

May 13, 2024



BLOG POST

Oracle Ushers in the AI Era With Oracle Database 23ai

How Oracle Ensures the Strategic Value of Oracle Database for Its Customers



Holger Mueller
VP & PRINCIPAL ANALYST



Produced exclusively for Constellation Research clients

INTRODUCTION

A new era in computing has begun with the wide adoption of generative artificial intelligence (GenAI). CxOs know that they need to have an AI strategy more than ever in 2024. There is a lot of uncertainty concerning how to adopt AI, but there is certainty about one thing: Data determines the availability and quality of AI. With petabytes of their critical transactional data stored in Oracle Database, CxOs want to know how they can leverage that data.

On May 2, 2024, Oracle unveiled Oracle Database 23ai at a LinkedIn live event featuring Oracle Founder, Chairman, and CTO Larry Ellison and Executive Vice President of Mission-Critical Database Technologies Juan Loiza. This blog post discusses the key innovations for enterprises and concludes with recommendations for CxOs on how to achieve what matters most today: Enterprise Acceleration.¹

KEY INNOVATIONS IN ORACLE DATABASE SUPPORT AI REQUIREMENTS

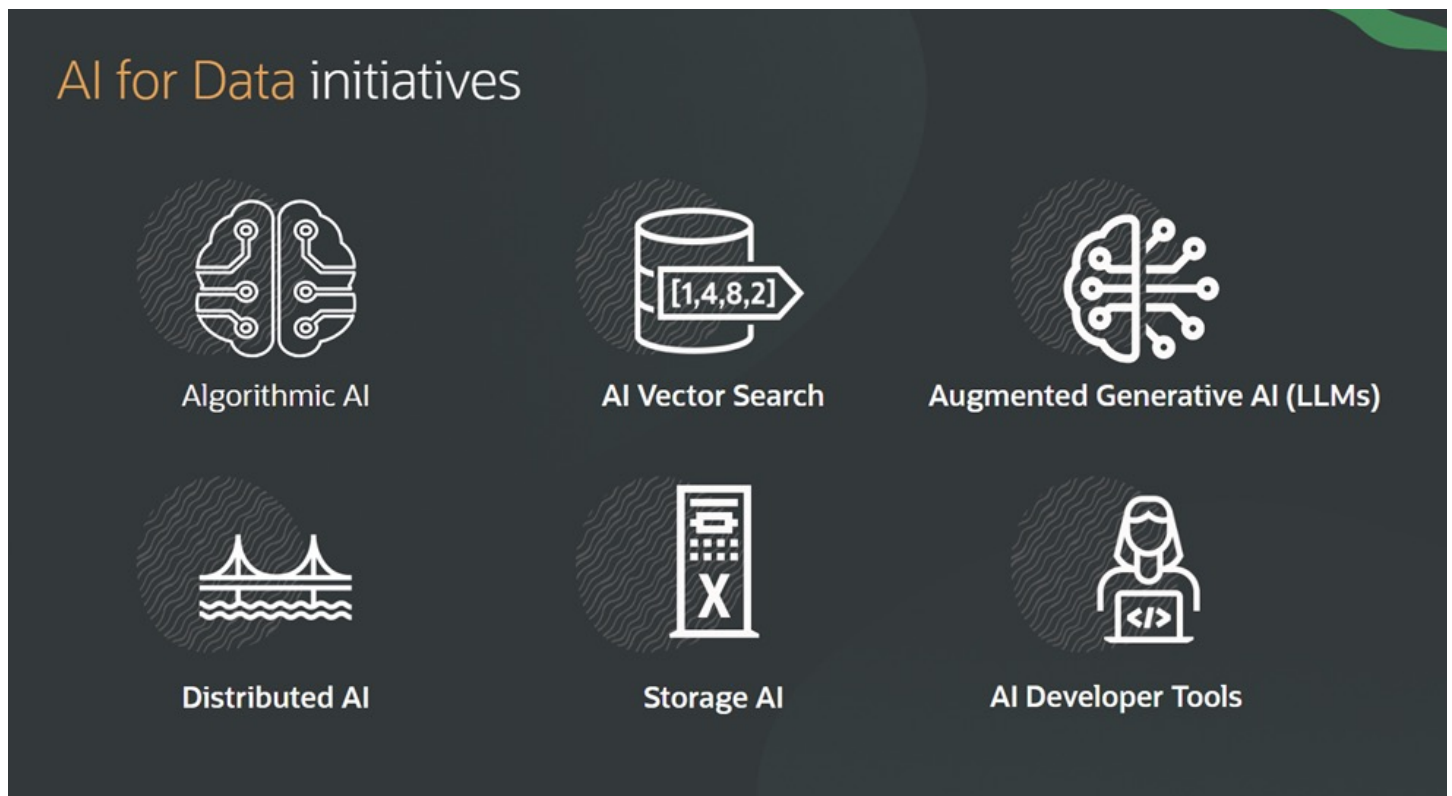
Oracle focused its new Oracle Database 23ai capabilities on six distinct functional pillars (see Figure 1), four of which stand out for discussion here.

