December 17, 2025

CR PULSE

Oracle Redefines the AI and Data Relationship With Oracle AI Database 26ai

The Architecture for AI Is Different When Your Customers Already Run the Most Scalable RDBMS in the Industry—and It Is Very Compelling

Holger Mueller
VICE PRESIDENT AND PRINCIPAL ANALYST

constellation
RESEARCH

Enterprises run on transactions, and this means that transactional databases are critical for enterprises. The ability to seamlessly scale transactional databases has always been a core challenge for vendors, and subsequently for their enterprise users. At the core of critical transactional systems are relational databases, which CxOs must plan to run for the foreseeable future—at least another decade to come—on-premises and/ or in the cloud.

At the same time, offerings must run on a global scale and in multiple clouds, and CxOs will try to consolidate the same or similar solutions across fewer database installations. Next-generation application[1] needs will be addressed in local cloud regions or on-premises for reasons including data residency, data sovereignty, compliance, and performance requirements. Consolidation delivers consistency, complexity reduction, licensing negotiation upsides, and more—all delivering a lower total cost of ownership (TCO), critical in times of perennial pressure on IT budgets.

With artificial intelligence (AI) being omnipresent since late 2023, database vendors have reacted, with almost all of them providing vector support for their offerings. At the same time, enterprises need to be able to run the workloads of up-and-coming agentic AI applications—often on an infrastructure/platform that cannot scale to the additional performance demands of AI-powered next-generation applications. This makes CxOs look for offerings that can combine both transactional database and vector scale, ideally in a single offering.

TRENDS

Customer checks from the BT150[2] and AI150[3] alumni show the following eight next-generation database trends for 2025 and beyond:

Scaling for AI application success is the No. 1 objective. AI-based applications need to leverage transactional data sources, and vectors have been added across the vendor landscape. But new agentic AI–powered applications will also create higher volumes of transactions than humans could, and database platforms need to scale for that as well as adjust to unpredictable workload spikes.

Scaling has a new dimension: the vector database. Traditional transactional database scalability has been a challenge for vendors from the start of the category, but this has been largely addressed and solved to different degrees. With vector databases emerging and having to operate in parallel with transactional databases, the scalability of the vector database becomes a second critical scalability dimension for enterprises. Alternatively, a single highly scalable universal database can run transactional and vector databases simultaneously.

Lowering complexity is critical. Database vendors need to solve the challenge of adding more and more capabilities but also need to reduce the complexity of their database offering. Automation and AI processes are critical for this, because humans can no longer scale at the speed and capacity needed for modern database operations. Practical AI operational use cases are table stakes going into 2026.

Solving cloud lock-in is within reach. Cloud lock-in has always been something CxOs want to avoid but rarely manage to. Being able to use a single database across a hybrid deployment model, public cloud and on-premises, with as close to real-time data synchronization as possible, is a key criterion for successful transactional database offerings. The same now applies to vector databases, which need to be offered across clouds and on-premises. Alternatively, a single universal database addresses real-time synchronization between transactional and vector databases and in a variety of deployment models.

Data residency requirements need to be addressed. Global single instances are good news for operational cost and insights but cannot satisfy data residency and compliance requirements. CxOs currently choose the location for their database in the geography where they expect the lowest violation fine. That practice cannot last forever, and vendors need to make it easier for CxOs to address data residency requirements.

Hybrid operation remains important for the near future. Data will remain on-premises for performance, statutory, and connectivity reasons for at least the next decade. CxOs want and need to run the same offerings both in public cloud and on-premises, with a strong preference for a single pane of glass and similar performance and governance across any deployment model.

High availability (HA) matters more than ever. With agentic AI–powered apps never going to sleep, always-on databases able to scale instantly to the workloads thrown at them are more important than ever. CxOs want to see manageable automated processes for scaling their enterprise workloads as needed, with little to no overhead on the people and license side.

Lower TCO is the true north. IT costs need to come down, and as much as vendors innovate and rightfully charge an innovation premium, TCO matters, and CxOs will look only at database offerings that lower their TCO over time.

