

White Paper

Cloud Databases: Unlocking Scalability and Flexibility for Businesses.

*Unparalleled scalability, data agility, and built-in AI functionalities
open doors for new data management approaches.*

 **Hitachi Vantara**

Kireet Kokala
Director of Hyperscaler
Data Science



The modern data landscape continues to evolve rapidly while aligning itself with user needs and core enterprise features. Powered by the Internet of Things (IoT) and Artificial Intelligence (AI), this evolution demands a novel approach to holistic database management.

The need is more pronounced as traditional on-premises solutions struggle to handle diverse data workloads' increasing volume, variety, complexity, and real-time requirements.

This historical trend, known as data gravity, specifies that applications need to reside near centralized data storage for efficient processing and analysis. However, the rise of cloud databases, alongside advancements in network technologies, continues to challenge the notion. Cloud databases offer unparalleled scalability, data agility, and built-in AI functionalities, opening doors for new approaches to data management.

On the one hand, hybrid cloud environments offer compelling solutions while managing databases across on-premises and cloud infrastructure presents its own challenges. To achieve successful alignment with their data strategy, businesses and enterprises must carefully consider how to navigate complexities, balance spend, and unlock the full potential of cloud databases.

Popular options are typically native offerings from hyperscalers like AWS, Azure, Google Cloud, and IBM, which emerged in the late 2000s as a response to the limitations of traditional solutions.

They offer significant advantages such as:

- Unparalleled Scalability: Effortlessly handle massive, exascale datasets required by modern AI applications.
- Enhanced Data Agility: Facilitate seamless data movement between on-premises and cloud environments for optimal analysis.
- AI-led Innovation: Embrace built-in AI functionalities and integration with leading governance frameworks, accelerating innovation for AI-driven applications.

This white paper explores the changing data landscape and the complexities of hybrid cloud database management, where cloud databases can be scalable, future-proof solutions, and Hitachi Vantara's role in helping AI-powered businesses.



Interestingly, 37% of surveyed enterprises find running AI applications on their current IT infrastructure a “significant” challenge⁶.

Hybrid Cloud and Cloud-Native Solutions

Over the past decade, the hybrid cloud has reemerged as a strong force in modern data management. It offers businesses and regulated entities a cost-effective and flexible approach. Businesses can easily extend their existing on-premises infrastructure by leveraging the agility, scalability, and often lower operational costs of the public and private cloud. This strategic shift is reflected in a recent study where a resounding 95% majority of organizations identified strategizing hybrid cloud as a top three priority on their tech roadmap¹.

Building upon the established foundation of traditional database solutions, cloud-native relational database management systems have surfaced as a preferred pillar in the technical community. Introduced in the late 2000s, options like Amazon Redshift (2012), Snowflake (2014), Microsoft Azure Synapse Analytics (2016), and Google Cloud's AlloyDB (2019) and AlloyDB Omni (2022) disrupted the market with innovative architectures explicitly designed for the cloud.

These new options offer compelling advantages for modern data management, particularly for workloads demanding scalability and agility. Cloud-native solutions excel at handling the vast datasets required by AI applications, often at a lower cost per storage unit compared to traditional options.³ Additionally, they boast on-demand scaling of storage capacity, leading to significant cost savings over traditional databases with fixed configurations.

Cloud-native solutions also enable seamless data movement between on-premises and cloud environments, as they help optimize analysis and workload distribution for businesses. Of high interest to enterprises, these capabilities empower them to leverage AI services for real-world applications:

- **Fraud Detection:** Several financial institutions leverage cloud databases and AI frameworks like PyTorch and Keras to analyze vast amounts of real-time transactional data, identifying and preventing fraudulent activities. For example, JPMC bank utilizes a fraud detection system developed using cloud-based TensorFlow to analyze millions of daily transactions, significantly reducing fraudulent activity.
- **Predictive Maintenance:** Multiple manufacturing companies use cloud databases and AI to analyze sensor data from machines, predict equipment failures, and schedule preventive maintenance. GE Aviation employs a predictive maintenance system that runs cloud-based machine learning to help analyze engine data and predict potential failures, preventing costly downtime.



30%

Recent cloud migrations have shown that they can reduce IT infrastructure costs by up to 30%⁵

Hybrid Cloud Databases Management Challenges

Users have many reasons to run their workloads on-premises like meeting regulatory and compliance requirements. However, they try to accomplish this process most cost-effectively while using cloud technologies. This makes the hybrid cloud a very interesting proposition! While cloud deployments offer flexibility and scalability, managing databases across disparate environments can be complex. Heterogeneity, wherein a mix of on-premises and cloud-based databases exist with varying tools and processes, can be cumbersome and error-prone. Optimizing performance for data access across the hybrid cloud, especially for demanding workloads, can be difficult. Additionally, maintaining data security and ensuring compliance with regulations across different cloud environments requires robust security strategies and consistent policies. This underscores the growing need for solutions and strategies that can effectively manage and optimize cloud database deployments.

