



What Impact Does AI Have on Your Oracle Multicloud Strategy?



by **Chris Hollies**
CTO of Capgemini's
UK Oracle Practice

Multicloud, Reconsidered

For most of the past decade, deciding where to run workloads across clouds was an IT decision, made on cost, resilience and regulatory fit. Most companies made it once and have not revisited it since.

Agentic AI changes the rationale. Where a company's Oracle data sits, and how easily AI can reach it, now determines how quickly AI turns into lower costs, faster products and better customer service. A good placement shows up on the P&L. A poor one produces another stalled pilot that never scales.

AI needs trusted data it can reach wherever that data sits, and Oracle Multicloud is what makes that possible without moving it. That argument is made in full elsewhere. This ebook asks the next question. Once the foundation is in place, what is it worth?

The pages that follow look at how the cloud decision has changed, what it takes to make data ready for AI, what sets apart the companies already getting results and why governance sets the pace once AI starts acting on its own. For companies already running Oracle across several clouds, the question now is how to get value from it.

01. The Cloud Decision That Shows Up on the P&L

Once agentic AI enters the picture, deciding where an Oracle workload runs stops being an infrastructure question and becomes a business one.

From IT decision to business decision

For years, deciding where to run an Oracle workload was a technology decision, based on skills, cost, resilience, regulatory obligations and existing standards. “The resulting business benefits were typically lower infrastructure costs, improved agility, and greater operational resilience,” says Chris Hollies, CTO of Capgemini’s UK Oracle Practice.

AI has changed the question. “Where customers were once asking ‘where should we run the application?’ they are now asking ‘where is the data that AI needs to access?’” says Al Walker, OCI Presales Lead at Capgemini North America. Moving large volumes of Oracle data between clouds is expensive, slow and hard to govern, so companies increasingly want that data to sit close to the AI services they use. Hollies frames the same change in business terms. “The question has shifted from where should this workload run to where can this workload create the greatest AI-enabled business value.”

Agentic AI is what makes this a board-level issue. “Most technology waves improve IT operations. Very few fundamentally change how work gets performed,” Hollies says. Agents do more than generate content. They carry out tasks inside business processes with limited human intervention, which can cut manual effort, speed up delivery and improve customer service. All of that depends on agents being able to reach data the business trusts.

Does that push companies toward more clouds or fewer? In practice, both. Walker sees companies using more clouds at the AI layer, to pick the best models from each provider, while consolidating at the data layer around “one trusted foundation for enterprise data, governance, security and metadata.” Hollies describes the result as a more deliberate kind of multicloud. “The organizations seeing the best results from AI are not necessarily using fewer clouds. They are using each cloud much more intentionally.” Walker also points to Oracle Multicloud, which extends Oracle’s AI Database across AWS, Azure and Google Cloud so a company can use them as a single environment.

What AI brings to the surface

Day to day, AI exposes problems companies had learned to live with. Duplicated data stores and complex integration pipelines were “largely invisible to business users,” Hollies says, until AI had to rely on them. The job of IT changes too. Beyond applications and databases, teams now manage AI models, agents and prompts, and need to know which agent accessed which data and which model produced which answer. As Walker puts it, AI turns multicloud “from primarily an infrastructure discussion into a data, governance and operations discussion.”

There is a prior consideration which Hollies raises. Before anyone asks how AI will change the P&L, he wants to know what business problem it is being asked to solve. “Are you trying to increase customer satisfaction? Are you trying to produce more widgets per hour? Are you trying to produce higher quality by X percent?”



Answer that, and the financial case follows from it. Start with the P&L, and a project tends to acquire a number before it has a purpose.

What does a strategy that has paid off look like a year from now? Hollies is wary of the obvious answer. “More pilots reaching production is not business value,” he says. “People can focus on how many things we drive into production. It does not necessarily mean it is driving any business value.” What he looks for instead is movement on the metrics a business already runs on. Walker lists the ones clients raise most often. Faster times to get shipments out. Faster times to fill open roles. Faster decisions. Lower cost across all of it. “That is what the measurable piece of the AI is designed to actually do,” Walker says. Hollies sums up the goal as “a repeatable capability to turn trusted enterprise data into measurable business value.”

—
“AI invites us to think beyond where data lives to what it can make possible.”

Kristian Toms

Vice President EMEA
Database Black Belts and
Multicloud Expert, Oracle

02. The End of the Trade-Off

Oracle Multicloud ends the old trade-off between moving data to AI and keeping it out of reach, which puts the focus on what modernized data is worth.

Ending the old dilemma

For years, companies with critical data in Oracle faced an uncomfortable choice. They could move that data to the cloud where their AI services ran and absorb the cost, latency and security risk, or they could leave it where it was and limit which AI services could use it. Oracle Multicloud removes much of that tension by placing Oracle Database inside AWS, Azure and Google Cloud, next to the AI services that need it.