Threshold Criteria

Constellation considers the following criteria for vendors to be considered in the category of next-generation databases. Most vendors must meet five out of the six criteria:

☑ **More than 25,000 production installations of critical applications**

☑ **Customers on five or more continents**

☑ **Vendor directly present with sales and service on five or more continents**

☑ **Public cloud deployment option**

☑ **Same database deployed on-premises and on public cloud (by vendor and/ or partner)**

☑ **New AI-triggered capabilities**

constellation
R E S E A R C H

VENDOR UPDATE

Oracle is one of the key vendors in the "Constellation ShortList™ Next-Gen Databases: RDBMS for On-Premises."[4] Constellation follows Oracle as a vendor as well as its customers, prospects, and partners regularly and attends Oracle AI World (formerly Oracle CloudWorld), Oracle's analyst meetings, and more. Constellation has seen the following updates from Oracle in the last 12 months—with a focus on databases:

- **Oracle simplifies AI adoption with Oracle AI Database 26ai.** Creating AI-powered next-generation applications for enterprises is hard, and CxOs are looking for platforms that make it easier for their enterprise and teams to build and operate AI-powered business processes. Oracle does this by providing new AI capabilities seamlessly architected into Oracle AI Database 26ai that deliver wide freedom of choice to build and deploy AI applications through support for Apache Iceberg, Model Context Protocol (MCP), industry-leading large language models (LLMs), popular agentic AI frameworks, and Open Neural Network Exchange (ONNX) embedding models.

- **Oracle AI Database continues its multicloud success.** After the company offered Oracle AI Database on Oracle Cloud Infrastructure (OCI), customer demand led Oracle and other hyperscalers to start on a mission to bring state-of-the-art, full-function Oracle AI Database, including Exadata and Real Application Clusters (RAC), to other clouds—beginning with Microsoft Azure and then, in quick succession, adding Google Cloud and most recently Amazon Web Services (AWS).[5] This helps enterprises that are already on these clouds for other purposes leverage their enterprise data in Oracle AI Database running natively on OCI with unified governance and support. Customers benefit from combining the unique capabilities that power their enterprise data with their preferred cloud apps. With the introduction of multicloud credits at Oracle AI World 2025, Oracle simplifies the operational burden for its customers operating in the multicloud.[6]

- **Oracle readies Oracle AI Database for agentic AI with vector capabilities.** Oracle has added native vector support to its database and thus provides AI Vector Search for an easy way to access Oracle AI Database–structured data as well as unstructured or semistructured data stored elsewhere, such as in Oracle Autonomous AI Lakehouse. Customers can also combine data from LLMs and internal data in Oracle AI Database, using retrieval-augmented generation (RAG) to obtain more contextually relevant answers to business questions. At Oracle AI World 2025, Oracle shared that it has successfully vectorized its internal customer data, providing a critical proof point for the scalability of its vector database.

- **Oracle goes beyond just RAG with AI agents in Oracle AI Database 26ai.** Oracle introduced AI Vector Search and RAG in Oracle Database 23ai, enabling customers to combine data from LLMs and internal data to obtain more contextually relevant answers to business questions. With Oracle AI Database 26ai, Oracle architected agentic AI directly into the database with AI agents that can use multistep workflows to plan and try multiple approaches, ask for more data, and take actions individually or in concert with other agents. To accomplish this, Oracle is integrating with all the popular app-tier agentic frameworks, including OCI Gen AI, Google Vertex AI, Amazon Bedrock, LangGraph, and CrewAI, among others. Oracle is also providing frameworks to enable customers to easily build, deploy, and manage AI agents, using Select AI Agent, Private Agent Factory, and SQLcl MCP server.

- **Oracle extends data access with Autonomous AI Lakehouse and native support for Apache Iceberg.** Customers can now run AI on *all* of their data anywhere it resides—in Oracle databases and on data platforms such as Databricks or Snowflake and multiclouds such as OCI, AWS, Azure, or Google Cloud—via native support for Apache Iceberg in Oracle AI Database 26ai and Autonomous AI Lakehouse. They can leverage unified hybrid vector search to combine AI Vector Search with relational, text, JavaScript Object Notation (JSON), knowledge graph, and spatial searches—enabling retrieval of related documents, images, videos, audio, and structured data. In addition, Autonomous AI Lakehouse provides a "catalog of catalogs" to simplify data discovery and access across multiple data platforms and clouds.

- **Oracle solves decade-old database challenges with Oracle Globally Distributed Exadata Database on Exascale Infrastructure** (Oracle Distributed Database on Exascale,[7] for short). Enterprises have struggled with scaling databases since the first transactional database. With its proven sharding capability with built-in

Raft-based replication, Oracle has enabled Oracle AI Database to run globally, and enterprises can localize the shards. This enables critical data residency requirements and—for the first time—enables enterprises to stay compliant with data residency regulations.

