



Cloud Tech (Q2 2024):

Funding rebounds amid IBM's
HashiCorp and RHEL AI plays

Keshawa Perera, CFA

SPEEDA
EDGE

What's inside?

This Insight covers quarterly updates on Cloud Technology and focuses on our coverage of the [Cloud-native Tech](#), [Cloud Optimization Tools](#), [Serverless Computing](#), and [Edge Computing](#) industry hubs.

Section 1: Industry outlook

- IBM acquires HashiCorp
- Red Hat launches RHEL AI

Section 2: Funding

- Funding surges, driven by edge computing

Section 3: Product updates

- Edge AI hardware

Section 4: Partnerships

- Simplifying AI edge deployments

Section 5: M&A

- Cloud-native Tech and Serverless Computing drive deals

Key takeaways

Funding

Edge Computing mega rounds stood out: Cloud Tech funding reached its highest in the past three years, raising USD 3.7 billion across 19 funding rounds (+4x QoQ and +1.5x YoY). Edge Computing accounted for ~82% of the total raised, led by [EdgeConneX](#)'s (USD 1.9 billion) and [DataBank](#)'s (USD 725 million) large-debt rounds to expand global data center capacity. Other Edge Computing startups [Yondr](#), [Hailo](#), and [Blaize](#) also raised USD 100+ million each, with a notable shift toward growth-stage rounds. Funding in Cloud-native Tech, Serverless Computing, and Cloud Optimization Tools saw significant gains, with Cloud Optimization Tools experiencing an exceptional 45x YoY growth.

Product updates

Red Hat launched RHEL AI: Red Hat Enterprise Linux AI ([RHEL AI](#)) enables the efficient, flexible deployment of AI models at the edge by providing optimized container images, hybrid cloud capabilities, and seamless integration with enterprise IT operations. It simplifies scaling AI across distributed edge environments, accelerating the adoption of AI-powered applications in resource-constrained, distributed settings.

Focus on edge AI hardware innovation: The edge AI landscape witnessed a flurry of new chip releases, with [EdgeCortex](#), [Hailo](#), and [Kneron](#) introducing advanced edge AI accelerators and servers. Incumbents like Intel and AMD also made significant advancements, with Intel launching edge-optimized [CPUs](#), [GPUs](#), and [FPGAs](#) and AMD unveiling its [Versal adaptive SoC](#). Alongside these hardware innovations, we noted numerous product announcements focused on enhancing Kubernetes adoption and capabilities in Cloud-native Tech.

Partnerships

Simplifying AI deployments and enhancing edge management: Collaboration between incumbents and disruptors aimed to simplify and enhance AI deployments and edge management. Noteworthy partnerships included [Vapor IO and VAST Data](#) integrating edge and core AI systems and Red Hat with [Intel](#), [Cloudera](#), [Veeam](#), and Nutanix to bolster AI, security, and virtualization. Additionally, [Dell partnered with Nvidia](#) to accelerate edge AI adoption.

M&A

IBM acquired HashiCorp: IBM's acquisition of HashiCorp for [USD 6.4 billion](#) was a strategic move to enhance its hybrid cloud and AI capabilities. By integrating HashiCorp's tools into its Red Hat OpenShift platform, IBM plans to offer customers greater flexibility in cloud deployments. The combination of IBM's expertise (including [Red Hat](#)) and HashiCorp's security and management solutions is expected to create a holistic platform tailored for the AI era, strengthening IBM's position in the rapidly evolving cloud and AI market.

Cloud-native Tech and Serverless Computing dominate activity: We tracked nine other M&A deals, primarily in Cloud-native Tech and Serverless Computing. Harness bolstered its product portfolio by [acquiring Split Software](#), further strengthening its acquisition-driven growth strategy. Other notable transactions included EQT's [purchase of WSO2](#), Timescale's [acquisition of PopSQL](#), and SUSE's [buyout of StackState](#).

Industry outlook

As the cloud and Edge Computing industries continue to evolve, several critical developments are shaping the landscape for GenAI adoption. Key players like Red Hat, SUSE, NVIDIA, AMD, and Intel have unveiled innovative products and platforms to enable the deployment of GenAI models across hybrid cloud and edge environments.

Red Hat's RHEL AI and OpenShift AI, for instance, aim to make open-source AI more accessible and enterprise-ready, while partnerships between companies like VaporIO, Dell, and Red Hat are positioning them to capitalize on the GenAI revolution. Additionally, cloud optimization tools are increasingly incorporating AI-based features, such as Red Hat's "Lightspeed" and [CAST AI's](#) AI Advisor, to enhance cloud resource management and governance.