Edge Computing and the Future of Hybrid Cloud

Public cloud providers increasingly recognize the need for real-time data processing and analysis closer to the source of data generation (viz., data origination). They're extending their reach beyond the cloud through various strategies:

- Hybrid cloud solutions (Ex, AWS Outposts/Snow series, Azure Stack, [Google Distributed Cloud](#)) bring some of their cloud capabilities closer to on-premises data centers, enabling businesses to leverage the benefits of cloud computing while maintaining control over sensitive data.
- Edge computing services (Ex, AWS Greengrass, Microsoft® Azure® IoT Edge, Google Cloud Edge TPU) empower businesses to process and analyze data closer to the network edge, where devices like sensors and IoT systems generate data. By saving egress fees, this option allows for faster insights and quicker decision-making without transferring vast amounts of data to the cloud.

These advancements, along with the capabilities of cloud-native databases and AI, are fostering a new paradigm in data management. Businesses can leverage a hybrid cloud approach that combines on-premises infrastructure with public cloud resources and edge computing, enabling them to unlock AI's full potential while gaining significant competitive advantages.

Hitachi Vantara's Perspective

Hitachi Vantara, a leading hybrid cloud infrastructure and services provider, is proud to be the first hardware partner to certify the [E-series](#) and [VSP One Block](#) platforms with Google's cloud-native database, [AlloyDB Omni](#). This collaboration strengthens our commitment to delivering best-in-class solutions for businesses running AI services and real-time analytics workloads in hybrid cloud environments. While Hitachi Vantara maintains strong partnerships with established database vendors like Oracle, SAP, Microsoft, and IBM, our collaboration with Google Cloud signifies our latest commitment to providing flexible options for our customers through AlloyDB Omni.

Hitachi Vantara also observed that the rise of Generative AI (GenAI) applications further emphasizes the need for flexible data management solutions. Our findings indicate that 78% of study participants cited a mix of on-premises and public cloud for building and utilizing GenAI solutions². This highlights the ability of cloud-native solutions to seamlessly integrate with both on-premises and cloud environments, catering to the specific resource requirements of GenAI workloads!



A study on optimizing hybrid cloud environments found that over 50% of organizations view optimizing their hybrid cloud technologies as their most important priority over the next 12 to 24 months.¹

The natural partnership between Hitachi Vantara and Google ensures our platforms seamlessly integrate with cutting-edge cloud-native solutions, enabling businesses to harness the full potential of the hybrid cloud.

Our partnership delivers a validated and powerful solution: E-series and VSP One Block storage optimized for Google's cloud-native AlloyDB Omni database. This combined offering enables businesses to readily navigate the changing data landscape and address various challenges of managing databases in a hybrid cloud environment.

By leveraging this solution, businesses can:

- Unleash the potential of AI and real-time analytics: Gain a competitive advantage through faster development cycles and more efficient data analysis for AI workloads.
- Simplify hybrid cloud deployments: VSP One Block seamlessly integrates with AlloyDB Omni, streamlining deployment and management across hybrid cloud environments.
- Optimize performance and scalability: VSP One Block's robust storage infrastructure ensures exceptional performance for even the most demanding database workloads.

AlloyDB Omni, a Postgres rdbms with vector database capabilities, boasts features specifically designed for AI and real-time analytics. By integrating VSP One Block with AlloyDB Omni, we empower businesses to accelerate AI Innovation. Also, leveraging the power of AlloyDB Omni for AI workloads resulted in faster development cycles and more efficient data analysis as seen in our lab testing. This is

