



Market Insight Report Reprint

Yellowbrick adopts Kubernetes, adds control plane for data warehousing in the distributed cloud era

April 26 2021

by **Matt Aslett**

With the recent addition of its Yellowbrick Manager and support for Kubernetes, Yellowbrick is positioning itself to address the growing requirements for data warehousing that is distributed across on-premises datacenters and multiple clouds.

451 Research

S&P Global

Market Intelligence

This report, licensed to Yellowbrick, developed and as provided by 451 Research, LLC, was published as part of our syndicated Market Insight subscription service. It shall be owned in its entirety by 451 Research, LLC. This report is solely intended for use by the recipient and may not be reproduced or re-posted, in whole or in part, by the recipient without express permission from 451 Research.

Introduction

We recently noted that data warehousing software and appliance provider Yellowbrick had adapted its offering to reflect the requirements for hybrid- and multicloud data warehousing by providing a consistent experience on-premises and in multiple clouds. Recently the company went one step further, delivering a preview of its Yellowbrick Manager functionality, which is designed to provide a unified control plane for managing data warehouse deployments distributed across multiple clouds and datacenters. The company has also adopted cloud-native architecture via Kubernetes to support that distributed cloud approach and has also delivered the latest version of its optimized hardware instances for private cloud. Known as Andromeda, the second generation of Yellowbrick's all-flash optimized architecture is said to deliver 3x performance improvements.

THE 451 TAKE

At the beginning of this year we outlined the case for adopting an infrastructure-agnostic data platform layer that can run on multiple clouds as well as on-premises datacenters. The ability to use the same database in multiple locations is already addressed by most database providers today but managing data across those multiple locations is easier said than done. Yellowbrick will be aiming to position the functionality being delivered by Yellowbrick Manager as a differentiator, alongside its Andromeda architecture and Kubernetes support, to help enterprises as they look to evolve their investments in hybrid IT and manage data consistently across distributed clouds that span private cloud datacenters, public cloud services and edge devices.

Details

Despite much consolidation over the years, the data warehousing sector remains crowded, with established incumbents jostling with emerging startups for attention. Providing differentiation is key to gaining traction. Yellowbrick, which was founded in 2014 and emerged from stealth four years later, has a couple of factors it can point to: a high performance all-flash data-warehousing appliance for on-premises, and availability as a service on the public cloud with asynchronous data replication to support hybrid IT and multicloud requirements. The company recently announced updates to Yellowbrick Data Warehouse that deliver enhancements related to both these factors, including the general availability of its Andromeda second-generation optimized hardware architecture, support for Kubernetes and a preview of Yellowbrick Manager for managing data across a distributed cloud and datacenter architecture.

Adopting Kubernetes as the core component of its cloud-native architecture better enables Yellowbrick to address portability across multiple environments – including running on Kubernetes in a customer's own datacenter, as well as public cloud Kubernetes offerings, and also on edge architecture. Kubernetes is also being built into the Andromeda optimized instances later this year. In the interim, Andromeda incorporates new 'Kalidah' scan accelerators for data scan rate improvements and also delivers network performance improvements, and improved linear scalability (up to 6PB across 40 nodes). The Andromeda instances are primarily positioned for on-premises private clouds but can also be managed as a service via AWS PrivateLink and Azure Private Link (with plans to do the same on Google Cloud Platform when an equivalent is generally available).


With its 5.2 release, Yellowbrick has also delivered the general availability of improved support for data lake architecture integration – specifically enabling data to be loaded into Yellowbrick Data Warehouse from tables in cloud object stores, such as Amazon S3, Azure Blob Storage and Google Cloud Storage, as well as MinIO object stores and the Hadoop Distributed File System. The loading of data from these object stores can be performed via SQL integration tools, or orchestrated via Yellowbrick Manager, which is now in preview and due to be generally available before the end of the year. Yellowbrick Manager is a significant addition to Yellowbrick Data Warehouse in that it is designed to provide a unified control plane to enable users to manage data distributed across multiple locations – including private cloud, public cloud or edge devices.

Initially aimed at data engineers and developers, with additional management console functionality to come aimed at database administrators, Yellowbrick Manager provides functionality to enable users to view the configuration and connection details of multiple Yellowbrick Data warehouse instances, as well as load and move data, and explore data and create queries with a built-in SQL editor. Yellowbrick Manager is designed to provide a unified management experience that complements the company's existing data loading and replication functionality, including the previously mentioned connectors for streaming or batch bulk inserts using Apache Spark and Apache Kafka, as well as asynchronous data replication across instances with support for read-only secondary instances.

While many operational and analytic database vendors offer their products for deployment on-premises as well as consumption via public clouds, 451 Research believes that the ability to manage data across these environments is becoming increasingly critical. As the figure below illustrates, more than three quarters (78%) of respondents to our Voice of the Enterprise: Data & Analytics, Data Platforms, 2021 agreed that the ability to run the same database in a hybrid IT environment is an important consideration for their organization when selecting a new data platform.

Importance of Database Support For Hybrid IT

The ability to run the same database in a hybrid IT environment (both on-premises and public cloud) is an important consideration for my organization when selecting a new data platform. n = 287



Source: 451 Research, Voice of the Enterprise: Data & Analytics, Data Platforms, 2021

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2021 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its Web sites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.