VULTR

Qdrant

Qdrant + Vultr for AI and ML Development

Vultr's cloud infrastructure and Qdrant's leading-edge vector database integration, designed for any AI and ML workloads.

VULTR.COM

Qdrant + Vultr: Vector Database for Next-Generation AI Projects

Vultr and Qdrant provide enterprises with advanced AI application benefits, including scalability, global reach, best price-to-performance ratios, data management, and privacy controls.

As a part of the Vultr Cloud Alliance, Qdrant offers a cutting-edge vector database crucial for developing advanced AI and ML applications. Running on Vultr's robust cloud infrastructure, this premier vector database technology supports businesses in implementing comprehensive AI strategies through a partnership that creates an integrated platform that delivers unrivaled performance, scalability, and cost efficiency for global enterprises.

Qdrant's technology enables the management of billions of vectors and supports the intricate matching of semantically complex objects, all built on Rust for the utmost performance and safety. Through this collaboration, businesses can harness the full capabilities of embeddings or neural network encoders for applications that require precise matching, searching, and recommending at scale.

Leveraging the strengths of Vultr's global infrastructure with Qdrant's Vector Database as a Service, organizations can build, deploy, and manage their AI operations effortlessly. This enhances data resilience and security and equips organizations with the agility and scalability needed to navigate the evolving demands of modern AI environments. Enterprises gain access to a comprehensive range of compute resources, integrated with Qdrant's top-tier vector database management, to optimize IT efficiency and operational effectiveness, ensuring a competitive edge in the digital economy. Utilizing Vultr's infrastructure – including Kubernetes clusters and Vultr Cloud GPUs, powered by NVIDIA – or cloud inference, enterprises can be fully equipped to handle AI/ML projects.

Vector databases on high-performance cloud infrastructure

Qdrant and Vultr redefine AI infrastructure by seamlessly integrating high-performance vector databases with scalable cloud solutions, offering unparalleled cost efficiency, optimized data management, customizable deployments, and accelerated AI processing for advanced applications.

Improved price to performance

Qdrant and Vultr dramatically cut cloud costs — up to 90% compared to traditional hyperscalers — while delivering powerful AI capabilities. The cost-performance advantage is essential for companies looking to maximize their technology investments without compromising quality.

Optimized vector database integration

Integrating Qdrant's vector database with Vultr's cloud infrastructure for advanced AI applications, combined with high-performance vector technology akin to search, enhances data analysis and processing capabilities, optimizing both performance and scalability.

Go further with hybrid cloud

With Qdrant and Vultr, businesses can deploy tailored vector databases within Kubernetes clusters, ensuring data sovereignty and privacy. Enjoy fully managed cloud services like monitoring, scaling, and backups while running data workloads alongside your main application code.

A solution optimized for AI deployment challenges

Performance bottlenecks

Vultr's GPU-accelerated Kubernetes Engine, powered by NVIDIA and integrated with Qdrant Hybrid Cloud, empowers developers to create advanced AI and machine learning systems capable of addressing complex challenges and efficiently deploying across diverse infrastructures.

Scalability challenges

Embrace scalability with Qdrant and Vultr's infrastructure, integrating Managed Kubernetes seamlessly for horizontal and vertical scaling, zero-downtime upgrades, and disaster recovery, ensuring confident deployment anywhere.

Productivity optimization

Vultr and Qdrant offer cutting-edge hardware acceleration and advanced data storage solutions, enhancing developers' productivity by enabling easy access and deployment of state-of-the-art AI technologies.

Our partner advantage

Access anywhere

Vultr's expansive infrastructure spanning across six continents and Qdrant's rapid vector search capabilities enable businesses to deploy AI solutions worldwide. This combination ensures low latency and adherence to local data regulations, optimizing user experience across different geographies.

Unified stack management

Seamlessly manage the entire AI application lifecycle in one integrated environment with Qdrant and Vultr, simplifying workflows and accelerating development.

Rapid deployment

Quickly deploy Qdrant Hybrid Cloud on Vultr with a one-click installation, streamlining the setup process and enabling rapid scalability.

Privacy driven efficiency

Empower your data management strategy with comprehensive sovereignty and privacy control across various data center regions. Seamlessly integrate managed solutions to run data workloads alongside your main application code, optimizing efficiency while maintaining privacy.

A better way to build AI and ML

Easy to use API

Vultr's API-first approach and Qdrant's OpenAPI v3 specification create highly customizable client libraries in almost any programming language, enhancing operational flexibility.

Enhance LLMs with RAG

Improve the accuracy of large language models with Qdrant-powered Retrieval-Augmented Generation (RAG) by integrating external information retrieval, supported by Vultr's scalable infrastructure.

GPU-accelerated Kubernetes engine

Supercharge AI and machine learning workloads with Vultr's GPU-accelerated Kubernetes engine, powered by NVIDIA, and Qdrant's efficient semantic search capabilities integrated into clusters for lightning-fast data processing and analysis. This powerful combination enhances performance for AI-driven applications.

VULTR **Qdrant**

Learn more about Qdrant and Vultr

Contact us at vultr.com to get started. →