

Unravel Newsletter: January 2020

Greetings from Unravel,

Hope you have had a restful holiday and wish you very Happy New Year!

In this month's newsletter, we start with discussing a much awaited and requested feature: multi-cluster support in Unravel. Then, we will discuss two promotions for getting started on Unravel cloud offerings, followed by an exciting announcement about Unravel joining the AWS Partner Network Global Startup Program. 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.

Multi-Cluster Support in Unravel

For the past few months, we have discussed with you and understood that for a variety of reasons including the following, there is an urgent need for adding multi-cluster support for Unravel:

- Analyze workloads from Test, UAT and Prod clusters in one place
- Comparing performance of apps as it goes through SDLC
- Single pane of glass into your business across all clusters
- Ease of management of one Unravel instance

Our engineering team has been hard at work to come up with the architecture and implementation for this. With this in place, here are some of the features and functionality you can expect:

- On various operational screens, you can choose to view metrics aggregated across clusters or choose to see the details of one of the selected clusters only. For example:

Select "All Clusters" to see aggregate metrics or select any one of the clusters being monitored to see its details

The screenshot shows the 'OPERATIONS' section of the Linnivel interface. It features several charts: 'FINISHED YARN APPLICATIONS' with success and failed counts, 'RUNNING YARN APPLICATIONS' with running and pending counts, and 'RESOURCES' with VCore and Memory usage. On the right, a 'RECENT EVENTS & ALERTS' list shows various application events. A dropdown menu is open over the 'All Clusters' button, listing several cluster names.

- On the Applications View, you can see the Cluster on which any given application ran by looking at the "Cluster ID" column, or you can choose to see the list of applications running on a particular cluster only by filtering on the corresponding Cluster ID.

In the "Applications View", Cluster ID column shows which cluster the Application ran on

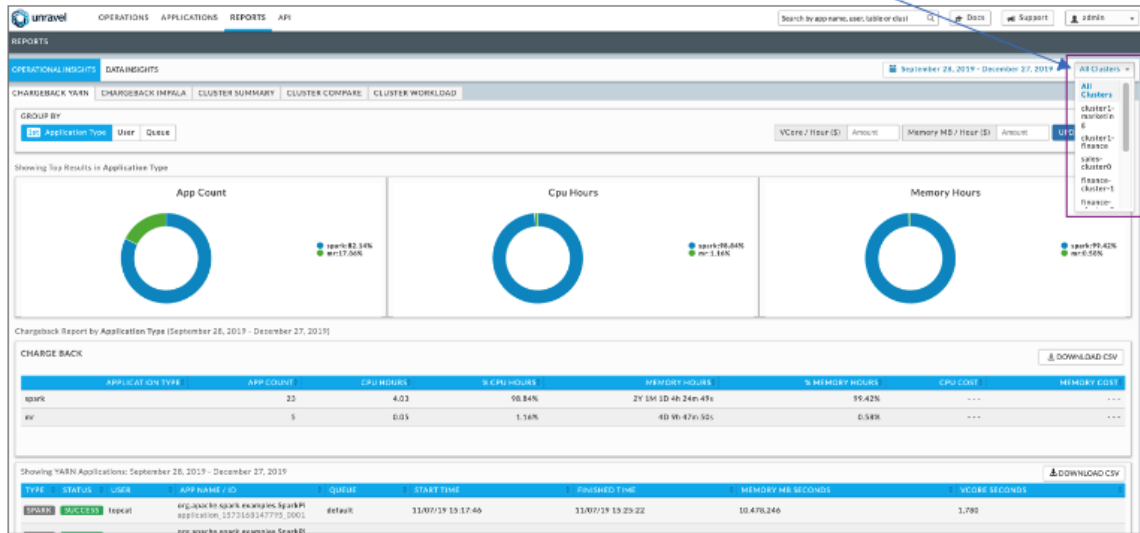
The screenshot shows the 'APPLICATIONS' view with a table of 28 results. The table has columns: TYPE, STATUS, USER, APP NAME / ID, START TIME, DURATION, READ, WRITE, CLUSTER ID, and QUEUE. A red box highlights the 'CLUSTER ID' column, which contains values like 'democluster-2', 'finance-cluster-1', and 'sales-cluster-0'.

