
NVIDIA ACQUIRES RUN:AI

STEVE MCDOWELL, CHIEF ANALYST
APRIL 24, 2024

CONTEXT

NVIDIA announced an [agreement to acquire](#) Run:ai, a startup specializing in chip management and orchestration software based on Kubernetes. The acquisition is part of CEO Jensen Huang's strategy to diversify Nvidia's revenue streams from chips to software.

The acquisition will bolster NVIDIA's DGX Cloud offering, allowing customers to access Run:ai's software through major cloud providers like AWS and Microsoft Azure or by purchasing NVIDIA servers for their data centers.

According to Israeli news outlet Calcalist, the deal is estimated to be valued at \$700 million, though NVIDIA has not officially confirmed the terms of the acquisition.

WHO IS RUN:AI

Run:ai's AI workload management and orchestration for AI applications is built on Kubernetes, a popular open-source system for automating deployment, scaling, and managing containerized applications. Run:ai's optimizes AI computing resources, making it easier for enterprises to manage and scale AI workloads across cloud, edge, and on-premises infrastructure.

The platform provides a single pane of glass for managing shared compute resources, allowing organizations to allocate, monitor, and control AI infrastructure efficiently. It supports key features such as resource pooling, where GPUs can be shared across different tasks and teams, and granular control over quotas, priorities, and clusters.

Run:ai's customers range from large enterprises to research institutions, spanning multiple industries. These organizations use Run:ai's platform to manage data-center-scale GPU clusters, orchestrate AI workloads, and ensure efficient utilization of computing resources. The platform also integrates with third-party AI tools and frameworks, offering flexibility and interoperability within diverse IT environments.

ANALYSIS

NVIDIA's acquisition of Run:AI expands NVIDIA's AI infrastructure offerings, complementing its existing platforms like DGX Cloud and DGX SuperPOD. Integrating Run:AI's capabilities with NVIDIA's offerings will create a seamless user experience, offering flexibility and improved performance in AI workloads.

The acquisition aligns with NVIDIA's strategy to evolve from a chip-focused company to a more comprehensive AI solutions provider. It strengthens NVIDIA's value proposition for enterprise customers and reinforces its competitive edge.

This acquisition makes perfect sense for both the NVIDIA and its customers – it's rare to see such a well-matched acquisition.

© Copyright 2024 NAND Research.

NAND Research is a registered trademark of NAND Research LLC, All Rights Reserved.

This document may not be reproduced, distributed, or modified, in physical or electronic form, without the express written consent of NAND Research. Questions about licensing or use of this document should be directed to info@nand-research.com.

The information contained within this document was believed by NAND Research to be reliable and is provided for informational purposes only. The content may contain technical inaccuracies, omissions, or typographical errors. This document reflects the opinions of NAND Research, which is subject to change. NAND Research does not warrant or otherwise guarantee the accuracy of the information contained within.

NAND Research is a technology-focused industry analyst firm providing research, customer content, market and competitive intelligence, and custom deliverables to technology vendors, investors, and end-customer IT organizations.

Contact NAND Research via email at info@nand-research.com or visit our website at nand-research.com.