

# Platform Spotlight

Quickly diagnose big data performance issues.



Platform Spotlight Performance Summary Screenshot

## Get a Correlated View of Your Applications

Pepperdata Platform Spotlight gives you full-stack observability of your big data clusters, hosts, queues, users, applications, and all relevant resources—in any cloud environment and on-premises.

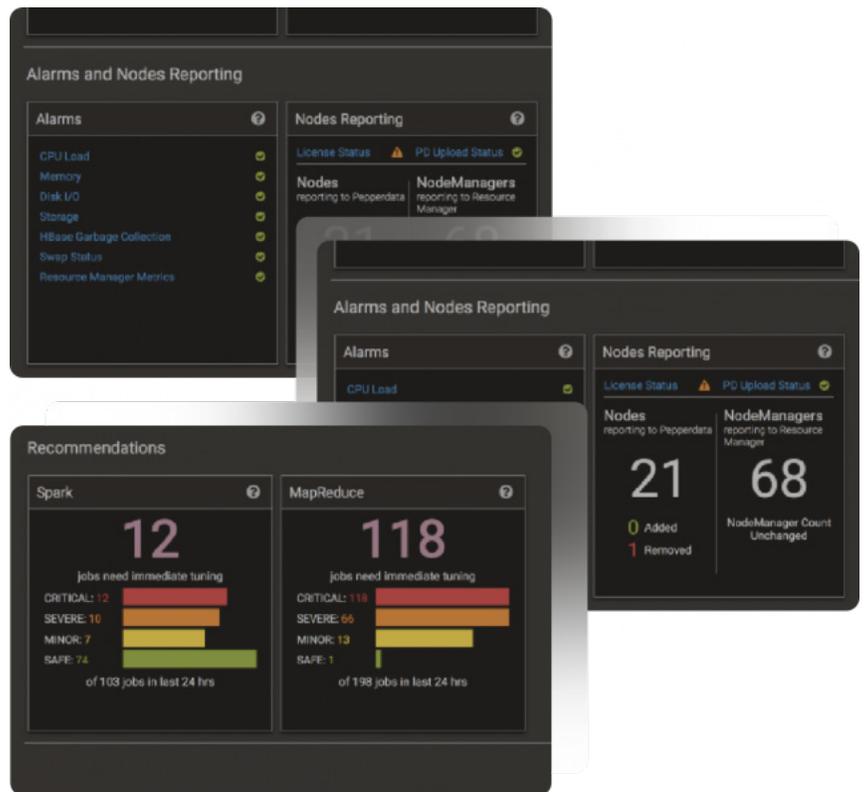
Understand and optimize infrastructure performance issues. Quickly diagnose big data performance issues up to 90% faster while making resource decisions based on user priorities and needs.

Using real-time and historical information covering system demand, wasteful users, and wasteful applications, Platform Spotlight provides custom alerts, so you can rapidly make resource decisions and diagnose performance issues. Platform Spotlight allows you to:

- Monitor supported technologies and frameworks, including Kafka, Impala, HBase, Hive, MapReduce, Kubernetes, Tez, Spark, Amazon EMR, and Google Dataproc.
- Identify the impact that applications have on cluster-wide resources such as CPU, disk, memory, network, NameNode, and HDFS.
- Improve troubleshooting by identifying application slowdowns caused by disk, network, or node failures/congestion, and quantify the impact the cluster has on applications.
- Use IT chargeback reporting to accurately attribute costs to users and business units.

**“Pepperdata provides insightful details and helpful metrics to monitor cluster performance. The user-friendly interface helps to navigate several details like CPU, memory, storage usage, I/O etc. Pepperdata also helps to optimize the cluster settings to ensure enhanced resource management. Pepperdata can integrate automated alerting if any cluster parameter shows alarming behavior. It can also provide features like AI-based detection of resource containment and alert accordingly.”**

Data Engineering Manager



*Platform Spotlight Performance Summary Screenshot*

## Optimize Workloads with Recommendations

Big data operators need to configure and size critical resources running on big data environments, but it's hard to get an accurate view without access to the right data. It's even harder when operating in complex cloud settings. To make this easier, Pepperdata provides actionable reporting, insights, and recommendations to rightsize containers, queues, and other resources.

Platform Spotlight does this by continuously collecting extensive unique metrics and data about your hosts, queues, users, applications, and other relevant resources, providing you with insights to quickly diagnose performance issues and make resource decisions. Platform Spotlight allows you to:

- Size big data queues appropriately based on criticality.
- Get self-service recommendations on container sizes and heap reservations.
- Identify growth trends and accurately forecast resource needs.