1. Bringing AI to the Data

Traditionally, data scientists have moved data into modelling environments to provide traditional analytical services such as predictive analytics. Not much changed with machine learning, and the same thinking has been applied more recently to AI projects. The dilemma is that moving data is cumbersome, expensive, and potentially even a data security risk. It is smarter to bring the AI to the data, and that is what Oracle is doing with Oracle Database 23ai.

With the help of PL/SQL, R, or Python, developers can run AI algorithms directly in the database, enabling “traditional” use cases such as product recommendations, customer churn, fraud detection,

Figure 1. Six Key AI Functional Areas in Oracle Database 23ai



Source: Oracle

and many more. This is important, because in reality not all AI use cases that move the needle in the enterprise toward Enterprise Acceleration are powered by large language models (LLMs) or GenAI.

2. Infusing AI With the Database via Vector Support and Search

With the introduction of a new “vector” data type, Oracle is making nonstructured information available for AI. But it is not enough to just create and store vectors; they also need to be searchable, a capability Oracle has added, called AI Vector Search, with Oracle Database 23ai. Oracle also brings in the full capabilities of the overall Oracle Database stack to vectors, with the Oracle GoldenGate 23ai product able to replicate vectors across different locations. The result is reuse of previously created vectors, saving on computing, consistency, and better and faster ways to leverage AI in the enterprise.

The main value, however, is that CxOs can use AI in the same place where their enterprise's critical transactional information resides: Oracle Database. Transactional information is hugely relevant for AI, because all key information of an enterprise is stored in transactions. The Oracle Database AI vector capabilities bridge the gap between transactional data and vectorized data, and with that allow enterprises to have better, more-relevant AI faster and at a lower cost, while enabling the critical grounding technique of retrieval-augmented generation (RAG).

3. Delivering Next-Generation Storage for AI

That AI is all about data is a well-known and accepted fact today. But data ultimately resides in storage, and the capabilities of the storage layer also determine the success of the AI generated from the data residing in the storage. Oracle is already familiar with this challenge, because transactional database performance also relies heavily on storage. Nonetheless, Oracle has revved up the underlying Exadata software with Exadata System Software 24ai. This adds support for vectors by allowing enterprises to offload vector processing to Smart Exadata Storage, parallelizing vector search and enabling faster performance for AI.

The speed and efficiency for automating AI-powered next-generation applications² is critical for enterprises, because speed means better customer and employee experience, better automation, and lower cost—all qualities CxOs want and need for their enterprises' success.

4. Running All Data Everywhere for AI

Enterprise data is distributed across multiple clouds and applications, on-premises, and at the edge. It is key to bring that data together to enable the best and most accurate and applicable AI automation, and for that, data needs to be brought together at least for creating AI vectors. And these vectors need to be generated with one single algorithm to ensure consistency, relevancy, and compatibility.