- **Oracle positions itself as offering the leading cloud for AI workloads.** Oracle has steadily added NVIDIA capacity for OCI, expanding its capex as part of free cash flow higher and higher—to even more than 100% in Q4 FY 2025.[8] This positions OCI as the "go-to cloud" for companies building LLMs, even counting OpenAI as one of its customers. At Oracle AI World 2025, Oracle unveiled the OCI Zettascale10 AI supercomputer, which provides an unprecedented 16 zettaFLOPS, as part of the Stargate partnership.[9] In a major move to diversify its compute offerings, Oracle has also partnered with AMD, using 50,000 AMD Instinct MI450 Series GPUs, making OCI the first public cloud with a supercluster powered by AMD at this scale.[10]

- **Oracle Acceleron delivers improved, next-generation network performance to OCI customers.** Networking—the combination of storage and CPUs/GPUs—has been a critical design point of and a major contributor to OCI's success. Oracle is not resting on its laurels and is building next-generation networking with Oracle Acceleron that combines dedicated network fabrics, converged network interface card (NIC), host-level zero-trust packet routing, and multiplanar designs to boost performance and availability— remarkably at no additional cost for Oracle customers.[11]

- **Oracle Fusion Cloud Applications build on Oracle's "only SaaS with a PaaS" differentiator—with AI.** Oracle already offers the largest and most mature enterprise application cloud suite across finance, human capital management (HCM), supply chain management (SCM), and purchasing. It is catching up in customer relationship management (CRM) and manufacturing, differentiating all of its offerings as coming from the only software-as-a-service (SaaS) vendor with its own platform as a service (PaaS)—and even infrastructure as a service (IaaS). The suitewide adoption of built-in AI is a key 2025 example of this differentiator. This effort culminated (for now) at Oracle AI World, where Oracle launched a lot of additional AI automation and agents as well as an instantly partner-adopted AI Agent Marketplace.[12]

- **Oracle's "chip to click" differentiating strategy advances beautifully.** With the Sun Microsystems acquisition in 2010, Larry Ellison put into action Oracle's long-term plan to provide an integrated suite from the CPU (or GPU) to the user click. At Oracle AI World 2025, Oracle showed that the entire Oracle stack has been empowered by AI—not bolted on, but native to the stack—further differentiating Oracle from its competitors at any level, from IaaS to PaaS to SaaS.

OFFERING ANALYSIS

With more than a dozen vendors in the relational database management system (RDBMS) market alone, customers are challenged in choosing a trusted partner. Following are the strengths and areas for improvement in Oracle AI Database 26ai.

Highlights

RDBMSs remain central to enterprises' functioning and gain even more prominence in the AI era, because they are the system of record of the enterprise—which is the foundation to keep LLMs honest, via RAG capabilities. Oracle's claim at Oracle AI World was that 94% of the Fortune Global 100 companies use Oracle Database[13] (see Figure 1).

- **A unique approach to AI that caters to Oracle's offering strengths and customer demands.** Whereas all of Oracle's competitors have taken an approach in which transactional data is moved to data lakehouses or separate vector databases—making real-time AI impossible—Oracle has taken the opposite approach. It uses data lakehouse technology for nontransactional third-party data, builds vectors on it, and brings those vectors back to its database. Oracle can do this from an offering perspective because it can scale in practically any dimension (in contrast to almost all competitors). Customers appreciate this approach,

because they can keep their precious transactional data in the same place—their existing RDBMS—with all the security, access, and scale demands being addressed by Oracle AI Database.

- **The easiest migration: no upgrade needed.** Enterprises are always wary about upgrading their critical infrastructure, and databases are no exception. Oracle addresses this concern perfectly by delivering Oracle AI Database 26ai to its customers as an update. The transition from Oracle Database 23ai is a simple update, not requiring any recertification of people, applications, and hardware for customers that already are on 23ai. Laudably, some of the AI features are also available for Oracle Database 19c, enabling customers on these older versions to participate in the AI boom.

