

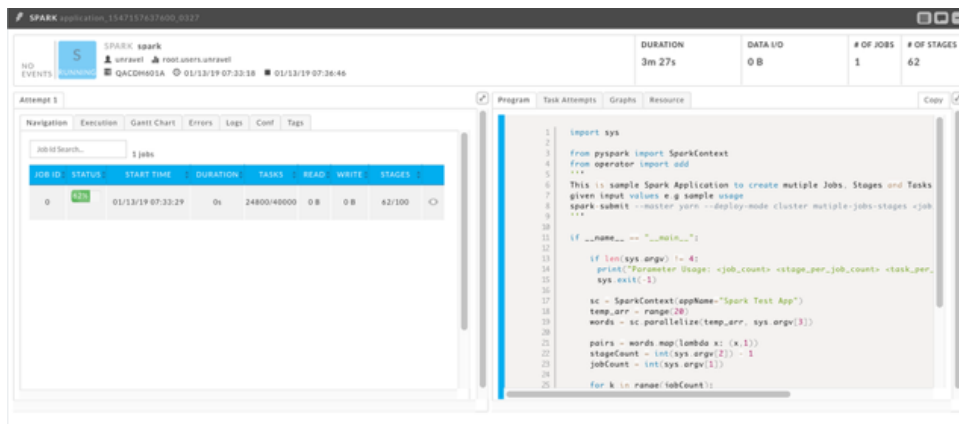
Unravel Newsletter: February 2019

We hope you have had a great start to the year so far. On our end, we have been busy and have some exciting news to share.

Unravel 4.5.0 is Available Now

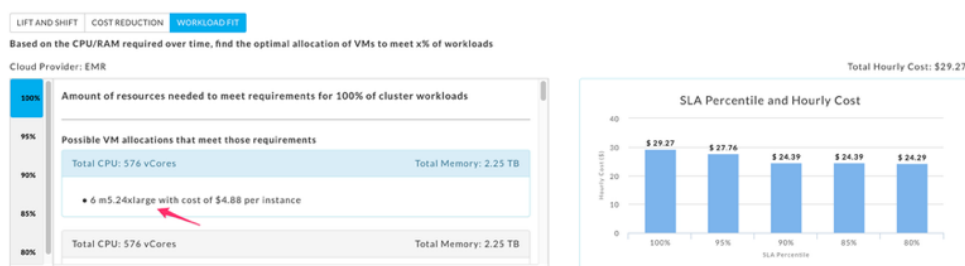
We got off to a running start this year with the release of Unravel 4.5.0 earlier this month. Some highlights of this release:

- **Applications**
 - Get *real-time* visibility into your *running* Spark jobs including *live* updates of their various stages.



Get real-time visibility into your running Spark jobs

- “Kill” or “move” your YARN applications via Unravel UI
- **Cloud Reports** – Looking to move your big data workloads to the cloud? Use Unravel’s Cloud Reports to help you in planning and managing this migration.



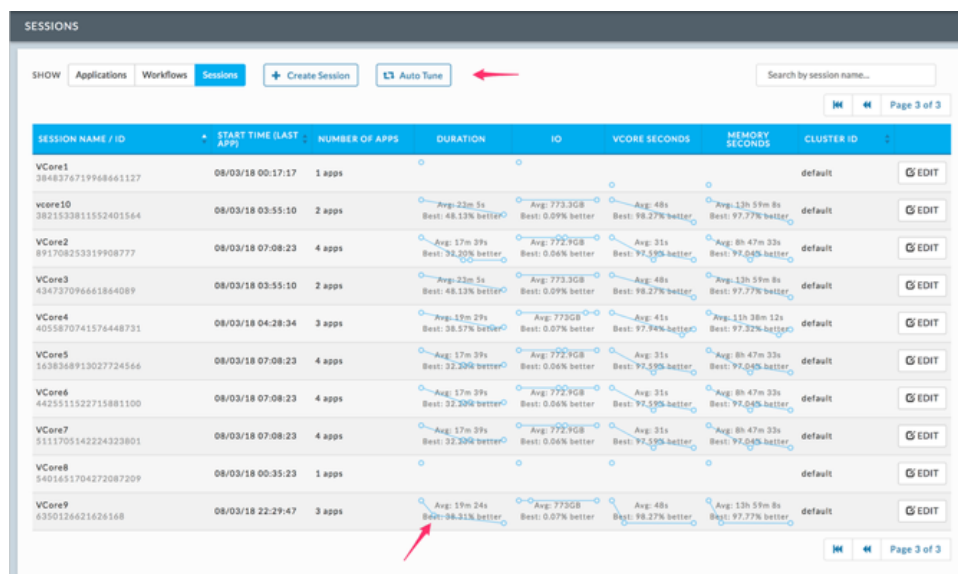
- **Insights** – New data, operator and SQL-level insights for Impala. New efficiency and failure events for Hive/Tez.
- **New API end points** – Retrieve app recommendations, Get finished and total no. of jobs that on a given date, Get status, error, logs, summary of individual apps.

In addition to the above new features, this release also includes several key enhancements and bug fixes, find more details in our latest [blog](#) and [release notes](#).

Don't forget to check out the awesome [Unravel User Guide](#) to help you get started in using the powerful new features and functionality!

Automated Tuning via “Sessions”

In case you missed it, back in September, the Unravel team added functionality for you to be able to auto-tune your applications (or get tuning recommendations). This feature called “Sessions”, allows you to run your application expressly and tune it for a specific goal e.g. meet a specific SLA or certain resource usage constraints.



