

# Platform Spotlight

Quickly diagnose Infrastructure performance issues and make resource decisions.

## Create a Holistic Source of Performance Truth Across Your Big Data Clusters

You use different big data monitoring solutions to understand your Hadoop clusters, but you still don't have an accurate picture. Pepperdata® Platform Spotlight continuously collects data about your big data clusters, hosts, queues, users, applications, and all relevant resources.

Platform Spotlight enables you to understand and optimize infrastructure performance issues, and to quickly diagnose Hadoop performance issues up to 90% faster while making resource decisions based on user priorities and needs.

**“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



## Get a 360° View of Your Platform

Platform Spotlight gives you a 360° view of your infrastructure and resource utilization, recommendations for optimization, automatic tuning, and alerting. Pepperdata cluster performance monitoring includes real-time and historical information about your clusters including system demand, abusive users, and wasteful applications. Drill down to analyze any big data application to understand its performance in the context of the cluster.

- Monitor technologies and frameworks including Kafka, Impala, HBase, Hive, MapReduce, Tez, Spark, IBM BigSQL, LLAP, EMR, and Dataproc.
- Identify the impact that applications have on cluster-wide resources such as CPU, disk, memory, network, NameNode, and HDFS.
- Identify which application slowdowns were caused by disk, network or node failures/congestion, and quantify the impact the cluster has on application.
- 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.
- Identify the impact of applications on resources such as CPU, disk, memory, network, NameNode, and HDFS.

## 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
- HDFS
- YARN
- MapReduce
- Kafka
- Tez
- Impala
- Hive
- HBASE

## Create and Recieve Alerts on Platform Health and Performance

Platform Spotlight enables you to create and receives alerts about events that interfere with performance.

- Create and receive alerts identifying root causes of performance issues and operational inefficiencies.
- Use data on activity such as duration, amount of data processed to alert on rogue applications and identify which applications will miss their SLAs.
- Achieve optimal application and cluster performance on multi-tenant systems.

## 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

With Pepperdata Application Spotlight, you can make informed capacity and resource decisions.

- 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—users and applications that are wasting system CPU and memory, and prioritize and tune them appropriately.
- Get accurate reports of which users and business units are the most costly based on resource utilization.

## 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](http://www.pepperdata.com)

**Send an email**  
[eval@pepperdata.com](mailto:eval@pepperdata.com)