Ending the trade-off does not make the data ready. Most ERP, finance and CRM systems “were designed to process transactions, not answer questions, drive insights or support autonomous agents,” says Abi Giles-Haigh, Analytics and Innovation Director at Capgemini. Modernizing that data, she says, “means making it discoverable, governed, trusted and consumable by AI, analytics and business processes.” In practice, that means clear owners for critical data, better quality, shared business definitions and consistent governance and security controls.

Where the value shows up

For the people using these systems, the benefit shows up in their daily work. A customer service agent gets real-time recommendations during a call. A finance team forecasts against live ERP data. A supply chain planner analyzes inventory without waiting for data to replicate. The value, Giles-Haigh says, shows up as “faster decisions, more responsive digital experiences and quicker time to insight.”

The first result can come quickly. Timelines depend on data quality and governance maturity, but Giles-Haigh suggests thinking “in terms of weeks not months” for an initial use case. The larger return comes after that. “The more important question is not how quickly the first use case appears, but how quickly the second, third and tenth use cases can be deployed,” she says. Because each new initiative builds on the same trusted data, every AI project that follows takes less time and costs less to deliver.

That is the real value of modernized data. It turns AI from a series of one-off projects into a capability the business can reuse.

“The more important question is not how quickly the first use case appears, but how quickly the second, third and tenth use cases can be deployed.”

Abi Giles-Haigh

Analytics and Innovation Director,
Capgemini

The Best-in-Class Playbook

A small share of companies already turn generative AI into measurable results, and what sets them apart is a handful of deliberate choices.

65%

Capgemini’s research projects that effective human-agent collaboration will lift human engagement in high-value tasks by that margin, with AI agents involved in most business tasks within three years.

1. Data first

They modernize their data foundations before scaling AI, treating data as a prerequisite rather than an afterthought.

2. Outcomes before models

They start with a business problem or a customer goal and choose the technology afterward.

3. Business change first

They redesign processes, governance and ways of working alongside the technical build.

4. An operating model that connects it all

“Many organizations possess impressive technology. Many possess talented employees. Fewer possess operating models that allow those capabilities to work together effectively,” says Abi Giles-Haigh, Analytics and Innovation Director at Capgemini.

5. Where to start

Measure business outcomes rather than the number of pilots, then pick a small number of high-value processes and redesign them around trusted data, governance and AI.

Sources: Capgemini Research Institute, “Rise of agentic AI: How trust is the key to human-AI collaboration,” 2025.
Quote from Abi Giles-Haigh, Analytics and Innovation Director, Capgemini.

03. Trust, Not Ticked Boxes

For a board, AI governance decides whether an AI project scales safely or gets shut down. Compliance is the smallest part of it.

From outputs to actions

Ask a board and a technical team about AI governance and you will hear two different concerns. “Boards worry about business continuity, regulatory exposure, reputation, customer trust and operational risk,” says Abi Giles-Haigh of Capgemini. Technical teams focus on the controls that manage those risks. Good governance turns the board’s appetite for risk into controls the technical team can put in place.

Often the governance decision that matters most is a simple one. Giles-Haigh describes a company whose AI forecasts disagreed because its business units drew on different data sources. Naming the ERP platform as the single authoritative source fixed it, and forecasting became part of planning. The outcome, she says, is “improved trust, wider adoption and faster deployment of AI-enabled processes.”

Oversight has to change once AI stops recommending and starts acting. When an agent starts and completes a task inside a live business process, the pause for review may not exist. “The governance model must evolve from reviewing outputs to governing actions,” Giles-Haigh says. “The question changes from ‘Was the answer correct?’ to ‘Was the action appropriate, authorized and traceable?’” In practice, that means clear limits, an escalation route and an audit trail.

Where humans stay in the loop

Capgemini’s own study found that trust in fully autonomous agents has fallen. The reason is maturity. As companies move from exploring AI to implementing it, they understand where a human belongs in the loop. Among organizations already implementing, 47% report above-average trust, against 37% of those still exploring.

Giles-Haigh is direct about the limits. “Agents still hallucinate, as incredible as they are now,” she says. “You can have a conversation with an agent and coax it into one point of view and then quickly get it to change its point of view.” She does not expect that to be solved soon. “That hallucination problem is not going away anytime soon.” Caution about full autonomy is therefore warranted.

Agents running across several clouds draw on many platforms and act across many applications. That makes one practice non-negotiable, which is clear accountability for data and decisions. A company must know which sources are authoritative and who is accountable for what each agent does. Without an answer to “Which source of data should this AI trust?” every other governance control becomes much harder to enforce, Giles-Haigh says.

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real

www.capgemini.com

Contact

Al Walker, OCI Presales Lead,
Capgemini North America

