

# On the Radar: Unravel provides APM for big data

---

Delivering an app-centric view of big data performance and utilization

Publication Date: 13 Oct 2017 | Product code: IT0014-003354

Tony Baer

---



## Summary

### Catalyst

An unexplored aspect of running and operationalizing big data clusters has been the overall performance of big data applications. While much of the spotlight has been on improving performance with new compute engines such as Spark, a number of other factors affect performance and reliability of big data applications. Unravel not only focuses on the compute engine or storage system, but also analyzes the overall performance of big data applications by tracking the end-to-end performance of all the system jobs, workflows, and pipelines in the cluster to isolate bottlenecks. Given that DevOps is used to monitor performance more from the point of view of the infrastructure, Unravel's challenge is educating its target audience on the need to take a more holistic applications-centric view across all components to maximize performance, reliability, and utilization.

### Key messages

- Unravel focuses on the entire big data stack, so data operations teams only need a single tool to analyze performance and troubleshoot problems.
- Unravel helps manage and optimize performance by using machine learning, AI (artificial intelligence), and predictive analytics to prioritize and automatically fix issues.
- As noted above, Unravel's challenge is educating its target market on the need to take an application-centric view to improve the performance of big data applications.

### Ovum view

Enterprises are continuing to scale their use of big data, which means there is more resource allocation, scheduling, and debugging to do, potentially making it harder to realize the potential benefits of big data systems. As one of the first tools on the market that can monitor and analyze the entire big data stack, Unravel is well placed to capitalize on the need for improving the performance of Hadoop clusters, but also faces a challenge for educating the market on how a holistic, APM-focused approach will help enterprises get greater value and faster results from their Hadoop investments.

## Recommendations for enterprises

### Why put Unravel on your radar?

Enterprises seeking to operationalize big data analytics could reap performance improvements by embracing Unravel's application-centric approach.

## Highlights

### Background

Unravel was founded in 2013 by a former Oracle executive and professor who was active in the open source Starfish project for providing self-tuning analytics for Hadoop. The cofounders developed what

became Unravel's tooling for taking an application-centric view of Hadoop performance while working on this project.

Unravel is an application performance management solution that looks at the whole big data stack. It provides a 360-degree, real-time view of all components of the big data ecosystem. It also correlates data from every layer of the stack to give a broad view of application performance, helping data teams optimize, troubleshoot, and analyze their big data systems. This means that instead of looking at several tools to pinpoint a problem, customers can do this with just one, and easily identify whether an issue is being caused by infrastructure, services, or applications.

Unravel has an operations dashboard, which provides an overview of the health and performance of an organization's entire big data stack. The system also reports on usage and access by users, applications, queues, and clusters. It provides reports that show how costs can be lowered by running applications more efficiently and allocating resources intelligently, freeing up cluster space by eliminating resource-wasting data and apps. It can also improve cluster utilization by managing pools and queues, increasing throughput by allocating the right size of pools and containers.

Unravel's overview of all applications and workflows is key to managing and improving performance. As problems can occur anywhere in the stack, there could be multiple potential reasons for failure; however, manually inspecting all the logs would be too time-consuming to be practical. Instead, Unravel uses machine learning, AI, and predictive analytics to optimize, troubleshoot, and analyze big data systems. It can prioritize and automatically fix problems – including performance bottlenecks and failures – helping to prevent new ones from happening. It can be configured to take actions automatically to solve recurring issues, keeping applications and clusters optimized. It offers "smart alerts," which understand an organization's cluster usage through machine learning and can accurately identify problems, sparing users from receiving a flood of incorrect alerts.

Unravel can manage storage too, monitoring files that have not been used in a while so that they can be removed to free up space. It can automatically suggest how to tier data sets based on how often they are used.

Unravel can be deployed on-premises, in the cloud, or a hybrid mix of both. While many early adopters have been using Unravel on-premises, the system is also effective in the cloud and can connect permanent data clusters with transient ones, correlating them all in one place. Unravel can capture logs from transient cloud clusters that would otherwise disappear, providing a useful trove of actionable information for improving performance in the cloud.

