

Unravel Newsletter: November 2019

Greetings from Unravel,

In this month's newsletter, we start with an exciting announcement about Unravel's new offering for Google Cloud Dataproc. Next, we share information about an Unravel AI Chatbot for simplifying Spark Performance Management, followed by a discussion of Forecasting Reports added to Unravel's portfolio of On-Demand reports recently. We will continue with some community highlights in the area of Data and Analytics and include links to recent recordings and upcoming events that you may find interesting.

New! Announcing Unravel for Google Cloud Dataproc

We are excited to announce the Unravel offering for Google Cloud Dataproc. Please check out the launch announcement: [Unravel Launches Performance Management and Cloud Migration Assessment Solution for Google Cloud Dataproc](#) and more information about this offering in the following links:

- <https://unraveldata.com/google-cloud-dataproc/>
- [Datasheet - Unravel for Google Cloud Dataproc](#)
- [Unravel Cloud Migration Sample Report for GCP](#)

With this offering, Unravel provides **Application Performance Management & Operational Intelligence** for Google Cloud Dataproc.

Value proposition:

- Better SLA: Run apps more reliably, improve performance
- Lower Costs: Run apps more efficiently
- Lower MTTR, Increase Productivity: Detect, troubleshoot and fix issues quickly
- Dashboards and Reports (usage, chargeback/show back, resource usage etc.)

Features and Functionality :

- Application Performance Management for Modern Data Workloads and Systems

- Automated Troubleshooting & Root-Cause Analysis
- Tuning Recommendations, Auto-tuning of Apps
- AI for DataOps
- Chargeback/Show back
- Resource and Cost Optimization

You are welcome to [get in touch](#) with us to try this out! As an early customer, your feedback will carry utmost weight in how this offering shapes up further.

Stay tuned for more exciting announcements in future!

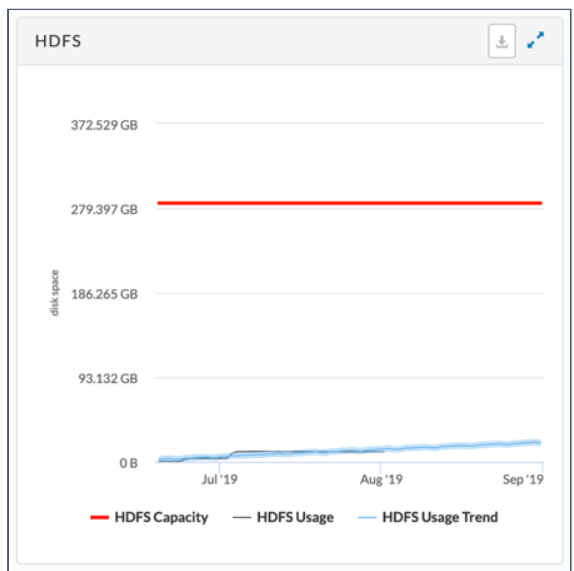
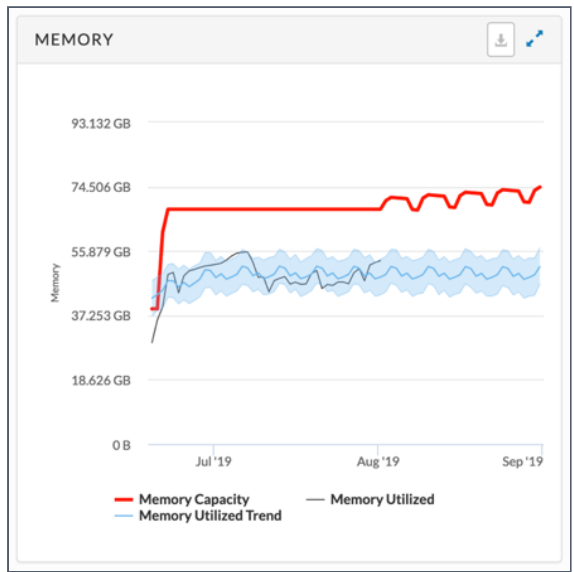

Unravel AI Powered Chatbot to simplify Apache Spark Performance Management

Watch Unravel CTO, Co-founder, and Duke University Adjunct Professor Shivnath Babu present a session about an AI Powered Chatbot to Simplify Apache Spark Performance Management [here](#). If you are interested in using such a chatbot in your organization or building one off of this, please do get in touch with us.

