by moving indexers into Kubernetes
- Maintaining *nix systems with the Splunk platform
- Maintaining Microsoft Windows systems with the Splunk platform
- Managing an Amazon Web Services environment
- Managing Azure cloud infrastructure
- Managing Splunk Cloud Platform knowledge objects
- Measuring memory utilization by host

- Monitoring NIST SP 800-53 rev5 control families
- Monitoring VMware virtualization infrastructure
- Optimizing Splunk knowledge bundles
- Optimizing systems and knowledge objects
- Performance tuning the forwarding tier
- Performance tuning the indexing tier
- Performance tuning the search head tier
- Recovering lost visibility of IT infrastructure
- Reducing log volume with SPL2 Linux/Unix templates for Edge Processor and Ingest Processor
- Reducing Windows security event log volume with Splunk Edge Processor
- Running a Splunk platform health check
- Sizing your Splunk architecture
- Troubleshooting high resource usage in Splunk Enterprise
- Using ingest actions to filter Windows event logs
- Using the Performance Insights for Splunk app
- Using the Splunk Cloud Monitoring Console effectively

Splunk Lantern articles for Splunk observability products

- Adopting monitoring frameworks - LETS
- Collecting Mac OS log files
- Creating SLOs and tracking error budgets with SignalFlow
- Maintaining *nix systems with Infrastructure Monitoring
- Maintaining Microsoft Windows systems with Infrastructure Monitoring
- Monitoring Kubernetes pods
- Speeding up root cause analysis with artificial intelligence
- Transforming unified network and Telco observability with Fabrix.ai and Splunk