Unravel® for Spark

Unravel® for Spark

A modern generation of Application Performance Management (APM) software has emerged to support analytical, machine learning, IoT, and artificial intelligence applications running on big data platforms such as Spark, Hadoop, Kafka, and NoSQL.

Unravel for Spark provides a comprehensive full-stack, intelligent, and automated approach to application performance management (APM) across the big data architecture. The Unravel platform helps you to analyze, optimize, and troubleshoot big data apps and operations in a seamless, easy to use, and frictionless manner.

Unravel for Spark bridges the gaps that exist in current fragmented approaches for big data APM and and operations management.

The **Unravel for Spark** provides a 360° view into the behavior of Spark applications on a single glass pane with all relevant and correlated information. Some of the unique capabilities, include:

Locate and analyze Spark app performance

- KPIs specific to an application, such as status, duration, data I/O, # of stages, tasks, etc.
- Drill down views from Spark jobs to stages to task execution; Gantt chart of stage timelines to bottlenecks, errors, task logs of drivers and executors, and configurations

AI driven Intelligence engine provides insights into

- Utilization of memory resources and Spark storage memory
- RDD Caching, CPU resource contention, and container resource utilization

AI driven recommendation engine to improve app efficiencies in simple English

- Spark executor memory/instances value
- Spark default parallelism new value

For failing and inefficient Spark Apps, Unravel provides

- Root cause analysis error view —metric and parameter tweaks to get the app back running
- Pro-active alerts to detect/eliminate rogue apps that can affect cluster performance, resource utilization and SLA requirements

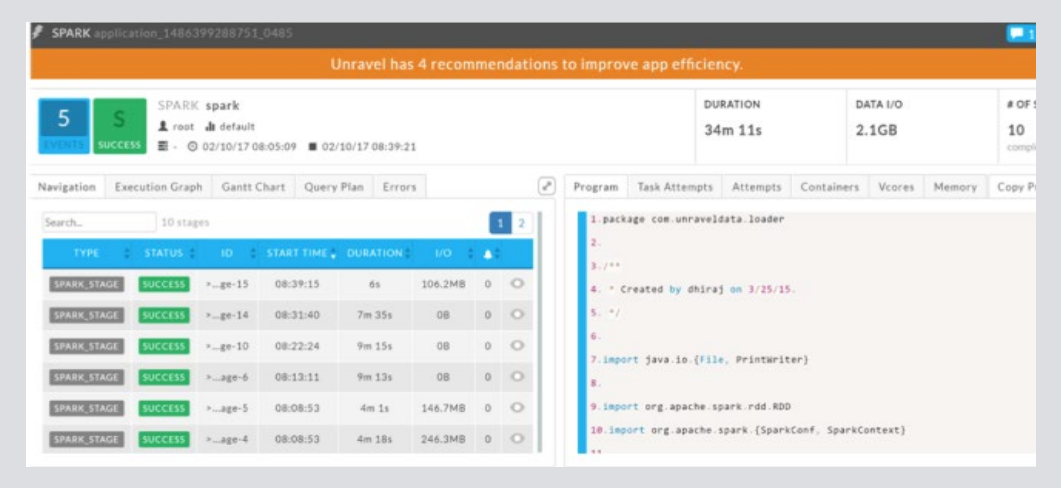

Business benefits of using Unravel

Unravel makes sure businesses can achieve their goals with big data spark applications. Some of the tangible benefits, include:

- Improved reliability to meet SLAs, ensuring no disruption in revenues
- Optimized resource utilization, resulting in lower infrastructure and project costs
- Improved productivity for all stakeholders

Other features include:

- Errors and warning messages
- Graph the resources the application/job/ stage consumed
- Histograms showing the distribution of map and reduce task duration, input and output size
- Donut graphs show the percentage of successful (green) and of failed (orange) tasks

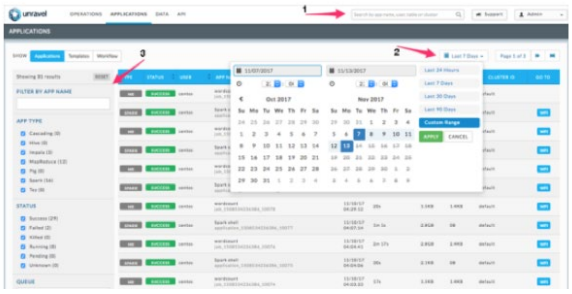

Search for your app(s) in a variety of ways

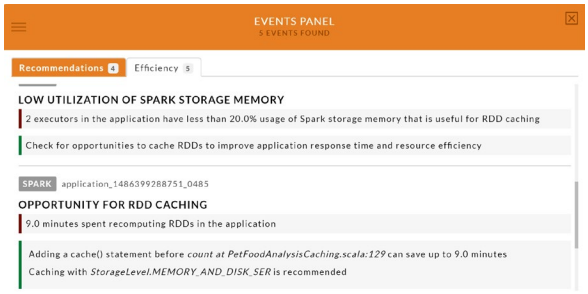

Events/Recommendations/Efficiency: query used too many reducers. Unravel shows that this query has one job using too many reduce tasks, and provides a recommendation for the configuration

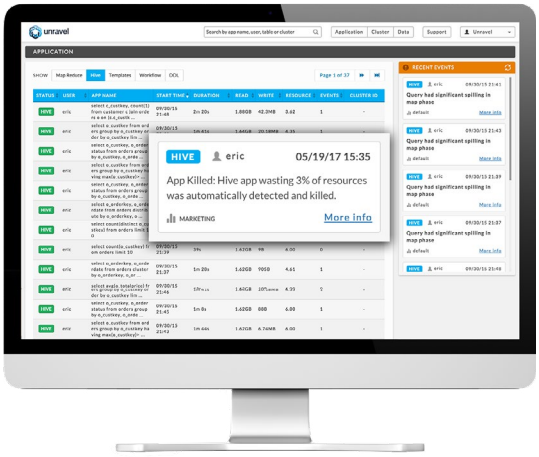

Enforcing automated policies and actions to meet SLA needs, diagnose and kill rogue apps

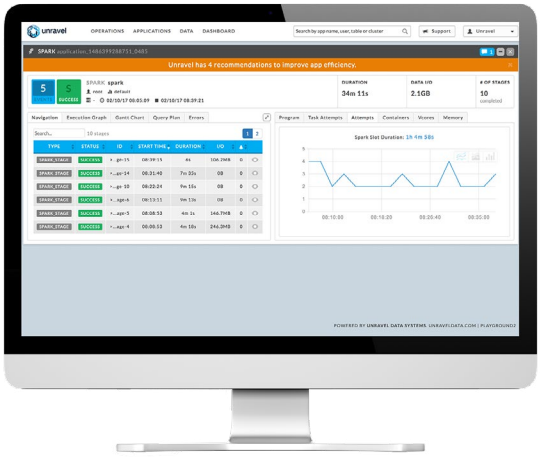

Built-in intelligence
KPIs, tasks, jobs, stages related to an application (left)
Attempts, containers, memory (right)