The Oracle solution to this challenge is of course Oracle GoldenGate 23ai, which brings together data from other clouds and applications and can feed that to Oracle Database 23ai instances to create

consistent AI vectors. Oracle GoldenGate 23ai can also propagate out the generated AI vectors for consistent uptake across the locations. This is of massive value for CxOs whose enterprises operate in the hybrid cloud and need to ensure consistent AI model and vector uptake across a large variety of deployments.

OTHER INNOVATIONS DELIVER WHAT CUSTOMERS NEED TODAY

With all the AI hoopla, it is easy to forget that Oracle Database is used day in and day out to power mission-critical workloads. Enterprises are expecting Oracle to provide more than just AI-centric innovation, and Oracle has delivered that with Oracle Database 23ai.

More specifically, Oracle is delivering the following distinctive capabilities (see Figure 2):

Figure 2. Oracle Database 23ai Overview



The image shows a presentation slide for Oracle Database 23ai. The slide has a dark background with white and orange text. On the right side, there is a laptop graphic with the Oracle Database 23ai logo on its screen. The logo consists of the word 'ORACLE' in red, 'Database' in white, and '23ai' in large white numbers. The slide text on the left lists the following points:

- Next-Generation Converged Database
- Available now
- Over 300 new features
- Thousands of enhancements
- The latest long-term support release of Oracle's flagship Database
- Existing customers get free upgrade

Source: Oracle

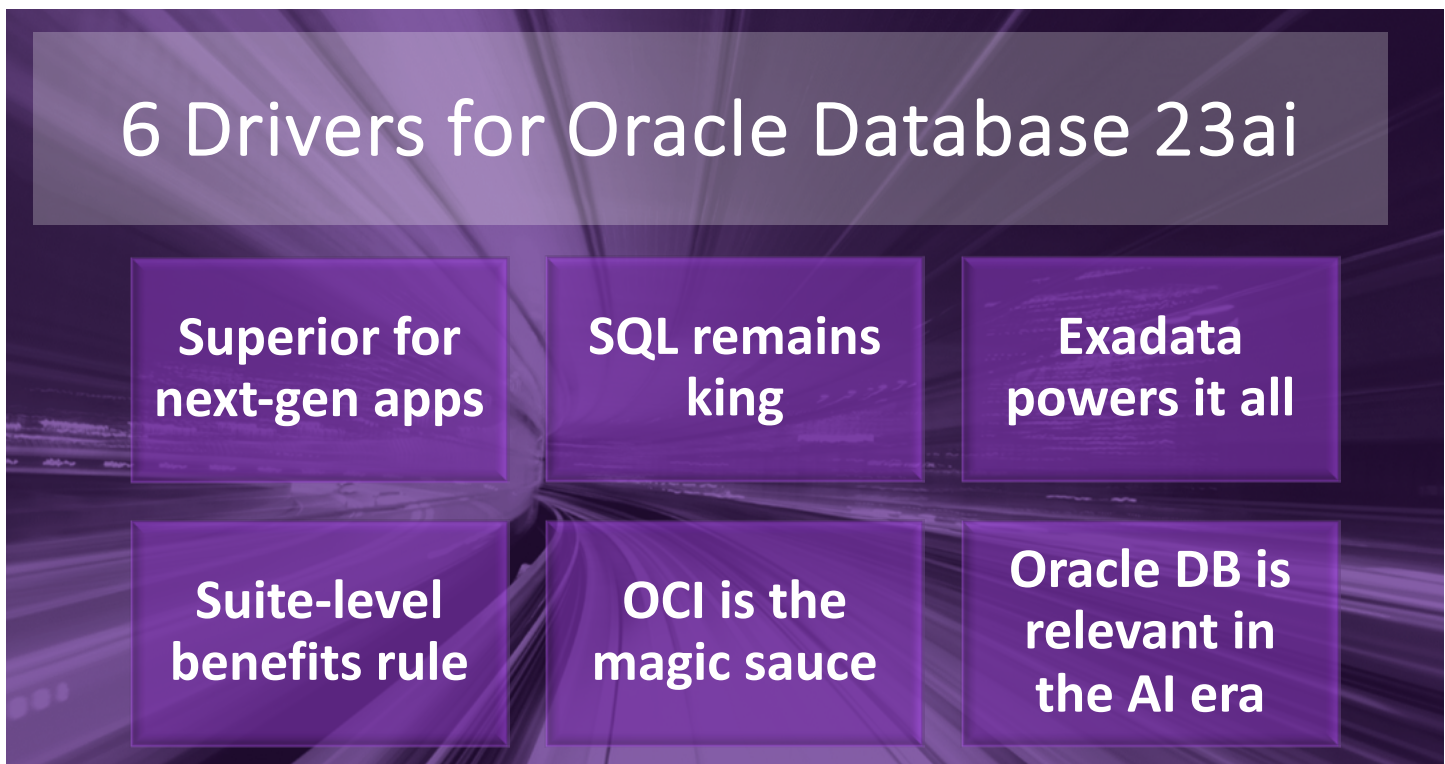
- **More than 300 new features shipped.** With more than 300 new features shipping with Oracle Database 23ai, there is a lot of innovation being delivered. That is value for enterprises, which can continue to rely on innovation being delivered for Oracle Database.
- **Thousands of enhancements delivered.** In a strong manifestation of Oracle listening to its customers, who can log enhancement requests, the vendor has delivered thousands of enhancements with the Oracle Database 23ai release.
- **Providing the new long-term support release of Oracle Database.** CxOs need stability to run their mission-critical workloads, so the length of support is key to how long they can achieve versional stability for their existing next-generation applications powered by Oracle Database. With Oracle Database 23ai being the new long-term support release from Oracle, there is more peace of mind for those running Oracle Database.
- **Free upgrades for existing customers.** In an unusual departure from typical enterprise practices, Oracle makes available new capabilities for free to its customers under annual support contracts. This is of tremendous value for CxOs, making it easier for them to justify the adoption and operation of Oracle Database via a massively improved total cost of ownership.