The screenshot shows the 'SESSIONS' interface in Unravel. It features a navigation bar with 'Applications', 'Workflows', and 'Sessions' tabs. Below the tabs are buttons for '+ Create Session' and 'Auto Tune' (highlighted with a red arrow). A search bar is located on the right. The main content is a table with columns: SESSION NAME / ID, START TIME (LAST APP), NUMBER OF APPS, DURATION, IO, VCORE SECONDS, MEMORY SECONDS, and CLUSTER ID. Each row represents a session and includes performance metrics (Avg and Best) for various categories like CPU, I/O, and memory, along with an 'EDIT' button. A red arrow points to the 'Auto Tune' button, and another red arrow points to the 'Best' value for the 'DURATION' column in the 'VCORE9' session row.

SESSION NAME / ID	START TIME (LAST APP)	NUMBER OF APPS	DURATION	IO	VCORE SECONDS	MEMORY SECONDS	CLUSTER ID	
VCORE1 3848376719968661127	08/03/18 00:17:17	1 apps					default	EDIT
vcore10 382153811552401564	08/03/18 03:55:10	2 apps	Avg: 23m 5s Best: 48.13% better	Avg: 773.3GB Best: 0.09% better	Avg: 48s Best: 98.27% better	Avg: 13h 59m 8s Best: 97.77% better	default	EDIT
VCORE2 891708253319908777	08/03/18 07:08:23	4 apps	Avg: 17m 39s Best: 32.20% better	Avg: 772.9GB Best: 0.06% better	Avg: 31s Best: 97.19% better	Avg: 8h 47m 33s Best: 97.04% better	default	EDIT
VCORE3 434737096661864089	08/03/18 03:55:10	2 apps	Avg: 23m 5s Best: 48.13% better	Avg: 773.3GB Best: 0.09% better	Avg: 48s Best: 98.27% better	Avg: 13h 59m 8s Best: 97.77% better	default	EDIT
VCORE4 4055870741576448731	08/03/18 04:28:34	3 apps	Avg: 18m 29s Best: 58.57% better	Avg: 775GB Best: 0.07% better	Avg: 41s Best: 97.94% better	Avg: 11h 38m 12s Best: 97.32% better	default	EDIT
VCORE5 1638368913027724566	08/03/18 07:08:23	4 apps	Avg: 17m 39s Best: 32.20% better	Avg: 772.9GB Best: 0.06% better	Avg: 31s Best: 97.19% better	Avg: 8h 47m 33s Best: 97.04% better	default	EDIT
VCORE6 4425511522713881100	08/03/18 07:08:23	4 apps	Avg: 17m 39s Best: 32.20% better	Avg: 772.9GB Best: 0.06% better	Avg: 31s Best: 97.19% better	Avg: 8h 47m 33s Best: 97.04% better	default	EDIT
VCORE7 5111705142224323801	08/03/18 07:08:23	4 apps	Avg: 17m 39s Best: 32.20% better	Avg: 772.9GB Best: 0.06% better	Avg: 31s Best: 97.19% better	Avg: 8h 47m 33s Best: 97.04% better	default	EDIT
VCORE8 5401651704272087209	08/03/18 00:35:23	1 apps					default	EDIT
VCORE9 6350126421626168	08/03/18 22:29:47	3 apps	Avg: 19m 24s Best: 98.31% better	Avg: 775GB Best: 0.07% better	Avg: 48s Best: 98.27% better	Avg: 13h 59m 8s Best: 97.77% better	default	EDIT

Use "Sessions" to run your application expressly and tune it for a specific goal

Unravel has developed an automated model-driven approach utilizing Machine Learning to provide you with the best recommendations and if you choose to, even automatically tune your application on your behalf. For more information on what “Sessions” are and how you can use them to run your applications more efficiently and reliably, please click on the following:

[Learn more](#) about auto-tuning your applications using “Sessions”

Webinars

- [Adapting DevOps Workflows for Your Big Data Apps](#). Find out what you need to know about the challenges DevOps teams face when supporting the big data stack.
- [Doing DevOps for Big Data? What you need to know for AIOps](#). Explore with us the promise AIOps has to create hyper-efficiency within DevOps teams as they struggle with the diversity, complexity, and rate of change across the entire stack.

Upcoming Events

- Join us at the [#StrataConf](#) in SF on March 28 at 11:50 am PT to learn how to easily root cause Big Data Stack App failures with an automated technique that uses deep learning techniques. Details here: <https://oreil.ly/2RHsBhp>.
- Join us at the [#SparkAISummit](#) in SF for a session on April 25 at 11:50 AM PT where our CTO, Shivnath Babu, will discuss how organizations can leverage [#AI](#) chatbots to enhance their [#Spark](#) performance mgmt. Details here: <https://bit.ly/2PRW42t>.

Community Highlights

- ["Architecting Modern Data Platforms"](#) is a new book from O'Reilly. This practical book covers details on how to build an end-to-end big data infrastructure both on-premises and in the cloud and successfully architect a modern data platform.

Resources

- New Blog from Unravel CEO Kunal Agarwal: [Fighting for precious resources on a multi-tenant cluster](#), technical insights for fighting for precious resources on a multi-tenant big data cluster
- [Learn more](#) about performance management for big data to better understand, improve and control performance of your full big data stack
- [Online Product Demo](#)
- [Unravel Partner Ecosystem](#)
- [Unravel Product Releases and Documentation](#)
- [Unravel Datasheet](#)

- [More Unravel News](#)

[Sign Up for 30-day Trial](#) | [Contact Us](#)

© 2019 Unravel - All Rights Reserved
2 Palo Alto Square, Suite 120, Palo Alto CA 94306
[Privacy Policy](#)