




## CUSTOMER SUCCESS STORY

# Vultr Selects Supermicro Servers with AMD EPYC™ CPUs and AMD Instinct™ GPUs to Deliver High-Performance Cloud Compute and AI Infrastructure

New Cloud Infrastructure Gives Customers a Full Range of Cloud Compute Services with the Latest Generation of AMD CPUs and GPUs



## INTRODUCTION

Vultr offers a wide range of cloud-based solutions to customers around the world. In addition to virtualized and bare metal systems, Vultr introduces the complete AMD AI infrastructure compute stack, featuring the latest ROCm software. The 4<sup>th</sup> Gen AMD EPYC processors and AMD Instinct MI325X and AMD Instinct MI355X GPUs power the new systems that Vultr is launching. This combination of CPUs and GPUs enables high-performance AI training and inference workloads. Vultr selected Supermicro servers to build out this infrastructure, giving the company a reliable, scalable foundation to expand services quickly and deliver consistent performance to users around the world.

As a global cloud infrastructure provider with 32 cloud data center regions, Vultr fosters innovation throughout the entire AI stack, creating an environment where a ready-to-deploy, flexible infrastructure is optimized for adaptability, energy efficiency, and performance.

As a cloud service provider, our customers' applications on the servers can vary depending on their core business application needs. To support an open AI and cloud computing ecosystem, Vultr works with Cloud Alliance partners and also offers a vibrant marketplace of third-party images and extensions that allow our customers to quickly and easily compose best-of-breed cloud stacks, which can be templated and provisioned globally. Combined with Supermicro's modular hardware platform, this makes it easy for Vultr to expand into new regions—enabling global infrastructure scalability without delay.

## INDUSTRY

Cloud Service Provider

## CHALLENGES

The AI infrastructure market is changing rapidly. Innovations are accelerating new options for building and scaling AI workloads globally. With these rapid developments, there are some key challenges:

- Open stack flexibility without lock-in
- Energy efficiency systems for cost-effective scale
- Optimizing and tuning performance across a unified CPU-GPU stack, with a reliable vendor with hardware optimized for both virtualized and bare metal deployments

Today, those aiming to build and scale AI infrastructure must prioritize openness, flexibility, energy efficiency, and optimized performance in their mission. Supermicro offers a modular building block architecture that provides the flexibility required to adopt emerging technologies swiftly. This system enables Vultr to deploy the latest infrastructure more quickly. Additionally, system-level and power-level design enhancements, including airflow and platform-level tuning, facilitate higher performance per watt, allowing Vultr to deliver more computing capacity for cost-effective scaling to meet the most demanding workloads. Supermicro serves as the foundation for Vultr's ability to achieve this.

Vultr needed to acquire and quickly make several different types of servers available to its demanding customers. Looking at the landscape, Vultr chose several different systems, all containing an AMD EPYC CPU and AMD Instinct GPUs.

## SOLUTION

Vultr chose several different Supermicro AMD-based servers for the recent installation, with the basic specifications listed in the table below.



**AS -8126GS-TNMR**



**AS -8125GS-TNMR2**

CPU	AMD EPYC 9575F	AMD EPYC 9534
Memory	3.0TB	2.3TB
GPUs	AMD Instinct MI325X	AMD Instinct MI300X

The combination of AMD EPYC CPUs with AMD Instinct GPUs in the same system offers an advantage because all of the processing power is engineered to work optimally from the same vendor.

## CHALLENGES

- Acquire State-of-the-Art Servers for new services offered to customers.
- Require high-performing CPUs and GPUs to offer customers additional options for AI training and inference.
- Deploy hardware quickly across multiple global regions while maintaining consistent performance.