WHY DOES IT MATTER?

Oracle Database 23ai is a critical release for enterprises. Here are the primary reasons it really matters (see Figure 3):

- **Superior for AI-powered next-generation applications.** The GenAI use cases that are popular right now are focused on the document-centric automation areas. Relevant information is stored in transactional databases, and the Oracle approach with AI vectors brings both together.
- **SQL remains king.** The world does not have enough AI specialists, does not have enough data scientists, and does not have enough developers. But more people than developers know SQL, and

Figure 3. Six Value Drivers for Oracle Database 23ai



Source: Constellation Research

with Oracle’s approach to AI, SQL can be used to build modern AI apps, thus democratizing access and allowing enterprises to build more AI applications sooner, faster, and cheaper.

- **Exadata powers it all.** Hardware matters again, especially in the era of AI. The Exadata design point of feeding data quickly to efficient CPUs fits perfectly into the GenAI design sweet spot. This has made it relatively easy for Oracle vis-à-vis its competitors to leverage Exadata as the foundation for Oracle Database 23ai.
- **Suite-level benefits rule.** Suites have won in enterprise software, because CxOs need to avoid becoming system integrators between different automation islands. With the AI vector capabilities, Oracle has added a new database type to its universal database.³
- **OCI is the magic sauce.** AI is powered by data, and data is regulated. Local presence matters for data residency and privacy—and Oracle Cloud Infrastructure (OCI) enables this. With Oracle

Globally Distributed Database and Oracle GoldenGate–based replication, enterprises can build AI-powered next-generation applications and stay compliant.

- **Oracle Database extends its enterprise relevancy for the AI era.** Enterprises are looking for their AI data platform. Cloud-powered lakehouses are a viable option, but with Oracle Database 23ai, Oracle is making a strong argument to keep data where it is in Oracle Database and power AI enterprisewide efforts from there.

