

Unravel Newsletter: September 2019

Greetings from Unravel,

In this month's newsletter, we start with an exciting announcement about Unravel's new offering for Azure Databricks. Next, we dive into Unravel's capabilities for monitoring HBase – a popular choice of NoSQL database for our customers. 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 Azure Databricks

We are excited to announce the Unravel offering for Azure Databricks. Please check out the launch announcement [here](#) and more information about this offering in the following links:

- [Learn more](#) about Unravel for Azure Databricks
- Watch the [video demo](#) of Unravel for Azure Databricks

With this offering, Unravel provides **Application Performance Management & Operational Intelligence** for Azure Databricks.

We have focused on Data Engineering use cases, Unravel Analyzes & Optimizes workloads on Job Clusters across all your Databricks Instances & Workspaces.

Features and functionality:

- Faster Diagnosis & Resolution of Application Issues
- Higher Cluster Resource Efficiency
- Unified View Across Workspaces & Instances
- Ensuring Reliable Apps in Production
- Usage Breakdown & Trending
- Speed up Applications
- Single Pane of Glass for Apps, Resources, Data, & Users

All the Databricks Job types: Notebook, Jar, spark-submit, and Python Task can be monitored by Unravel. This functionality is available on a single Unravel Deployment across all your Databricks Instances & Workspaces.

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.

Also, check out the [Unravel for Azure Databricks](#) section in the Unravel User Guide to help you get started.

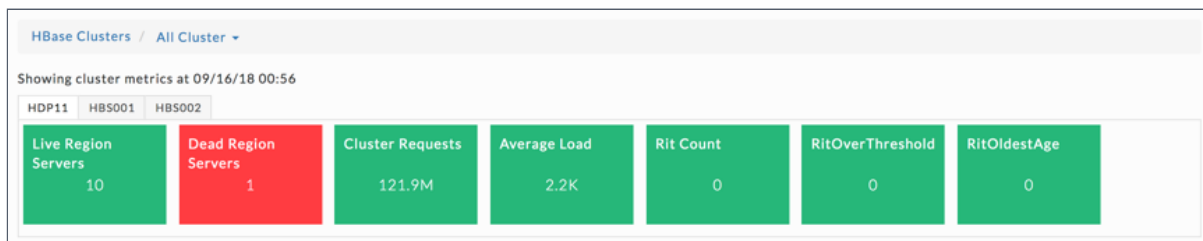
Stay tuned for more exciting announcements in future!

Monitor and Manage your HBase clusters using Unravel

Apache HBase is a distributed column-oriented database built on top of the Hadoop file system. A wide-column data store modeled after Google Bigtable, this is quite a popular NoSQL system used by many. When it comes to operating and maintaining HBase clusters, there are multiple factors that can affect the performance and health and it isn't simple to deal with those issues. While most cluster managers provide high level metrics, Unravel collects and surfaces detailed data and metrics to help you quickly identify root causes of issues and improving performance of the cluster. Some quick highlights of the Unravel solution for HBase:

Summary view of all Clusters

From the above interface, you can now choose to drill down to a particular cluster. As you may notice, throughout the UI, there is clear coding to indicate the health and status of each component.