## Pepperdata Experience

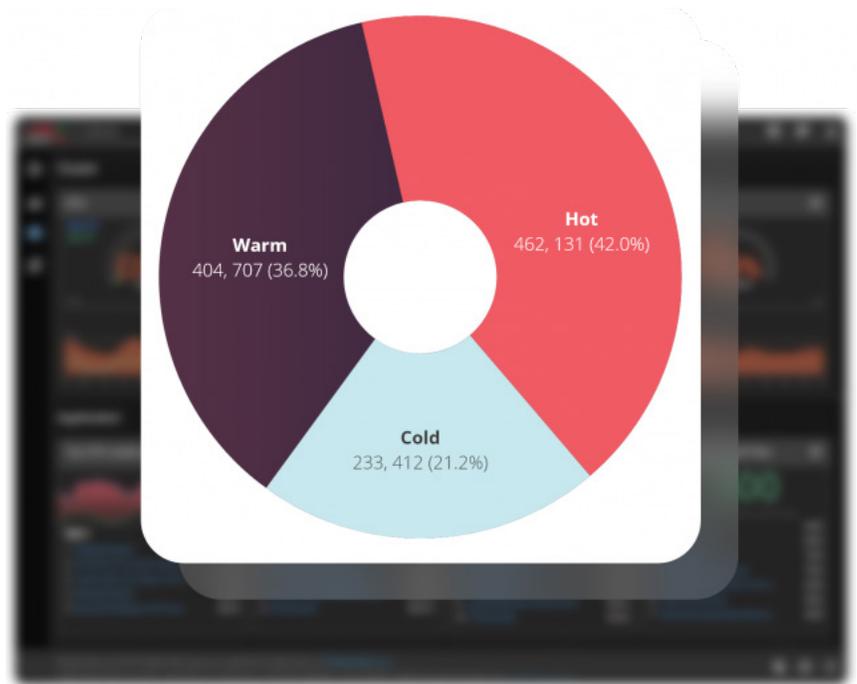
- We continuously monitor over 275 customer production clusters.
- These clusters have over 50,000 hosts spanning all big data distributions and hardware configurations.
- We monitor more than one billion jobs and collect over 900 trillion data points annually.

## Pepperdata Expertise

- We have extensive experience with Fortune 1000 customers.
- Our vast real-time proprietary data informs best practices and subsequent 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
- Kubernetes
- Tez
- Impala
- Hive
- HBASE



*Platform Spotlight Data Temperature allows you to get the age and size of your HDFS files.*

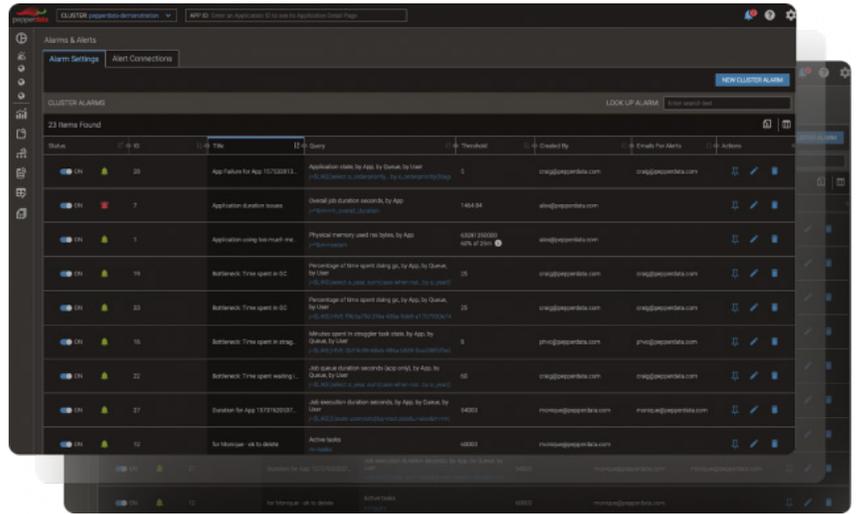
## Store Your Data More Efficiently with HDFS Data Temperature

You probably don't know the exact age of every file you have stored in HDFS, or which files should be archived or moved to SSDs based on how frequently (or infrequently) they are accessed. With Platform Spotlight, you can generate a summary of policy definitions and rolled-up data on the state of HDFS, including the total number of files and where they fall within the policy spectrum. Included is a cold file breakdown that details each policy's coverage and each file's state relative to the governing policies. This breakdown includes the exact file names and sizes for each temperature. With the Pepperdata Data Temperature report, you can:

- Optimize your data and reduce your storage spend.
- Get a comprehensive view of how much movement is required to adhere to policy.
- Get the age and size of your HDFS data files.
- Access exact file names for each temperature.
- See which files don't match their current policy based on access times.

# Support Enterprise Security Requirements

- Role-based access control (RBAC)
- Single sign-on (SSO)
- Support for LDAP integration
- No PII collected



Platform Spotlight Alerts Summary Dashboard

## Create and Receive Alerts on Platform Health and Performance

Platform Spotlight enables you to create and receive alerts about events that degrade performance. You can:

- Quickly perform root cause analysis (RCA) of performance issues and operational inefficiencies.
- Use data on activity such as duration and 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.

## About Pepperdata

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

©2021 Pepperdata Inc. All rights reserved. Pepperdata and the Pepperdata logo are trademarks or registered trademarks of Pepperdata Inc. All other trademarks are the property of their respective owners. Pepperdata reserves the right to change this document without notice. To ensure you have the latest version of this document, visit [www.pepperdata.com](http://www.pepperdata.com).



3945 Freedom Circle, Suite 920  
Santa Clara, CA 95054

**Start a Free Trial**

[www.pepperdata.com/trial](http://www.pepperdata.com/trial)

**Send an Email**

[eval@pepperdata.com](mailto:eval@pepperdata.com)