These developments suggest a growing focus on making GenAI more scalable, cost-effective, and user-friendly across diverse cloud and edge computing setups, potentially removing key bottlenecks.

IBM expands cloud and AI capabilities with HashiCorp acquisition and RHEL AI launch

IBM acquired HashiCorp

IBM announced the acquisition of [HashiCorp](#), a leading provider of multi-cloud infrastructure automation solutions, for USD 6.4 billion.

The acquisition aligns with IBM's goal of creating a holistic hybrid cloud platform designed for the AI era. By integrating HashiCorp's cloud-agnostic infrastructure provisioning and management tools, such as Terraform, into its Red Hat OpenShift platform, IBM aims to provide customers with greater flexibility in their cloud deployments. Additionally, HashiCorp's Vault, a security product for managing sensitive data, can be leveraged to drive revenue growth in IBM's existing accounts and introduce the broader HashiCorp portfolio to new customers.

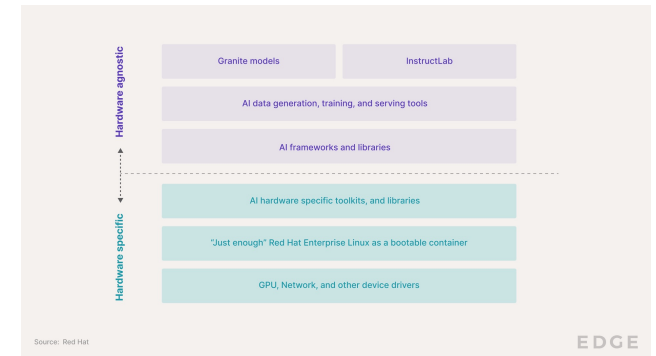
The acquisition also addresses some challenges HashiCorp has faced in its transition to a SaaS-based subscription model, as IBM's resources and expertise can help the company refocus on its core strengths and deliver value to customers more effectively.

Red Hat launched RHEL AI

Red Hat Enterprise Linux AI (RHEL AI) is a new platform from Red Hat that makes it easier for businesses to use and develop AI technologies. It combines Red Hat's popular Enterprise Linux operating system with open-source AI models and tools, allowing companies to run and customize AI applications easily on their existing IT infrastructure.

RHEL AI includes AI models called "Granite" from IBM Research, which can be fine-tuned by domain experts using the InstructLab project. It also provides tools for deploying AI models across on-premises servers and public clouds. By making AI more accessible and customizable, RHEL AI aims to accelerate the adoption of AI across various industries.





Red Hat Enterprise Linux AI development stack



RHEL AI aims to make AI development more accessible, customizable, and scalable

Red Hat's recent announcements provided solutions across the AI application development value chain.

Stages	Open-source AI models	Model fine-tuning and enhancement	Deployment and optimization	Production and scaling
Red Hat solutions	<p>RHEL AI includes the open-source licensed Granite family of large language models (LLMs) developed by IBM Research.</p> <p><i>This provides access to high-quality, transparent AI models as a foundation for applications.</i></p>	<p>Integrates with InstructLab, which provides tools for domain experts to fine-tune and enhance the LLMs.</p> <p><i>This enables customization of the AI models to specific use cases without requiring extensive data science expertise.</i></p>	<p>RHEL AI provides optimized, bootable RHEL container images for deploying the Granite LLMs and InstructLab tools across hybrid cloud environments.</p> <p><i>This supports hardware acceleration from AMD, Intel, and NVIDIA to ensure efficient performance of the AI models in production.</i></p>	<p>Integrates with Red Hat's OpenShift AI platform, which enables running the AI models and InstructLab tools at scale across distributed cluster environments.</p> <p><i>This allows enterprises to more easily manage the lifecycle and integration of AI models with their cloud-native applications.</i></p>

Products	 Granite LLMs	 InstructLab	 RHEL container images	 Red Hat OpenShift AI
	<p>Open-source licensed LLMs from IBM Research for language generation and code assistance</p>	<p>Open-source tools to fine-tune and enhance LLMs using a knowledge and skills taxonomy, based on IBM's LAB methodology</p>	<p>Lightweight, bootable versions of Red Hat Enterprise Linux optimized for running applications in containers</p>	<p>A hybrid machine learning operations platform that enables running RHEL AI models and InstructLab at scale across distributed cluster environments</p>