ADVICE FOR CXOS

Constellation has the following recommendations for CxOs considering Oracle Database 23ai:

- 1. Understand current and future AI workloads of your enterprise.** Don't get caught in the hype of document-centric AI use cases. These are important but will not move the needle for an enterprise. The combination of GenAI and transactional data, on the other hand, *will* move the needle, so look at these future use cases and workloads. That requires considering the to-be AI operational setup that will power your enterprise for at least the next decade.
- 2. If you are an Oracle Database customer, upgrade.** With Oracle making the upgrade to Oracle Database 23ai free for existing customers, upgrading is a no-brainer. Even if it is for pilot, training, or proof-of-concept projects, it is key to get going. With 23ai being a supported version long-term, it makes that move even easier to justify. Plus, Oracle offers free versions for developers to build applications.
- 3. If you are an Oracle prospect, evaluate.** There are fewer and fewer high-performance transactional databases for mission-critical applications. The 2020s are the first decade since the 1980s where there are fewer than a handful of transactional mainstream databases available. Transactional workloads will not go away, so non-Oracle customers cannot ignore Oracle's relative strength in the market. This means CxOs should evaluate the new database version—even if they have concerns regarding Oracle. Bits and bytes don't know about past experiences but power today's use cases.

- 4. Move to the cloud as soon as possible.** AI lives in the cloud. AI projects need to run in the cloud, ideally where data resides, because data gravity is real. The sooner your enterprise moves to the cloud, the more AI-ready it is. CxOs cannot dawdle in regard to AI.
- 5. Practice commercial prudence.** As always, CxOs need to practice commercial prudence when it comes to platform decisions. One-time costs, ongoing costs, capex versus opex, and lock-in effects are the key areas to consider before making platform decisions. Database platform decisions are no exception to the need for commercial prudence in all phases of purchase, adoption, and usage cycle.

MY POV

Enterprises need to formulate their AI strategy and execute it as soon as possible. We are now seeing the responses in products from pre-AI era vendors—such as Oracle with Oracle Database 23ai. Oracle has a long practice of adding suffixes to its database release numbers: It used to be an *i* (internet), then became a *c* (cloud), and now it is, of course, *ai*. But product naming apart, Oracle's approach of adding vector search and storage into their converged database engine is not only relevant and valid but also extends the relevancy of Oracle Database for the next decade and beyond. Furthermore, it allows enterprises to keep their mission-critical transactional processing in the same place, which is a huge relief. The power of the Exadata platform in combination with the OCI architecture and presence makes Oracle Database even more attractive and relevant for CxOs.

In the long run, the verdict is still out on how the AI data platform will be powered. Right now, it is the cloud-based lakehouse versus the AI vector-powered transactional databases. From an Oracle Database perspective, Oracle is in the latter camp and making a strong innovation pitch to enterprises. The good news for CxOs is that Oracle can play it both ways by including a potential future lakehouse offering in OCI. This gives CxOs the ultimate peace of mind in making this key platform decision, because Oracle can emerge as a winner in both outcomes. It is a good time to be an Oracle customer in general and an Oracle Database customer more specifically.

RELATED RESEARCH - ORACLE-RELATED

For the Constellation ShortList™ for next-generation computing platforms, see: Holger Mueller, “Constellation ShortList Next-Generation Computing Platforms,” February 9, 2022. <https://www.constellationr.com/research/constellation-shortlist-next-generation-computing-platforms-3>

For the Constellation ShortList on RDBMS on-premises, see: Holger Mueller, “Constellation ShortList Next-Gen Databases: RDBMS for On-Premises,” Constellation Research, February 9, 2022. <https://www.constellationr.com/research/constellation-shortlist-next-gen-databases-rdbms-premises-5>