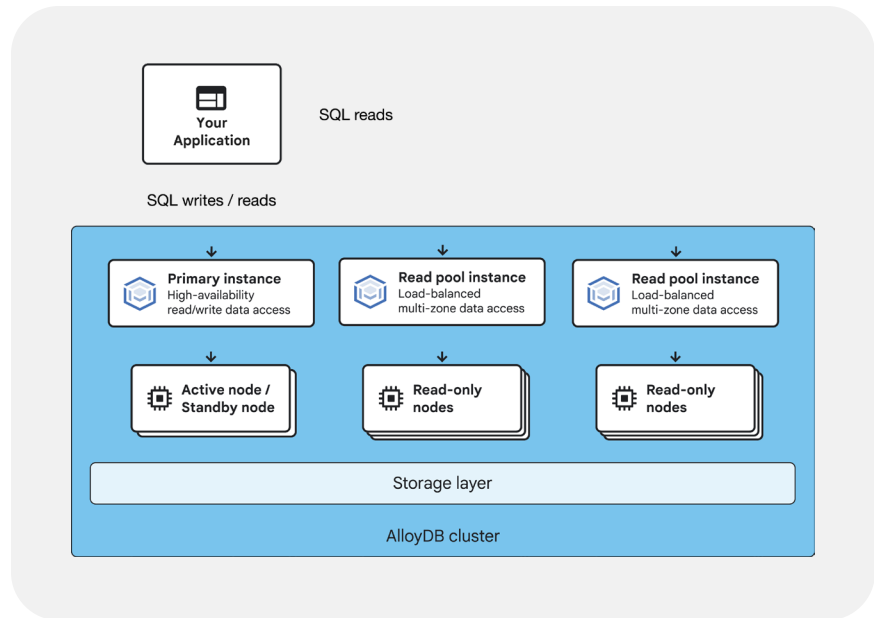


Fig. 1 Google Cloud AlloyDB Reference Architecture

significant news as research indicates that next-generation AI applications stand to add up to \$4.4 trillion annually to the global economy.³ VSP One Block's storage capabilities streamline data management across hybrid environments, ensuring data security and accessibility for critical AI applications.

Through rigorous benchmarking, we've validated that VSP One Block delivers optimal performance for AlloyDB Omni deployments.

Here are the key takeaways from our findings:

- Superior Performance with 32 LUN Configuration: Distributing data across 32 LUNs within VSP One Block significantly outperformed a single LUN configuration in various workloads. This highlights VSP One Block's ability to optimize storage access for demanding database applications like AlloyDB Omni.

- Alignment with Google Cloud's Benchmarks: The benchmarking results closely mirrored the requirements outlined in [Google Cloud's documentation for AlloyDB Omni](#). This ensures compatibility and expected performance when deploying AlloyDB Omni on VSP One Block in a hybrid cloud environment.

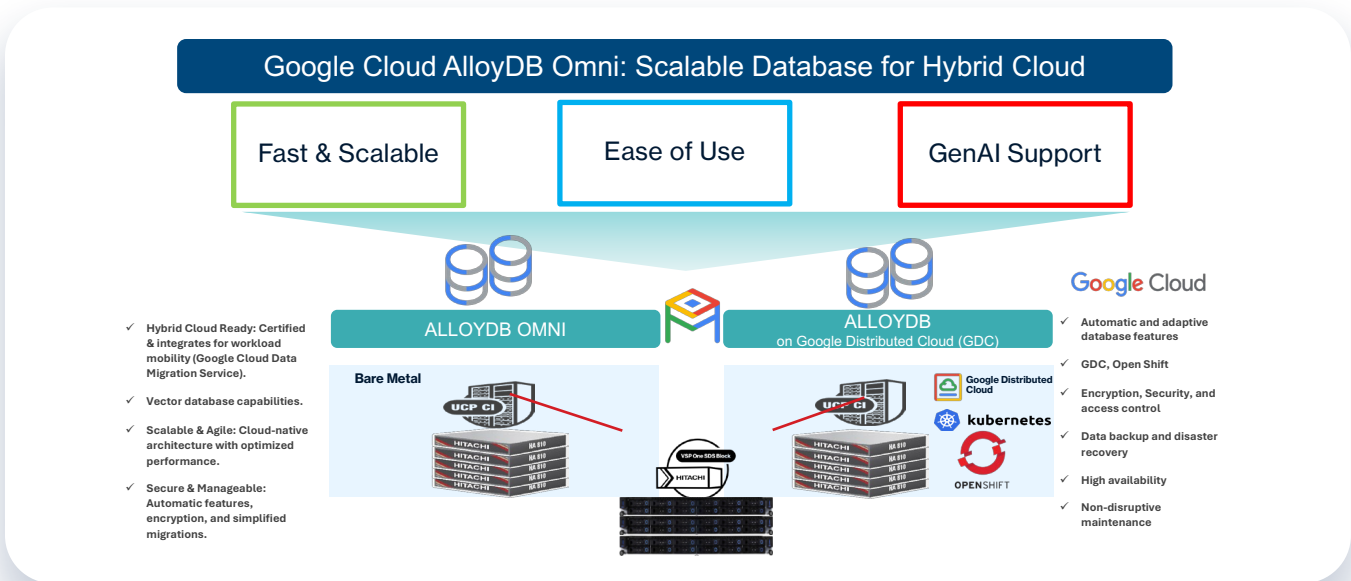


Fig. 2 Scalable Vectorized Database for Hybrid Cloud

The modern data landscape is undergoing an overhaul. Cloud databases are becoming increasingly popular due to their advantages in scalability, agility, and overall ease of management.