- **The power of the integrated suite in a single database.** Oracle has been delivering its converged database for more than a decade, making Oracle AI Database the one place to store vector; text; spatial; JSON; Extensible Markup Language (XML); and, of course, relational data. With this, Oracle AI Database is the one database for running AI/machine learning (ML), analytical, blockchain, distributed, graph, Internet of Things (IoT), online transaction processing (OLTP), and streaming workloads—all in one place. And it can run on-premises; in containers; on AWS, Azure, and Google Cloud—as well as, of course, on Oracle Cloud; and at the customer site with Oracle Cloud@Customer and OCI Dedication Region. With that, it is what Constellation refers to as a universal database.[14]

- **Innovation built on a stable foundation.** Oracle AI Database 26ai is basically built on Oracle Database 23ai with the addition of more than 300 new features and capabilities. Major innovation areas have, not surprisingly, involved AI-first capabilities such as vector (data type, memory management, indexing) support for GraphQL table functions for SQL and more—as well as JSON capabilities (e.g., JSON-to-Duality Migrator, JSON replication support, and more).

- **Natural language to work with data.** With Oracle AI Database 26ai, Oracle is also adding natural-language support to its database. Natural language is critical, because using it is the way humans communicate with each other. Now they can interact the same way with Oracle AI Database, which will lead to further democratization of the data inside Oracle databases with less technology-literate business users. Of course, even technology-fluent users will take advantage of natural language, because it accelerates task completion for them.

- **A better approach to caching.** Because the performance of enterprise applications has always been critical, caches have been introduced for many decades. But almost all caches gear their creation and operation to the application level of the technology stack, not the database level. With Oracle AI Database 26ai, Oracle changes that via Oracle True Cache, a lightweight read-only diskless Oracle AI Database that acts as a cache. As such, it behaves like a data cache that receives back-end changes in real time for its cache and makes it consistent as of a point in time. All Oracle SQL, vector, JSON, spatial, and graph query capabilities are available via True Cache. Exadata can also significantly accelerate AI vector queries by offloading them to intelligent storage.

- **Better security with Oracle SQL Firewall.** SQL injection has been one of the most prominent types of malware attack in the database space. With Oracle AI Database 26ai, Oracle introduces Oracle SQL Firewall, which effectively blocks unauthorized SQL and SQL-injection attacks. The result is greater peace of mind for the CxOs running Oracle Database.

Figure 1. Oracle AI Database 26ai
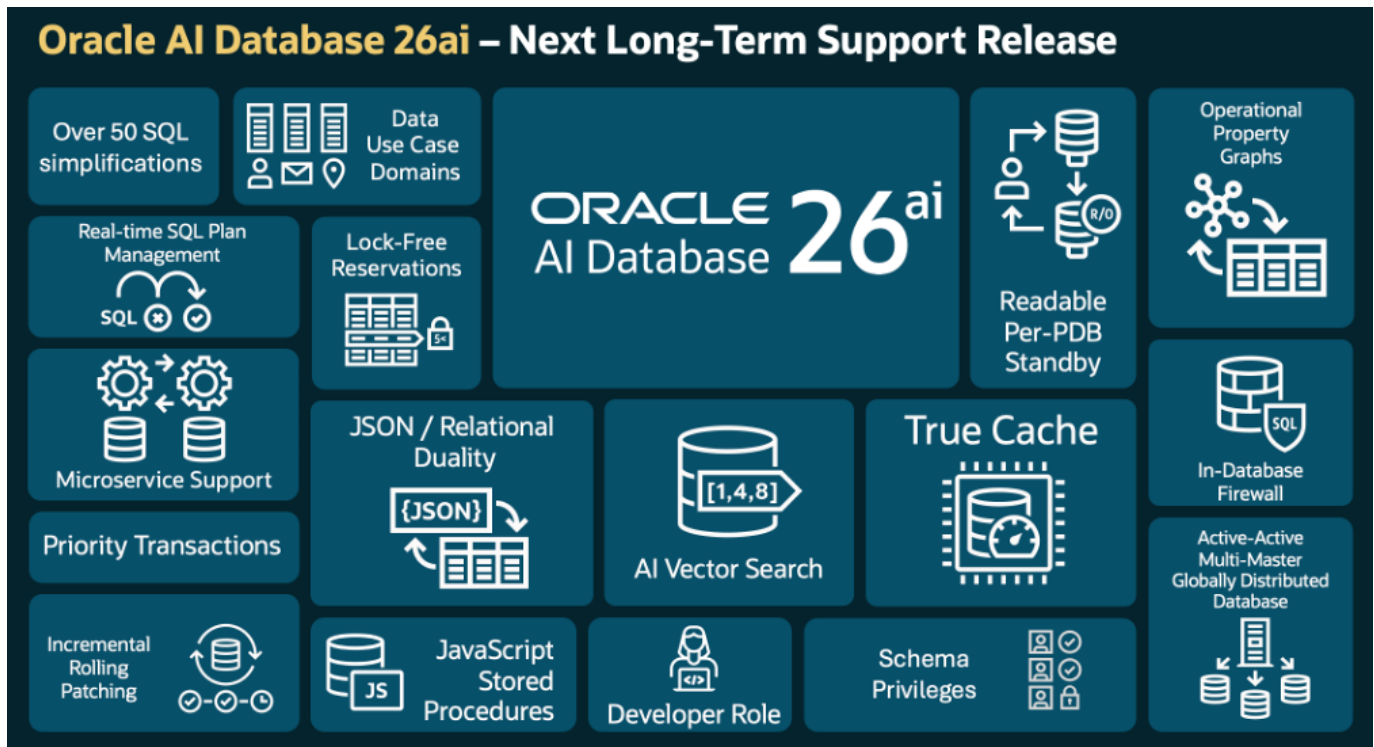
Source: Oracle