Unravel Forecasting Reports

Several of our customers requested features that would help them see the trending of their resource usage and to predict future needs so they could use Unravel for capacity planning. As a result, we added a set of Forecasting Reports to Unravel's portfolio of On-Demand reports.

Specifically, reports help you monitor CPU, Memory and HDFS disk capacity usage and carry our capacity planning for future. Unravel uses historical usage data to extrapolate capacity need trends. Using this, you can more effectively plan for additional resource needs. Next, some example screenshots of these reports are shared:



(Top to bottom) Forecasting - CPU, Memory, HDFS

Check out [this section](#) in the Unravel User Guide to help you get started with using these reports.

Keep looking at the newsletters for upcoming exciting additions to this report or get in touch with us directly to get your feature requests in!

Conference Sessions & Recordings

- Check out [How to get started with Unravel for Amazon EMR from the AWS Marketplace](#).
 - Check out [How to get started with Unravel for Azure HDInsight from the Azure Marketplace](#).
-

Upcoming Events & Webinars

- Join us at the [Phoenix Data Conference](#), in Phoenix on Nov 09 where Carlos Morillo, Unravel Solutions Engineer will present a session on Spark application performance management for Azure Databricks.
 - Join us at [Big Data LDN](#), Olympia, London, Nov 13 - Nov 14 at the Unravel Booth. Also looking forward to seeing you on 11/13, at 13:10 where Chris Santiago, Unravel Solution Engineering Manager will be presenting a session on Effective Cloud Migration Strategies for your Modern Data Applications. More details [here](#).
 - Join us at [AWS re:Invent](#) where our CTO and Co-founder Shivnath Babu, will be presenting a session on Architecting High-Performance Data Pipelines when Migrating to AWS.
-

Community Highlights

- [Introducing Beekeeper - Managing stale data in the cloud](#): Beekeeper is a new open-source tool from Expedia for cleaning up unreferenced data in Apache Hive tables. It listens for metastore changes (like changes to table or partition locations) and periodically deletes abandoned data in S3. The github repo has instructions for trying it out using docker.
- [Getting Started with Rust and Apache Kafka](#): A post on the Confluent blog describes porting a non-trivial Kafka application from Clojure to Rust. The author describes the tradeoffs between the various Rust libraries for Kafka, how to extend a client to support Avro records/Schema Registry, and shares some

benchmarks comparing Clojure and Rust performance/memory usage.

- [Scaling ETL: How data pipelines evolve as your business grows:](#) Data drives so much of what people at Grubhub do. Their engineers like to think of the data in two ways: source data - data coming from a single point / user, and analytics data - compiled data coming from multiple points and multiple sources. This is the story of their data pipelines.
- [Improving Amazon Redshift Performance: Our Data Warehouse Story:](#) Engineers at Udemy share their history with Amazon Redshift and highlight some of the pitfalls and lessons learned along the way. They also discuss some of the tools they use to ensure that they're getting the most out of their data warehouse.

Resources

- [Learn more](#) about Unravel.
- [Online Product Demo](#)
- [Unravel Partners](#)
- [Unravel Product Releases and Documentation](#)
- [Unravel Datasheet](#)
- [More Unravel News](#)



[Contact Us](#). [Sign Up for 30-day Trial](#). [Take Unravel for a Test Drive](#).

2019 Unravel. All Rights Reserved. 2 Palo Alto Sq, Suite 120, Palo Alto, CA 94306