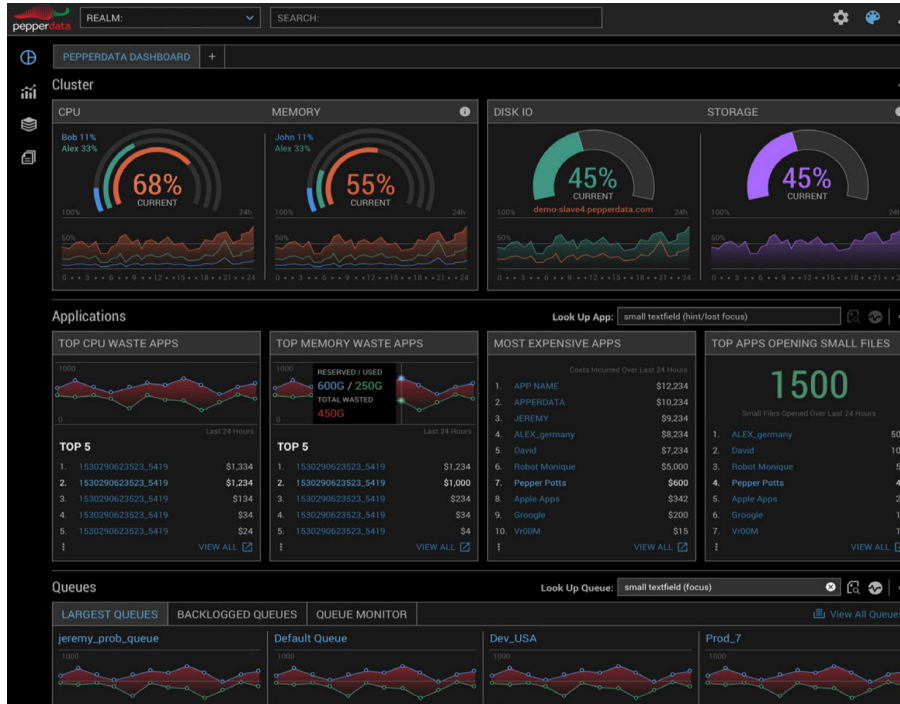




Platform Spotlight



You're constantly looking at different tools to understand the performance of your clusters, to manage and monitor resource capacity, maximize your existing infrastructure investment, and forecast resource needs, but it's impossible to have an accurate view without access to the right data. Your operators are challenged with configuring and sizing critical resources running on multi-tenant clusters with mixed workloads, and you receive alerts without enough detail to isolate and resolve problems. Improving the performance of your clusters and successfully managing capacity requires an understanding of dozens of performance metrics and tuning parameters.

Pepperdata Platform Spotlight continuously collects extensive unique data—that nobody else collects—about your hosts, queues, users, applications and all relevant resources, providing you with a 360° cluster view to quickly diagnose performance issues and make resource decisions.

“Before Pepperdata, we experimented with various approaches to solve our performance issues and we could never see deep enough into the cluster to get insight. Platform Spotlight shined a bright light into our Hadoop environment and provided the detailed data that helped us isolate and resolve the problem.”

Geovanie Marquez, Philips Wellcentive

Platform Spotlight

GET A 360° VIEW OF YOUR PLATFORM

Create a holistic source of operational and performance truth across your clusters.

- **Single source of operational and performance truth.** Access real-time and historical data about the cluster, including system demand, abusive users, and wasteful applications, as well as queue and container sizes. Drill down or zoom out to analyze any application to understand its performance in the context of the entire multi-tenant cluster.
- **Complete instrumentation of your big data ecosystem.** Monitor any process including Kafka, Impala, HBase, Hive, MapReduce, Tez, Spark, and LLAP.
- **Quantify the impact applications have on the cluster.** Identify the impact that applications have on cluster-wide resources such as CPU, disk, memory, network, NameNode, and HDFS.
- **Quantify the impact the cluster has on applications.** Identify which application slowdowns were caused by disk, network or node failures/congestion.
- **Attribute accurate costs to users and business units.** Get detailed usage data of every resource, which when combined with your cost model, generates the most accurate reports to identify the most costly users and business units.

CREATE AND RECEIVE ALERTS ON PLATFORM HEALTH AND PERFORMANCE

Create and receive alerts that identify root causes of performance issues and operational inefficiencies.

- **Identify who is blowing up the cluster.** Track any exhaustible resource, including CPU, memory, and IO to alert on rogue users.
- **Identify which applications will miss their SLAs.** Use data on activity such as duration, amount of data processed, or other metrics to alert on rogue applications.
- **Identify under-performing or failing hardware.** Identify resource contention and get ahead of issues before they become critical.

SUPPORT ENTERPRISE SECURITY REQUIREMENTS

- **No PII collected**
- **Role-based access control**
- **Single sign-on**
- **Support for LDAP integration**

ACCURATELY REPORT ON THE PLATFORM WITH 360° INSIGHTS

Make informed capacity and resource decisions.

- **Accurately plan for growth.** Access to reports with all the relevant real-time and historical information about the cluster—including system demand, users, and applications—on one cohesive capacity planning report.
- **Identify wasteful use of system resources.** Identify users and applications that are wasting system CPU and memory, thus enabling you to prioritize and tune them appropriately.
- **Attribute accurate costs to users and business units.** Get accurate reports of which users and business units are the most costly based on resource utilization.

PEPPERDATA EXPERIENCE

- We continuously monitor over 275 customer production clusters.
- These clusters have over 50,000 nodes spanning all big data distributions and hardware configurations.
- We monitor over one billion jobs and collect over 900 trillion data points annually.

PEPPERDATA EXPERTISE

- We have extensive experience with Fortune 1000 customers.
- Our vast real-time proprietary data helps inform best practices and platform tuning.
- Our extensive experience with Fortune 1000 clusters provides insight to help you select the best hardware and software stack for your applications.
- Our experts help you squeeze the most performance out of your existing capacity and accurately forecast capacity needs

PEPPERDATA SUPPORT

- Amazon Web Services
- Azure
- Google Cloud Platform
- Qubole
- Spark
- MapReduce
- Kafka
- Tez
- Impala
- Hive
- HBASE
- HDFS
- YARN

ABOUT PEPPERDATA

Pepperdata provides analytics performance management that guarantees SLAs and reliability for your infrastructure and apps.



19409 Stevens Creek Blvd., Suite 260
Cupertino, CA 95014 | 408.475.0590

[Schedule a demo](#)
www.pepperdata.com

[Send an email](#)
eval@pepperdata.com