Opportunities for Improvement

Oracle AI Database 26ai was built on Oracle 19c, a solid and proven foundation that is already well established in the world's largest organizations, operated by skilled administrators, and used by millions of developers. That makes Oracle AI Database 26ai a well-trusted, industry-hardened offering with a robust underpinning technology. Oracle will now have to show that it can scale its vectors to enterprise-scale with early adopters to prove the validity of its approach of running AI while keeping the data inside its RDBMS. At Oracle AI World, Oracle executives mentioned repeatedly that Oracle has created a vector of its complete customer data—but size, performance, and other specifics have so far not been shared.

Note: Full strengths, weaknesses, opportunities, and threats (SWOT) analyses for this category are available as an advisory call for research-unlimited clients.

BUYING CONSIDERATIONS

According to Constellation's reference checks, customers and prospects choose Oracle Database for the following reasons:

- Want to consolidate on a single database for all types of data and workloads, including vectors

- Value scalability, security, and reliability highly

- Need to deploy Oracle AI Database on multiple clouds

- Are Oracle Fusion Applications customers

- Want to leverage NVIDIA as their AI platform

- Appreciate Oracle AI Database's portability and deployability with the highest identicality[18] from on-premises to all public clouds

- Prefer Oracle's built-in AI approach, leaving transactional systems in place and enabling AI via vectors

RECOMMENDATIONS

Constellation's executive network of CxOs from the AI150[16] and BT150[17] lists recommends the following best practices:

Enable Enterprise Acceleration[18] in the agentic AI era. Speed and agility are survival factors for the enterprise in the twenty-first century—even more so with the emergence of AI-powered next-generation applications. Unifying cloud and former on-premises data creates the data foundation for the AI era; now the ability to keep all transactional data in the database is a unique differentiator for Oracle AI Database 26ai that will help enterprises accelerate and become more agile in the agentic AI era.

Start building AI-powered next-generation applications. Oracle makes it easy to build AI-powered next-generation applications: The transactional data stays in place, and for the initially selected use cases, the unstructured data is made available via vectors. This enables CxOs to build AI-automated workflows, from the smallest to the largest scale, in their enterprise—a perfect approach to learning small and going live big, as needed.

Update for the transition to Oracle AI Database 26ai. For Oracle customers, there is no reason to wait. They are likely on Oracle AI Database already, due to their scalability and security requirements. Keeping the data there for agentic application automation makes sense, because scalability and security headaches do not have to be repeated with another platform/system. Therefore, CxOs need to update as soon as possible.

Consolidate on Oracle Database 19c. Enterprises running older versions of Oracle Database (such as 11g, 12c, 12c release 2, and 18c) should consolidate these as soon as possible on 19c. Technology disruption is not nice to those who wait, and the sooner the enterprise can move to 26ai from 19c, the better shape it will be in for the agentic AI era.

Consolidate your data. Better insights are fostered by correlation of more high-quality data. Bringing Oracle AI Database transactional data together and adding other data sources, such as via use of LLMs or from lakehouses, create a powerful foundation for business insights and applications in the era of AI.