## SOLUTION

- Supermicro A+ Servers with AMD EPYC CPUs and AMD Instinct GPUs
- Air-Cooled Systems with the latest AMD Instinct GPUs
- Vultr uses Supermicro servers with AMD EPYC CPUs and Instinct GPUs to deliver fast AI/ML compute, easy deployment, and global scalability

Vultr offers a range of services to end users, from bare-metal access to the systems mentioned above to virtualized instances. Depending on their workloads, users can choose which type of system they want to use, and Vultr engineers work closely with customers to determine the most optimal solution.

The Supermicro servers Vultr selected, with AMD EPYC processors and AMD Instinct GPUs, are easily integrated into an existing data center environment. These servers with system-level and power-level design enhancements, including airflow and platform-level tuning, enable higher performance per watt, allowing Vultr to offer more computing capacity to support cost-effective scaling for the most demanding workloads. Supermicro is the foundation for Vultr's ability to achieve this.



## BENEFITS

By standardizing on Supermicro servers with AMD EPYC™ CPUs and AMD Instinct™ GPUs, Vultr is able to deliver high-performance cloud compute and AI infrastructure that is easy to access, cost-effective, and available globally. Vultr brought these systems online in weeks, not months, due to their use of aircooling, while still acquiring the most performant systems available. The combination of Supermicro hardware and AMD CPUs and GPUs gives Vultr customers access to reliable, high-speed infrastructure for AI training, inference, and compute-intensive applications without needing to manage the hardware themselves. With the new systems, Vultr can offer state-of-the-art solutions to a wide range of customers:

- High-frequency cloud compute instances with the 4th Gen AMD EPYC™ processors, helping customers achieve faster results for compute-heavy tasks.
- Accelerated AI development with quick access to Instinct GPUs and a full AI-native stack on Vultr cloud, spin up instances in minutes to cut training time and accelerate model delivery.
- Reliable, high-performance cloud compute with AMD EPYC CPUs and AMD Instinct GPUs, engineered for consistent training, inference, and HPC performance, with full workload isolation.
- Choice of virtualized or bare metal environments, customers can choose flexible virtual machines or dedicated servers based on workload needs, all running on the same standardized infrastructure.
- Vultr delivers the best price-to-performance cloud infrastructure with full compliance, data residency, and on-demand sovereign cloud, globally available with consistent performance and automated deployment.

## BENEFITS

- Broader AI, HPC, and high-frequency compute offerings with Vultr
- High-End AMD-Based Offerings
- Supermicro's comprehensive solutions for existing data centers





*Vultr is at the forefront of offering high-performance systems for AI and HPC workloads. Our range of systems from Supermicro includes the latest AMD EPYC CPUs and AMD Instinct GPUs. We can give users full control over their computing environment or work closely with them to ensure a seamless experience as requested. The AMD CPUs and GPUs are performing as expected, and we continue to expand our offerings in our 32 data centers worldwide."*

– **Kevin Cochrane**  
CMO, Vultr

## SUPER MICRO COMPUTER, INC.

Supermicro is a global leader in high performance, green computing server technology and innovation. We provide our global customers with application-optimized servers and workstations customized with blade, storage, and GPU solutions. Our products offer proven reliability, superior design, and one of the industry's broadest array of product configurations, to fit all computational need.

[www.supermicro.com](http://www.supermicro.com)

## VULTR

Vultr is on a mission to make high-performance cloud infrastructure easy to use, affordable, and locally accessible for enterprises and AI innovators around the world. Vultr is trusted by hundreds of thousands of active customers across 185 countries for its flexible, scalable, global Cloud Compute, Cloud GPU, Bare Metal, and Cloud Storage solutions. Founded by David Aninowsky and self-funded for over a decade, Vultr has grown to become the world's largest privately held cloud infrastructure company.

<https://www.vultr.com/>



©Super Micro Computer, Inc. Specifications subject to change without notice. All other brands and names are property of their respective owners. All logos, brand names, campaign statements and product images contained herein are copyrighted and may not be reprinted and/or reproduced, in whole or in part, without express written permission by Supermicro Corporate Marketing.