Funding

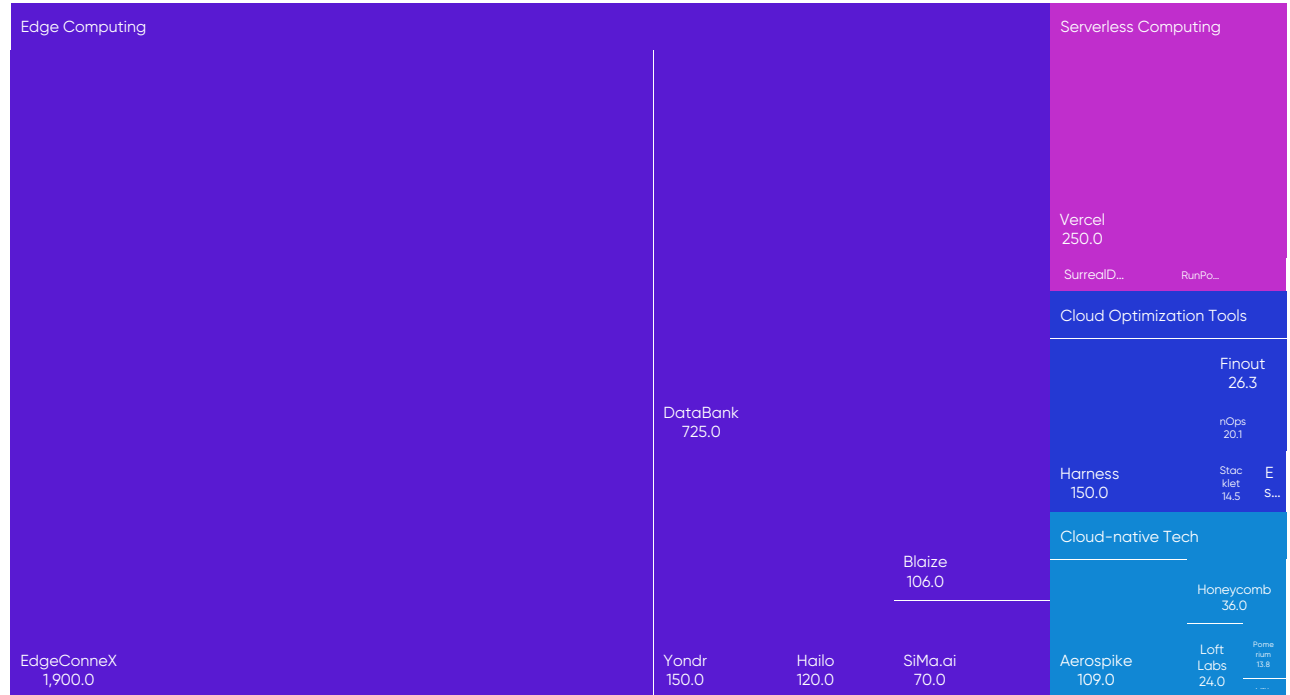
Analyst Take: Cloud Tech startups raised USD 3.7 billion during Q2 2024. This figure significantly exceeded the average quarterly funding of USD 1.3 billion over the past three years. The resurgence in cloud tech funding mirrored the broader [uptick](#) in global venture funding during this period.

Mega rounds from Edge Computing startups stood out, with [EdgeConneX](#)'s USD 1.9 billion and [DataBank](#)'s USD 725 million debt rounds to expand global data center capacity comprising ~70% of total Q2 funding. There was also a notable shift toward growth-stage funding, with [Yondr](#), [Hailo](#) and [Blaize](#) each raising USD 100 million+. Yondr also plans to use the funds for data center expansion, while Hailo, and Blaize focus on advancing GenAI at the edge.

Other notable rounds included [Vercel](#)'s USD 250 million in Series E funding to advance AI-powered website generation and security products and [Harness](#)' USD 150 million in debt funding to enhance its platform with GenAI and acquire [Split Software](#).

Edge Computing mega rounds stand out; Cloud Optimization Tools makes a comeback

Cloud Tech funding mix (Q2 2024)
(USD millions)

