



Solution Brief

TrueNAS V-Series

Performance for Generative AI and Autonomous Agents

Purpose-built for the next intelligence frontier, the TrueNAS V-Series provides the low-latency foundation for training Generative AI and deploying autonomous agents, turning data into action without the friction of a steep learning curve.

Features

Tri-Mode Storage Flexibility

Consolidate storage choices in a single box with 24 front-loading bays supporting SAS HDD and Gen4 NVMe SSDs.

Massive Scalability

Support for up to 20 PB of NVMe storage or up to 40 PB of HDD storage with added expansion shelves.

Next-Gen Performance

Includes 4 rear Gen5 NVMe drive bays for high-velocity data caching.

Extreme Network Velocity

Support for high-speed fabrics up to 400 GbE.

Advanced Data Management

Features Fast Deduplication, block copy, and RDMA support as part of the TrueNAS Enterprise suite.

Benefits

High Availability (HA)

Ensure uninterrupted operations with dual controllers featuring automated rapid failover.

Frictionless Versatility

Easily adapt to changing workloads with tri-mode bays that support both hybrid and all-flash configurations.

Simplified Portfolio

Benefit from a versatile platform and an intuitive interface that requires no sharp learning curve.

Enterprise Efficiency

Advanced features like inline compression, deduplication, snapshots, and replication come standard, reducing management costs.

Contact

Phone: +1 (866)-TRUENAS

Email: sales@truenas.com

Web: www.truenas.com

Technical Powerhouse for AI Workloads

- **Extreme Network Velocity with 400 GbE:**
Support high-speed fabrics up to 400 GbE to keep GPU clusters fully saturated during AI training and inference.
- **Ultra-Low Latency via RoCE:**
Integrated RDMA with RoCE enables ultra-low-latency communication for AI agents processing feedback and executing complex workflows.
- **Next-Gen NVMe over Fabric (NVMe-oF):**
NVMe-oF (TCP/RDMA) extends local flash performance across the network, enabling up to 20 PB of high-speed NVMe access.
- **High-Velocity Caching:**
Four rear Gen5 NVMe bays act as cache for hot data, accelerating access to model weights and agent memory logs.
- **Massive Tri-Mode Scalability:**
Consolidate entire AI pipelines in one system scaling to 40 PB of HDD or 20 PB of NVMe storage.

Enterprise Benefits for the Intelligence Era

- **High Availability (HA) for Persistent Agents:**
Dual controllers and rapid failover provide 24/7 reliability, ensuring autonomous systems never lose critical memory or mission uptime.
- **Frictionless Data Management**
TrueNAS Enterprise features standard deduplication, compression, and snapshots to slash the cost of storing massive, redundant AI datasets.
- **Versatile & Adaptive:**
This tri-mode platform adapts seamlessly from high-capacity HDD data lakes to all-flash performance tiers for any modern workload.
- **Secure and Compliant:**
TrueSecure offers hardware-accelerated encryption and FIPS 140 options to protect proprietary models and sensitive AI training data.

The V-Series isn't just storage; it's the infrastructure that allows your Generative AI to dream and your AI agents to act.