In the "Applications View", select only the applications that ran on specific cluster(s) by filtering on those

This screenshot shows the same 'APPLICATIONS' table as above, but with a filter dropdown menu open over the 'CLUSTER' column. The dropdown menu lists several cluster names, including 'cluster1-finance', 'cluster-7317', 'finance-cluster-1', 'sales-cluster-0', 'democluster-2', 'finance-cluster-0', and 'sales-cluster-1'.

- Various Reports support the same functionality as well, you can select "All Clusters" to see aggregate metrics across clusters.

Chargeback Report: Select "All Clusters" to see aggregate metrics or select any one of the clusters being monitored to see its details separately



Cluster Compare Report: Select "All Clusters" or a specific cluster only



Promotions for getting started with Unravel Cloud Offerings

Unravel is offering some exciting promotions for you to try out two of our Cloud offerings:

- Try Unravel on Amazon EMR for 14 days with no obligation to buy, and we will credit you \$2000 in AWS credits which you may redeem for your AWS account. [Start here.](#)
- Try Unravel on Azure HDInsight for 14 days with no obligation to buy, and we will credit you \$2000 in Azure credits which you may redeem for your Azure account. [Start here.](#)

Unravel Joins AWS Partner Network Global Startup Program

We are excited to announce our membership in the Amazon Web Services (AWS) Partner Network (APN) Global Startup Program.

The APN Global Startup Program is a unique “white glove” support and go-to-market program for selected startup APN Partners, allowing members to build on their AWS expertise, better serve shared customers, and accelerate their growth. To be selected for the APN Global Startup Program, Unravel had to meet

predefined criteria, including a clear, demonstrated product market fit for an innovative enterprise tech product, be backed and recommended by a top-tier venture capital firm, and demonstrate a strategic commitment to building their AWS and cloud expertise. Check out the complete blog [here](#).

Keep subscribing to these newsletters for upcoming exciting Unravel features and functionality or [get in touch](#) with us directly to get your feature requests in!

Conference Sessions and Recordings

- Watch the recording of this webinar by Unravel and AWS: “[Application Performance Management & Operational Intelligence for Amazon EMR](#)”
- Watch the recording of this webinar by Unravel and AWS: [Look Before You Leap: Migrating On-Premises Hadoop to AWS](#)
- Check out this webinar Recording: [Best Practices: Troubleshoot and Optimize Spark Data Pipelines with Unravel](#)
- Check out this webinar Recording: [Unravel Demo - Big Data Application Performance Management](#)
- 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 Hy’s Steakhouse in Toronto for Toronto [IT Leaders & "Big Data" Professionals Luncheon \(sponsored by Microsoft Azure & Unravel Data\)](#), in Toronto, Canada, January 30, 2020.
- Join us at the AWS Spheres for [Dinner with AWS and Unravel + Amazon Spheres Tour](#), in Seattle, February 6, 2020.
- [Debugging and Demystifying Spark with Unravel](#) - Join Dave Goodhand, Head of Customer Success EMEA, Unravel Data on 30 Jan 7:30 pm PT to learn about uncovering and understanding the key datasets, metrics, and best practices needed to develop mastery with Spark performance management with Unravel Data.

Community Highlights

- [Big Data Quality at Criteo](#): Criteo writes about how they enforce data quality for the 450 PB that they have in their Hadoop clusters. Since Hadoop is very flexible/lenient in the data it supports, they run over 7,000 data quality checks per day using statistics from Hive tables and custom queries.
- [Dynein: Building an Open-source Distributed Delayed Job Queueing System](#): Airbnb has open sourced Dynein, their job queuing system that

they use for offloading tasks from the main request path and performing other asynchronous operations. It uses DynamoDB as a scheduler for future jobs and SQS for queuing—the post describes how this is built in a highly scalable way.

- [Running and Deploying Elasticsearch on Kubernetes](#): An overview of Kubernetes and how to deploy Elasticsearch on Kubernetes. This is a great introduction to many of the core concepts of Kubernetes (e.g. Deployment, Pod), including those that are important for running a stateful service (e.g. StatefulSet, PersistentVolumeClaim). It also shows how to configure accounts for ES using Kubernetes RBAC and use the helm package manager for deploying to Kubernetes.
- [A Billion Rows A Second](#): Wouldn't it be great if you could load a 1 Terabyte data file in Python instantly, reading only the parts you need, with a strategy that smart kernel programmers have spent decades optimizing? The answer is Vaex.

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](#).

2020 Unravel. All Rights Reserved. 2 Palo Alto Sq, Suite 120, Palo Alto, CA 94306