## Current position

Headquartered in Menlo Park, California, the company has received \$22m in three rounds of equity funding. The company has grown threefold over the past year, and customers include Sony, Autodesk, Box, Yellow Pages, banks, and healthcare firms. Unravel initially focused on support for Hadoop, and spent a year developing this. It then added support for Spark, and recently introduced support for Kafka to show the pipeline from end to end. Amazon EMR and Impala are also supported.

Unravel is now turning its attention to NoSQL systems and aims to launch support for NoSQL by the end of 2017. The company will look at HBase and Cassandra initially, as these are the biggest NoSQL databases that run alongside Hadoop. It will then turn its attention to other NoSQL databases.

Unravel sells direct to the market. It uses customer case studies, conferences, and word of mouth to raise awareness about the product. Its prime challenge is getting data-driven developers to embrace

the more holistic views that are common among application developers, with the rationale that as their analytics programs become operational, they must pay attention to service levels and performance.

## Data sheet

### Key facts

**Table 1: Data sheet: Unravel**

<b>Product name</b>	Unravel	<b>Product classification</b>	APM for Hadoop clusters
<b>Version number</b>	4.0	<b>Release date</b>	March 2017
<b>Industries covered</b>	All	<b>Geographies covered</b>	North America, EMEA
<b>Relevant company sizes</b>	Midsized to large	<b>Licensing options</b>	Subscription
<b>URL</b>	<a href="http://www.unraveldata.com">www.unraveldata.com</a>	<b>Routes to market</b>	Direct
<b>Company headquarters</b>	Menlo Park, California, US	<b>Number of employees</b>	50

Source: Ovum

## Appendix

### On the Radar

On the Radar is a series of research notes about vendors bringing innovative ideas, products, or business models to their markets. Although On the Radar vendors may not be ready for prime time, they bear watching for their potential impact on markets and could be suitable for certain enterprise and public sector IT organizations.

### Further reading

*On the Radar: Wavefront by VMware offers cloud-native APM*, IT0022-001093 (October 2017)

*Market Radar: Cloud-native Application Performance Management*, IT0014-003329(September 2017)

### Authors

Tony Baer, Principal Analyst, Information Management

[tony.baer@ovum.com](mailto:tony.baer@ovum.com)

Beck Thompson, Associate Analyst

## Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at [consulting@ovum.com](mailto:consulting@ovum.com).

## Copyright notice and disclaimer

The contents of this product are protected by international copyright laws, database rights and other intellectual property rights. The owner of these rights is Informa Telecoms and Media Limited, our affiliates or other third party licensors. All product and company names and logos contained within or appearing on this product are the trademarks, service marks or trading names of their respective owners, including Informa Telecoms and Media Limited. This product may not be copied, reproduced, distributed or transmitted in any form or by any means without the prior permission of Informa Telecoms and Media Limited.

Whilst reasonable efforts have been made to ensure that the information and content of this product was correct as at the date of first publication, neither Informa Telecoms and Media Limited nor any person engaged or employed by Informa Telecoms and Media Limited accepts any liability for any errors, omissions or other inaccuracies. Readers should independently verify any facts and figures as no liability can be accepted in this regard – readers assume full responsibility and risk accordingly for their use of such information and content.

Any views and/or opinions expressed in this product by individual authors or contributors are their personal views and/or opinions and do not necessarily reflect the views and/or opinions of Informa Telecoms and Media Limited.

## CONTACT US

[www.ovum.com](http://www.ovum.com)

[analystsupport@ovum.com](mailto:analystsupport@ovum.com)

## INTERNATIONAL OFFICES

Beijing

Dubai

Hong Kong

Hyderabad

Johannesburg

London

Melbourne

New York

San Francisco

Sao Paulo

Tokyo