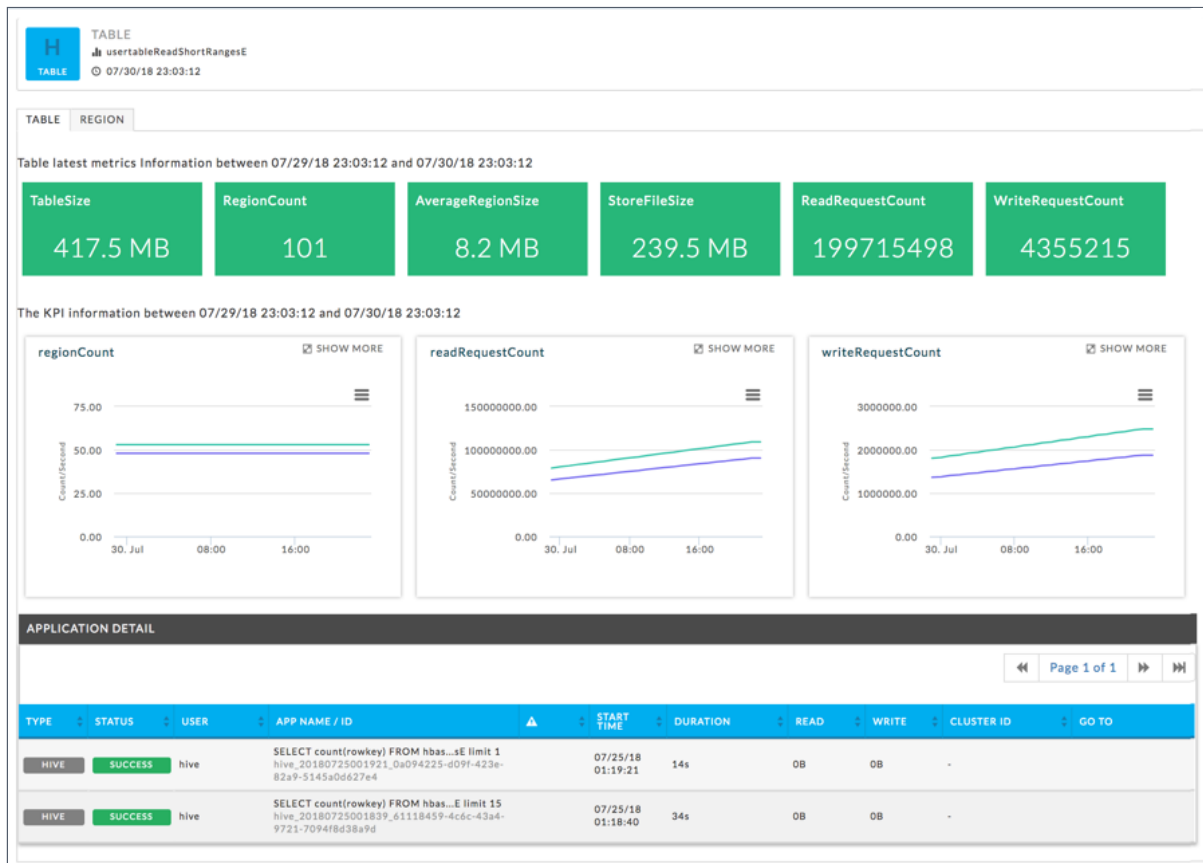
Region Servers View

This lists the cluster's region servers, with their KPIs and health, various more details graphs and charts are available for each of the KPIs.

The table shows region server metrics for various servers. It includes a search bar for region servers and a table of metrics:

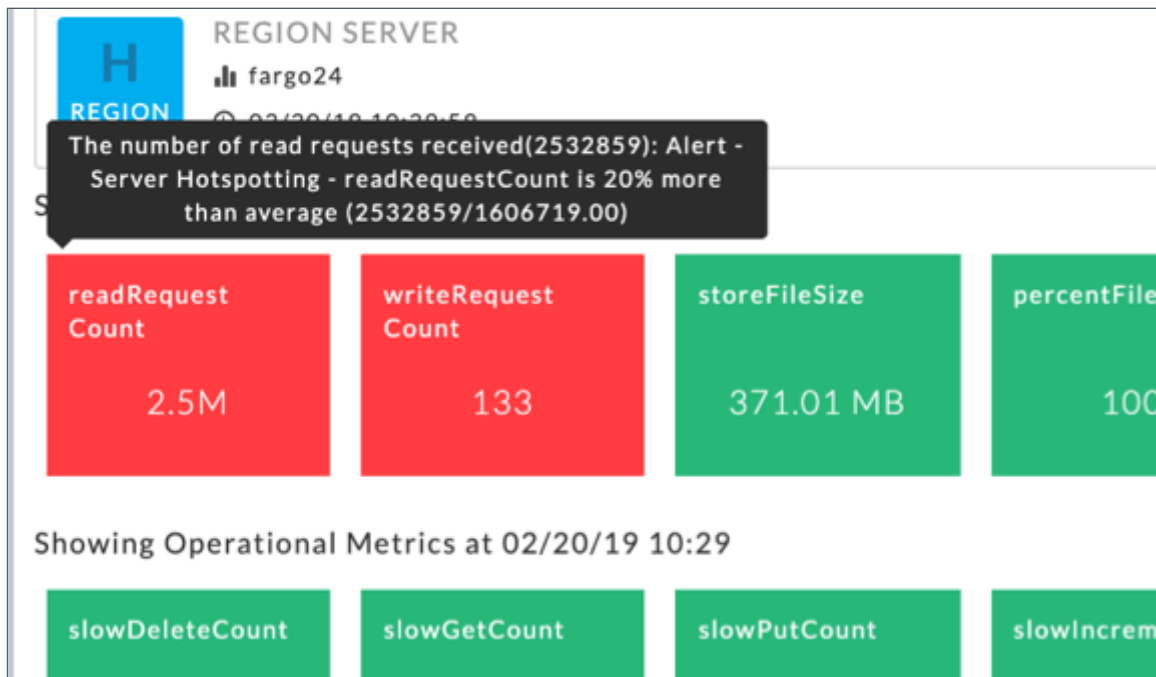
REGION SERVER NAME	READ REQUEST COUNT	WRITE REQUEST COUNT	STORE FILE SIZE	PERCENT FILESLOCAL	COMPACTION QUEUE LENGTH	REGION COUNT	HEALTH
fargo1	0	13006200	282.2 MB	0	0	2342	Good
fargo7	0	10620600	281.6 MB	99.26	0	2035	Good
fargo8	0	13395400	296.7 MB	99.47	0	2342	Good
fargo3	0	10079400	253.8 MB	97.84	0	1987	Good
fargo6	0	12860800	292.7 MB	99.15	0	2340	Good
fargo5	55773439	11502068	324.7 MB	99.38	0	2009	Bad
fargo2	0	12960400	283.8 MB	0	0	2341	Good
fargo4	0	12918200	294.0 MB	98.54	0	2340	Good
fargo9	0	13089200	288.0 MB	99.67	0	2343	Good
fargo16	0	9599200	236.2 MB	98.03	0	1776	Good