For more on Oracle Exadata, see: Holger Mueller, “Oracle Exadata X9M and Oracle Exadata Cloud@Customer X9M Enable Rapid Enterprise Acceleration,” Constellation Research, September 29, 2021. <https://www.constellationr.com/research/oracle-exadata-x9m-and-oracle-exadata-cloudcustomer-x9m-enable-rapid-enterprise>

For more on the Oracle MySQL introduction, see: Holger Mueller, “There Is a HeatWave for Databases: Oracle MySQL Database Service Merges OLTP and OLAP,” Constellation Research, June 3, 2021. <https://www.constellationr.com/research/there-heatwave-databases-oracle-mysql-database-service-merges-oltp-and-olap>

For more on Oracle MySQL, see: Holger Mueller, “Oracle Switches to Autopilot, Turns Up the Heat With MySQL Database Service,” Constellation Research, September 16, 2021. <https://www.constellationr.com/research/oracle-switches-autopilot-turns-heat-mysql-database-service>

For information on Oracle Application Express (Oracle APEX), see: Holger Mueller, “Oracle Makes Its Great Database Even Better—and Adds Low Code to It,” Constellation Research, February 10, 2021. <https://www.constellationr.com/research/oracle-makes-its-great-database-even-better-and-adds-low-code-it>

For information on logging and observability, see: Holger Mueller, “Oracle Brings Holistic Observability to the Enterprise,” Constellation Research, November 19, 2020. <https://www.constellationr.com/research/oracle-brings-holistic-observability-enterprise>

For information on Oracle Exadata Cloud Service, see: Holger Mueller, “Oracle Exadata Cloud Service X8M Brings Leading Elasticity to Enterprise Needs,” Constellation Research, October 19, 2020. <https://www.constellationr.com/research/oracle-exadata-cloud-service-x8m-brings-leading-elasticity-enterprise-needs>

For information on Oracle in DevOps, see: Holger Mueller, “Oracle Boosts DevOps Velocity With Oracle Linux Automation Manager,” Constellation Research, December 13, 2021. <https://www.constellationr.com/research/oracle-boosts-devops-velocity-oracle-linux-automation-manager>

RELATED RESEARCH CONTINUED

For information on Oracle's new JSON capabilities, see: Holger Mueller, "New Offering Launch—Oracle Boosts Developer Velocity With Autonomous JSON Database," Constellation Research, August 13, 2020. <https://www.constellationr.com/blog-news/new-offering-launch-oracle-boosts-developer-velocity-autonomous-json-database>

For information on Oracle Dedicated Region and Oracle Autonomous Database, see: Holger Mueller, "Oracle Raises the Stakes in the Next-Gen Compute Platform Market," Constellation Research, July 22, 2020. <https://www.constellationr.com/research/oracle-raises-stakes-next-gen-compute-platform-market>

For more on Oracle Database at Deutsche Bank, see: Holger Mueller, "Oracle Affirms Its Strategic Role in Banking," Constellation Research, July 20, 2021. <https://www.constellationr.com/research/oracle-affirms-its-strategic-role-banking>

ENDNOTES

¹ Mueller uses the term Enterprise Acceleration to refer to the need for enterprises to move faster and become more agile. For more, see: Holger Mueller, “Why the C-Suite Must Embrace Enterprise Acceleration,” May 2, 2019. <https://www.constellationr.com/research/why-c-suite-must-embrace-enterprise-acceleration>

² Mueller uses the term next-generation applications to describe applications that use a combination of AI, big data, cloud, and build applications across seven distinct use cases. Find more on next-generation applications here: Holger Mueller, “The Era of Infinite Computing Triggers Next-Generation Applications,” June 1, 2018. <https://www.constellationr.com/research/era-infinite-computing-triggers-next-generation-applications>

³ Mueller defines a universal database as one integrated database for multiple database use cases, effectively becoming the one and only database for all enterprise needs. The alternative is the suite of specialized databases. Oracle refers to its database as a converged database. For more, see: Holger Mueller, “The Universal Database Versus a Suite of Specialized Databases,” April 5, 2021. <https://www.constellationr.com/research/universal-database-versus-suite-specialized-databases>

ANALYST BIO

Holger Mueller

Vice President and Principal Analyst

Holger Mueller is a vice president and principal analyst at Constellation Research. He provides guidance for the fundamental enablers of the cloud, IaaS, and PaaS, with forays up the tech stack into big data, analytics, and SaaS. Mueller provides strategy and counsel to key clients, including chief information officers, chief technology officers, chief product officers, investment analysts, venture capitalists, sell-side firms, and technology buyers.