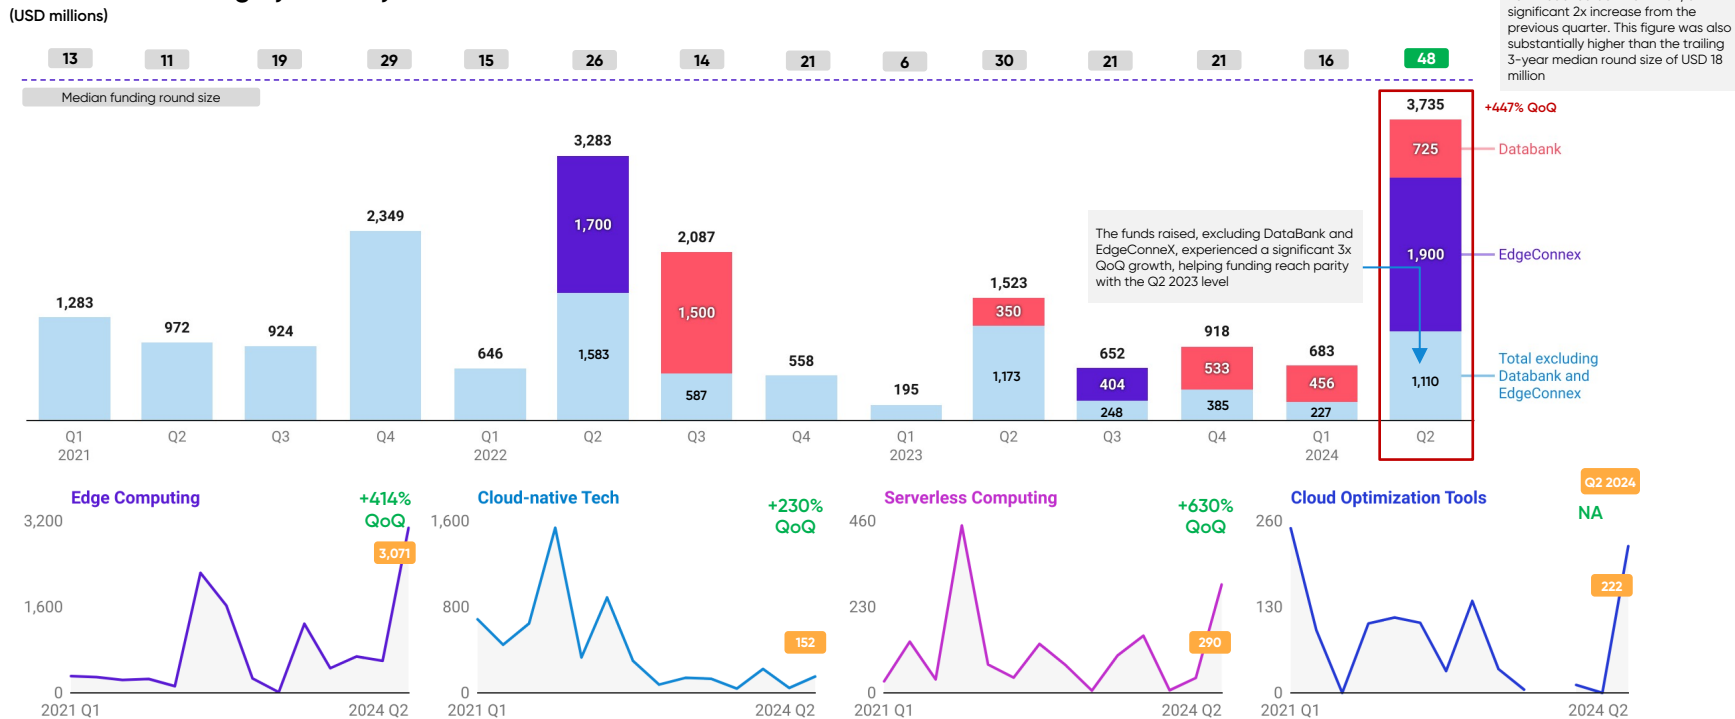


Source: Funding data powered by Crunchbase

Funding

Funding records a three-year high, led by EdgeConnex and Databank

Cloud Tech funding by industry (USD millions)




Funding rounds that involve multiple industry hubs are allocated to the hub with the most significant overlap in scope

Source: SPEEDA Edge • Funding data powered by Crunchbase

Funding

Growth rounds lead the charge, early and seed stages also sees uptick

Funding by stage
(USD millions)



Major funding rounds in Q2 2024

	EdgeConnex Debt	USD 1,900 mn	●
	DataBank Debt	USD 725 mn	●
	Vercel Series E	USD 250 mn	●
	Harness Debt	USD 150 mn	●
	Yondr Debt	USD 150 mn	●

	Finout Series B	USD 26 mn	●
	Loft Labs Series A	USD 24 mn	●
	nOps Series A	USD 20 mn	●
	SurrealDB Series A	USD 20 mn	●
	Stacklet Series B	USD 15 mn	●

	RunPod Seed	USD 22 mn	●
	Espresso AI Seed	USD 11 mn	●

- Cloud-native Tech
- Cloud Optimization Tools
- Edge Computing
- Serverless Computing

Source: Funding data powered by Crunchbase

Product updates

Analyst Take: IBM's (Red Hat) release of RHEL AI underscores its focus on hybrid cloud as a key enabler of AI. Unlike other tech giants developing proprietary approaches, IBM positions itself as an "open hybrid/multi-cloud provider" that can integrate with competitors' cloud services (AWS, GCP, Azure). By combining infrastructure automation, a leading enterprise Linux distribution, and AI capabilities, IBM aims to provide a holistic solution for enterprises looking to harness AI while maintaining control over their IT environments.

