Application Spotlight

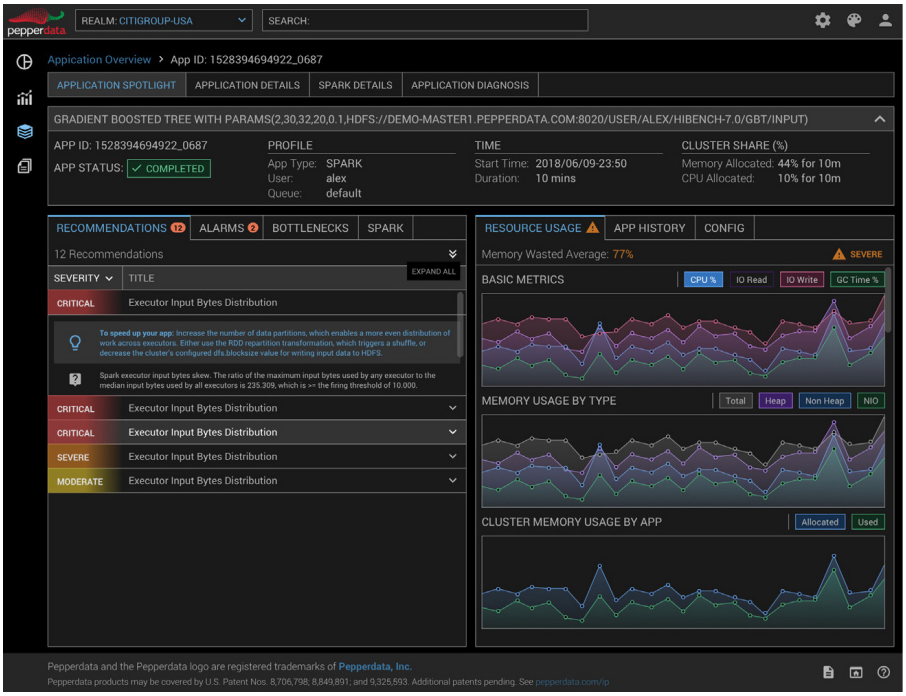

Application Spotlight

You're constantly looking at different tools to understand the performance of your applications to ensure SLAs, avoid failure, improve efficiency, and monitor resource capacity, and you still don't have an accurate picture because you don't have access to the right data. You submit queries and wait. You have no insight into when a query or app starts running, if there are delays, or whether it will finish. Developing fast and efficient applications requires an understanding of dozens of performance metrics and tuning parameters.

Application Spotlight provides you a 360° view of all this data in one place, so you can gauge performance in the context of the entire cluster, quickly diagnose application performance issues up to 90% faster, and improve overall efficiency. Pepperdata also provides automatic tuning for recurring applications, delivers job-specific recommendations and allows you to set up alerts on specific behaviors and outcomes to avoid the risk of failure.

Av 360° APPLICATION VIEW

Create a holistic source of application performance truth across your clusters.

Self-service access to your application performance data. Profile and optimize application performance via recommendations and key performance indicators.

Complete instrumentation of your application. Understanding exactly what CPU and memory resources it requested, what it needs, what it used, and what it wasted.

Quantify the impact that the cluster has on your applications. Identify the impact that the queue congestion, bottlenecks, and hardware failures have on application performance.

At APPLICATION TUNING

Improve performance and efficiency of your applications.

Single source for application performance data presented in context of the cluster. Self-service access to all of the data on your applications in one place. The ability to distinguish whether performance issues were caused by your application or other applications on the cluster.

AI-driven optimization of recurring applications. Recurring applications account for a substantial portion of workload. These applications typically need to meet constantly changing SLAs. Pepperdata auto tunes configurations to optimize resource utilization or runtime.

Ar APPLICATION RECOMMENDATIONS

Achieve optimal application performance.

Change configuration parameters to optimize performance. Get self-service recommendations on data partitioning and serialization.

Tune CPU and member reservations based on actual consumption. Get self-service recommendations on container sizes and heap reservations.

Change queue selection or launch time based on cluster activity. Identify the best queue and launch time for applications based on current workloads to improve runtime and maximize performance.

Aa APPLICATION ALERTING

Create and receive alerts about events that interfere with application performance.

Identify cluster bottlenecks. Use Pepperdata's rich real-time data to identify resource bottlenecks, including CPU, memory, and IO.

Identify application bottlenecks. Use Pepperdata's data to pinpoint straggling tasks or poor parallelization that can significantly impact runtime.

Identify applications at risk of missing their SLAs. Use Pepperdata's data to alert on duration, amount of data processed or other milestone.

In 360° INSIGHTS

Make informed capacity and resource decisions.

Accurately plan for growth. Access reports with all relevant real-time and historical information about the cluster including system demand, users, and applications, one cohesive capacity planning report.

Identify wasteful use of system resources. Identify users and applications wasting system CPU and memory, enabling you to prioritize and tune them appropriately.

Attribute accurate costs to users and business units. Get accurate reports on which users and business units are the most costly based on resource utilization.

> "Visibility into the metrics and being able to visualize Java metrics was critical. Before we installed Pepperdata, we didn't have a way to access JVM metrics or drill down to the level of granularity that was needed to effectively find and resolve issues the way we do now."
>
> *Jesse Escobedo, Rubicon Project*

PEPPERDATA EXPERIENCE

- We continuously monitor over 250 customer production clusters.
- These clusters have over 30,000 nodes spanning all Big Data distributions and hardware configurations.
- We monitor over 550 million jobs annually and collect over 600 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

- Pepperdata supports clusters running on-premise and in the cloud and is certified on all Big Data distributions:
  - Spark
  - MapReduce
  - Kafka
  - Solr
  - Tez
  - IBM BigSQL
  - Impala
  - Hive
  - Kubernetes
  - HBASE
  - HDFS
  - YARN

PEPPERDATA BIG DATA SUCCESS ELEMENTS:

	APM PRODUCTS				EXPERIENCE		EXPERTISE	
	Av 360° APPLICATION VIEW	**At** APPLICATION TUNING	**Pv** 360° PLATFORM VIEW	**Pt** PLATFORM TUNING	250 + PRODUCTION CLUSTERS	30K + NODES	**Bp** BEST PRACTICES	**Pp** PERFORMANCE PLANNING
In 360° INSIGHTS	**Ar** APPLICATION RECOMMEND-ATIONS	**Aa** APPLICATION ALERTING	**Pr** PLATFORM RECOMMEND-ATIONS	**Pa** PLATFORM ALERTING	550M + JOBS ANNUALLY	600t + DATA POINTS ANNUALLY	**Cp** CAPACITY PLANNING	**Ad** ARCHITECTURE DESIGN

amazon web services · APACHE SOFTWARE FOUNDATION · MAPR · IBM · Hortonworks · cloudera

pepperdata

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

Schedule a demo
www.pepperdata.com

Send an email
eval@pepperdata.com