As a non–Oracle Database user, consider Oracle AI Database 26ai. In enterprises that do not use Oracle Database but are at their capacity limits for their current transactional database systems, CxOs need to consider switching to Oracle. The main driver is that enterprises need to plan for agentic AI success, which will result in more transactions. Already-maxed-out back-end systems will not be able to scale to run an agentic AI enterprise.

Take a stance on commercial prudence. No matter the vendor and offering, enterprises must make sure they pay for value. This is valid with Oracle as well, and CxOs must pay attention to ensure that licenses and services for implementation are providing their enterprise with an attractive TCO.

ENDNOTES

[1] 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

[2] Constellation Research Business Transformation 150™, Constellation Research, 2024. constellationr.com/business-transformation-150

[3] AX100 2024, Constellation Research, 2024. constellationr.com/work-constellation/ax-leaders/ax100-2024

[4] Holger Mueller, "Constellation ShortList™ Next-Gen Databases: RDBMS for On-Premises," Constellation Research, February 12, 2025. constellationr.com/research/constellation-shortlist-next-gen-databases-rdbms-premises-8

[5] This CR Pulse was triggered by Oracle's making Oracle Database@AWS generally available on Amazon Web Services (AWS).

[6] For more about Oracle Multicloud Universal Credits, see the press release "Oracle Introduces Multicloud Universal Credits," Oracle, October 15, 2025. https://www.oracle.com/news/announcement/ai-world-oracle-introduces-multicloud-universal-credits-2025-10-14/

[7] For more on Oracle Distributed Database on Exascale, see Holger Mueller, "Oracle Solves Decades-Old DBMS Challenges and Future-Proofs Oracle Database," Constellation Research, August 27, 2025. https://www.constellationr.com/research/oracle-solves-decades-old-dbms-challenges-and-future-proofs-oracle-database

[8] For more details, see Oracle Corporation Q4 Fiscal 2025 Financial Results, page 7. s23.q4cdn.com/440135859/files/doc_financials/2025/q4/Q425_Form8K_Exhibit99-1_Earnings_Release_Tables_Final.pdf

[9] For more details, see the press release "Oracle Unveils Next-Generation Oracle Cloud Infrastructure Zettascale10 Cluster for AI," Oracle, October 14, 2025. https://www.oracle.com/news/announcement/ai-world-oracle-unveils-next-generation-oci-zettascale10-cluster-for-ai-2025-10-14/

[10] For more details, see the press release "Oracle and AMD Expand Partnership to Help Customers Achieve Next-Generation AI Scale," Oracle, October 14, 2025. https://www.oracle.com/news/announcement/ai-world-oracle-and-amd-expand-partnership-to-help-customers-achieve-next-generation-ai-scale-2025-10-14/

[11] For more, see the press release "Oracle Introduces New Cloud Networking Capabilities for Any Workload," Oracle, October 14, 2025. https://www.oracle.com/news/announcement/ai-world-oracle-introduces-new-cloud-networking-capabilities-for-any-workload-2025-10-14/

ENDNOTES CONTINUED

[12] For more, see the press release "Oracle Launches Fusion Applications AI Agent Marketplace to Accelerate Enterprise AI Adoption," Oracle, October 15, 2025. *https://www.oracle.com/news/announcement/ai-world-oracle-unveils-ai-data-platform-empowering-customers-to-innovate-in-the-ai-era-2025-10-14/*

[13] This statistic was quoted in numerous keynotes (e.g., in Juan Loaiza's keynote on October 15, 2026, at Oracle AI World in Las Vegas).

[14] For more on universal databases in general, but more specifically the difference between a universal database and a suite of specialized databases, see Holger Mueller, "The Universal Database Versus a Suite of Specialized Databases," Constellation Research, April 5, 2021. *https://www.constellationr.com/research/universal-database-versus-suite-specialized-databases*

[15] By *identicality* Mueller means next-generation computing platforms that enable organizations to run the same code assets across clouds and on-premises without any code changes because the technology stacks are identical.

[16] "AX100 2024," Constellation Research, 2024. *https://www.constellationr.com/work-contellation/ax-leaders/ax100-2024*

[17] "Constellation Research Business Transformation 150™," Constellation Research, 2024. *https://www.constellationr.com/business-transformation-150*

[18] 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*

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 & Veittinger, 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   *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-suites, boards of directors, and vendor clients. Our mission is to identify, validate, and share insights with our clients.

Organizational Highlights

- 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*

ℹ️ *info@ConstellationR.com*   ❓ *sales@ConstellationR.com*

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