This shift requires a hybrid cloud approach that combines on-premises infrastructure with public cloud resources to effectively manage data sprawl while supporting modern workloads demanding real-time analytics and AI applications.

Nevertheless, managing databases across these disparate environments also presents several challenges. Optimizing performance, ensuring data security and compliance, and navigating heterogeneous environments are all hurdles that organizations must overcome. Hitachi Vantara's platforms, combined with Google Cloud's AlloyDB Omni, provide a compelling solution that addresses these challenges and unlocks the full potential of cloud-native databases in a hybrid cloud environment.

We provide compelling solutions that address the challenges outlined in the whitepaper and unlock the full potential of cloud-native databases in hybrid cloud environments. This integrated solution offers a robust foundation for managing your data across disparate environments and empowers you to leverage the power of AI and real-time analytics. Stay tuned for our upcoming publications for a deeper dive into the specifics of ETL for AI within a hybrid cloud strategy. In the meantime, to explore how this cloud database solution can address your specific needs and help you thrive in today's data-driven world.

Hitachi Vantara and Google Cloud are well-positioned to help you unlock the power of AI and real-time analytics in your hybrid cloud environment. Contact us today and visit our website to learn more about our certified and optimized solution for AlloyDB Omni on E-series and VSP One Block.

A recently concluded survey found that over 80% of organizations believe hybrid IT environments benefit their ability to manage applications and data.⁴



Why Hitachi Vantara and Google's Partnership Makes Sense

A powerful combination of reliable data storage and leading-edge cloud technologies has led to a symbiotic partnership. Hitachi Vantara, a leader in data storage solutions, and Google Cloud, a pioneer in cloud computing, have joined forces to address many challenges. Hitachi Vantara's proven hardware infrastructure ensures smooth data access for AI applications running on Google Cloud's cutting-edge cloud databases. Conversely, it offers secure and scalable cloud solutions designed for AI and real-time analytics. This strategic alliance enables businesses to bridge the gap between on-premises data and the cloud, simplifying hybrid cloud management and unlocking the full potential of AI. Customers can expect faster AI development cycles, efficient data analysis, and robust security – all of which are crucial for becoming AI-powered businesses.

About The Authors

Kireet Kokala is the Director of Hyperscaler Data Science. He helps lead the strategic implementation of Data Science and advanced analytics initiatives, focusing on hyperscalers like AWS, Azure, and Google Cloud Platform.

Bjoern Rost is a Product Manager for Google Cloud Databases who has architected, implemented, and tuned many data warehouses and relational transaction systems around the world.

Peer reviewers: Alice Chau, Dilip Modi, Michael Hay, Spencer Tidwell, Ted DuBois.

About Hitachi Vantara

Hitachi Vantara is transforming the way data fuels innovation. Hitachi Vantara provides the data foundation the world's leading innovators rely on.

The company helps customers build the foundation for sustainable business growth through data storage, infrastructure systems, cloud management, and digital expertise. To learn more, visit www.hitachivantara.com.



- 1 Enterprise Strategy Group Custom Research Report, Market Opportunity Research: Hybrid Cloud Buyer Insights, April 2023. <https://www.techtarget.com/esg-global/>
- 2 Enterprise Strategy Group Custom Research Report, Hitachi AI Infrastructure: Buyer Insight Research, February 2024. <https://www.techtarget.com/esg-global/>
- 3 GigaOm Research. (2022). A study on cloud-native database economics.
- 4 McKinsey & Company. (April 2024). The economic impact of artificial intelligence. <https://www.mckinsey.com/mgi/overview>
- 5 McKinsey & Company. (April 2023). Projecting the global value of cloud: \$3 trillion is up for grabs for companies that go beyond adoption. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/projecting-the-global-value-of-cloud-3-trillion-is-up-for-grabs-for-companies-that-go-beyond-adoption>
- 6 Nutanix. (April 2024). The State of Enterprise Cloud Adoption. <https://www.nutanix.com>



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
hitachivantara.com | community.hitachivantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
hitachivantara.com/contact

© Hitachi Vantara LLC 2024. All Rights Reserved. HITACHI and Pentaho are trademarks or registered trademarks of Hitachi, Ltd. All other trademarks, service marks and company names are properties of their respective owners.

HV-BTD-WP-Unlocking-Scalability-and-Flexibility-for-Businesses-28June24-A