


VergelQ: On-Premises AI for Every Enterprise

Private, Generative AI & VMware Alternative in One Software Solution



Fast Facts:

- **Instant AI Deployment:** Integrated directly into VergeOS, deploying fully operational AI environments in minutes.
- **No Specialized Expertise Needed:** Designed for immediate usability without requiring AI-specific training or configurations.
- **Completely Private and On-Premises:** Ensures absolute data control and security within your own data center.
- **Vendor-Neutral GPU Support:** Compatible with standard GPUs, avoiding lock-in and reducing hardware costs.
- **Built-in VMware Alternative:** Provides a comprehensive VMware exit strategy along with integrated generative AI capabilities.

Secure, private AI integrated into VergeOS

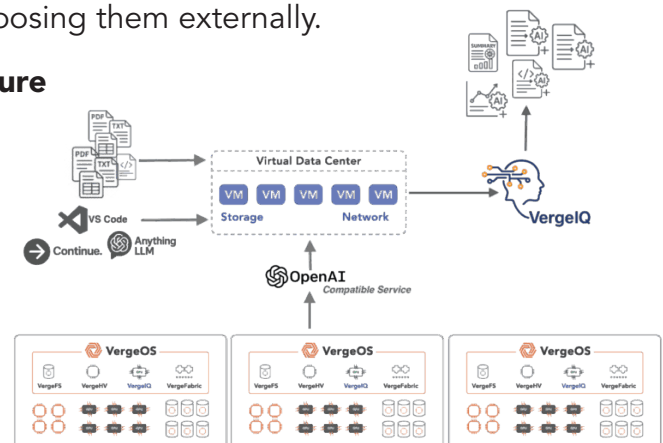
VergelQ makes **AI a native resource within VergeOS**, just like storage (VergeFS) and networking (VergeFabric). It enables enterprises to leverage fully private, on-premises generative AI environments with the same simplicity, efficiency, and control they expect from their core infrastructure. VergelQ enables organizations to maintain complete control over sensitive documents, proprietary code, and internal datasets, while generating immediate insights, tailored content, and actionable intelligence—all securely within their own data centers.

What Can you Do With VergelQ?

VergelQ enables your enterprise to use private AI capabilities within your own secure infrastructure. Within minutes of VergeOS installation, you'll gain practical, actionable insights while maintaining complete data privacy and control.

Value for the Data Center:

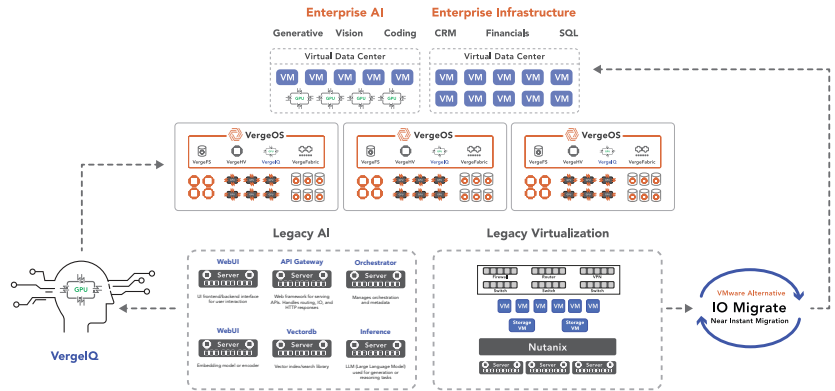
- **Rapidly Analyze Documents** — Securely upload PDFs, DOCX, TXT, or HTML files to produce summaries and actionable insights instantly.
- **Audit Proprietary Code** — Quickly review and optimize internal codebases without exposing them externally.
- **Automate Infrastructure Tasks** — Generate accurate scripts to automate and simplify infrastructure management beyond VergeOS.



VergelQ: On-Premises AI for Every Enterprise

Value for the Enterprise:

- **Generate Content** — Efficiently create documentation, knowledge articles, and policies from your secure datasets.
- **Real-Time Infrastructure Insights** — Directly query your infrastructure for immediate, actionable operational data and insights.
- **One Install Solves Three Problems** — Exit VMware, Modernize Infrastructure, and deliver AI to the organization. At your pace, all at once, or one at a time.



Why Choose VergelQ for On-Premises AI?

VergelQ uniquely combines private generative AI capabilities with a seamless VMware alternative—all within your data center. Explicitly designed to simplify infrastructure management, improve resource efficiency, and ensure complete security and compliance, VergelQ empowers enterprises to embrace generative AI confidently.

- **Immediate Usability** — Deploys automatically as part of VergeOS, delivering secure generative AI environments quickly, without complex configurations or training.
- **Flexible, Vendor-Neutral Infrastructure** — Leverage any standard GPU hardware and popular Large Language Models (LLMs), including LLaMa, Qwen, Phi, Gemma, avoiding vendor lock-in and reducing costs.
- **Optimized Resource Management and Efficiency** — Intelligent GPU pooling and dynamic resource allocation provide near-bare-metal performance, enhancing infrastructure efficiency and scalability.