There is also a Tables View that lists all the tables associated with the cluster, their KPIs and the table's health. And for every table in the system, you can get to the specific list of applications associated with that table.



Server Hotspotting

Based on collecting and analyzing detailed data and metrics, an instance where Unravel has detected and alerted on the occurrence of Server Hotspotting:



[Raushendra Kumar](#), Senior Software Engineer at Unravel, recently published a blog illustrating how you can use Unravel's HBase Application Program Manager (APM) to debug issues in your HBase cluster and improve its performance. Check it out: [End-to-end monitoring of HBase databases and clusters](#).

Also check out the [HBase](#) section in the Unravel User Guide to help you get started with monitoring HBase clusters in Unravel. Also see the [HBase Alerts and Metrics](#) section for HBase metrics descriptions and list of alerts with suggested actions as you start using this.

Conference Sessions & Webinar Recordings

- Watch this webinar by Aengus Rooney, Head of Solution Engineering - International, Unravel - [How to Optimize Spark Data Pipelines on Azure Databricks](#). Develop an understanding of the performance dynamics of modern data pipelines and applications. In this session, you will learn about uncovering and understanding the key datasets, metrics, and best practices needed to develop mastery with Spark performance management on Azure Databricks.
- Check out this video - [Founder Real Talk Episode #27 with Kunal Agarwal, Co-founder & CEO of Unravel Data](#), where Glenn Solomon, Partner at GGV Capital talks to Unravel Co-founder and CEO Kunal Agarwal.
- Check out this video where ClearSense CIO Charles Boicey explains why you'd be out of your mind to monitor your big data environment without Unravel - [ClearSense on keeping their big data promises in Healthcare](#).

Upcoming Events & Webinars

- Th, Sep 5: [Unravel for Cloud Migration](#) – Join Abha Jain, Director, Products, Unravel to learn about Unravel’s features and functionality that can help you migrate your big data environment to the cloud in a faster and more cost-efficient way.
 - Join us the Unravel Booth at the [Strata Data Conference](#), in New York, Sept 23-26.
 - Join us at the Unravel Booth at the [Kafka Summit](#), in San Francisco, Sept 30 - Oct 1.
 - Join us at the Unravel Booth at [AWS Summit Global Summit](#), in Toronto on Oct 3.
-

Community Highlights

Some interesting articles in the area of Data and Analytics:

- Check out [Databricks Academy](#) to go through trainings and learn more about Databricks.
 - A well-regarded book to develop in-depth knowledge in [HBase: HBase: The Definitive Guide: Random Access to Your Planet-Size Data](#). This book shows you how Apache HBase can fulfill your needs of scalable storage solution to accommodate a virtually endless amount of data. As the open source implementation of Google's Bigtable architecture, HBase scales to billions of rows and millions of columns, while ensuring that write and read performance remain constant. Many IT executives are asking pointed questions about HBase. This book provides meaningful answers, whether you're evaluating this non-relational database or planning to put it into practice right away.
 - [Python Data Transformation Tools for ETL](#) - This post covers some of the available tools out there to get you started on Extract Transform Load pipelining in Python - starting from Pandas, then Dask, modin, petl and PySpark.
 - [Inside Facebook's New Robotics Lab, Where AI and Machines Friend One Another | WIRED](#) - The social network has a plan to merge the worlds of artificial intelligence and real-world machines, so that both may grow more powerful. This article shows how they use data and AI in the new Silicon Valley HQ.
-

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.](#)

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