Prior to joining Constellation Research, Mueller was VP of products for NorthgateArinso, a KKR company. He led the transformation of products to the cloud and laid the foundation for new business-process-as-a-service (BPaaS) capabilities. Previously he was the chief application architect with SAP and was also VP of products for FICO. Before that he worked for Oracle in various management functions—on both the application development (CRM, Fusion) and business development sides. Mueller started his career with Kiefer & Veitinger, which he helped grow from a startup to Europe's largest CRM vendor from 1995 onward. Mueller has a Diplom-Kaufmann degree from the University of Mannheim, with a focus on information science, marketing, international management, and chemical technology. A native European, Mueller speaks six languages.

 [@holgermu](https://twitter.com/holgermu)  constellationr.com/users/holger-mueller  linkedin.com/in/holgermueller

ABOUT CONSTELLATION RESEARCH

Constellation Research is an award-winning, Silicon Valley–based research and advisory firm that helps organizations navigate the challenges of digital disruption through business model transformation and the judicious application of disruptive technologies. Unlike the legacy analyst firms, Constellation Research is disrupting how research is accessed, what topics are covered, and how clients can partner with a research firm to achieve success. Over 350 clients have joined from an ecosystem of buyers, partners, solution providers, C-suite, boards of directors, and vendor clients. Our mission is to identify, validate, and share insights with our clients.

Organizational Highlights

- Named Institute of Industry Analyst Relations (IIAR) New Analyst Firm of the Year in 2011 and #1 Independent Analyst Firm for 2014 and 2015.
- Experienced research team with an average of 25 years of practitioner, management, and industry experience.
- Organizers of the Constellation Connected Enterprise—an innovation summit and best practices knowledge-sharing retreat for business leaders.
- Founders of Constellation Executive Network, a membership organization for digital leaders seeking to learn from market leaders and fast followers.

 www.ConstellationR.com

 [@ConstellationR](https://twitter.com/ConstellationR)

 info@ConstellationR.com

 sales@ConstellationR.com

Unauthorized reproduction or distribution in whole or in part in any form, including photocopying, faxing, image scanning, emailing, digitization, or making available for electronic downloading is prohibited without written permission from Constellation Research Inc. Prior to photocopying, scanning, and digitizing items for internal or personal use, please contact Constellation Research Inc. All trade names, trademarks, or registered trademarks are trade names, trademarks, or registered trademarks of their respective owners.

Information contained in this publication has been compiled from sources believed to be reliable, but the accuracy of this information is not guaranteed. Constellation Research Inc. disclaims all warranties and conditions with regard to the content, express or implied, including warranties of merchantability and fitness for a particular purpose, nor assumes any legal liability for the accuracy, completeness, or usefulness of any information contained herein. Any reference to a commercial product, process, or service does not imply or constitute an endorsement of the same by Constellation Research Inc.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold or distributed with the understanding that Constellation Research Inc. is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional person should be sought. Constellation Research Inc. assumes no liability for how this information is used or applied nor makes any express warranties on outcomes. (Modified from the Declaration of Principles jointly adopted by the American Bar Association and a committee of publishers and associations.)

Your trust is important to us, and as such, we believe in being open and transparent about our financial relationships. With our clients' permission, we publish their names on our website.

San Francisco Bay Area | Boston | Colorado Springs | Denver | Ft. Lauderdale | New York Metro
Northern Virginia | Portland | Pune | San Diego | Sydney | Washington, D.C.