Meanwhile, edge AI chips from companies like [EdgeCortex](#), Hailo, [Kneron](#), Intel, and AMD enable powerful GenAI capabilities on resource-constrained edge devices. These chips accelerate LLMs and multi-modal AI workloads, enabling hybrid cloud-edge architectures that leverage the strengths of both. Software-driven approaches, particularly from EdgeCortex, are democratizing GenAI access by allowing the deployment of edge AI capabilities on existing hardware systems while providing flexible, updatable edge AI solutions that can continuously evolve to meet changing needs.

Edge AI ecosystem energized by wave of hardware releases

SPEEDA Edge quarterly reports | Cloud Tech - Q2 2024 GenAI/AI integration

Cloud tech product updates - Q2 2024

	Version upgrade A release of new product that includes improvements, bug fixes, and new features	New product A new product that provides unique functionalities previously not offered	Feature add-on Introducing new functionalities or capabilities to the existing solutions
Cloud-native Tech	OpenStack VMware Redpanda Data Grafana Labs Lighbend Kubernetes SUSE Docker CAST AI Mirantis Isovalent HAProxy Technologies	Nvidia Riverbed Technologies RedHat SUSE Microsoft Kong Tetrate Mirantis Datadog	Rafay Systems Logicmonitor Pure Storage Alphabet SingleStore
Cloud Optimization Tools	CAST AI Upbound	FinOps Foundation	
Edge Computing	Nerdio	Kneron Advantech Gcore Intel Hailo Nokia AMD Edgecortex Vertiv Zenlayer Akamai Flexnode	Fastly Cloudflare StackPath
Serverless Computing		Pinecone SurrealDB ZEDEDA Databank	

Source | SPEEDA Edge **EDGE**

Red Hat democratizes open-source AI with RHEL AI and more; disruptors continue to add GenAI-based features

- Cloud-native Tech
- Cloud Optimization Tools
- Edge Computing
- Serverless Computing

Products to facilitate Gen AI workloads

● Is bolstering its AI and machine learning capabilities across its products. The company has launched [RHEL AI](#), a foundation model platform combining RHEL, open-source Granite models, and InstructLab to simplify working with open-source AI. **Red Hat OpenShift AI** now enables deploying models to edge environments via single-node OpenShift, allowing both predictive and GenAI on a single platform to reduce costs. Additionally, [Podman AI Lab](#) provides developers with local workspace to build GenAI apps using sample models before deployment to OpenShift/Kubernetes.

● Unveiled [NVIDIA NIM](#), a service offering inference microservices for deploying GenAI applications

● Announced its AI strategy and [SUSE AI](#) solutions, a new vendor- and LLM-agnostic GenAI platform

● Launched a series of new [chips and FPGAs](#) aimed at expanding AI capabilities to the network edge, including several edge-optimized CPUs and GPUs and the Agilex 5 Series FPGAs

● Expanded its [Versal adaptive SoC](#) portfolio with AI Edge Series Gen 2 and Prime Series Gen 2

GenAI-based product enhancements

● [Expanded](#) Lightspeed AI across OpenShift and RHEL to boost GenAI-powered user assistance

● Unveiled GenAI-based [upgrades](#) for Nerdio Manager for Enterprise and Nerdio Manager for MSP, enhancing both platforms

● [Introduced](#) new anomaly detection features using machine learning (ML) to identify threats like ransomware, malicious behavior, and denial-of-service attacks






● Its LM Envision platform now includes [LM Cost Optimization](#), an AI-based feature to help manage the complexities and costs of cloud infrastructure





● Launched an AI-powered [observability platform](#). It uses AI to process and analyze information to identify root causes and initiate automated corrections

● Launched an [AI Advisor](#) in CAST Highlight to give users quick insights and recommendations for managing their application portfolios more effectively

EdgeCortix, Hailo, and Kneron propel edge AI forward with innovative deployment platforms





● Edge Computing ● Serverless Computing






Company	Product	Description
	SAKURA-II	<ul style="list-style-type: none">Unveiled the SAKURA-II edge AI accelerator to handle GenAI tasks efficiently and effectively, especially in restrictive environments at the edge
	Hailo-10	<ul style="list-style-type: none">Launched the Hailo-10 high-performance GenAI accelerator for lower-latency AI models to run locally at the edge without relying on cloud-based servers
	KNEO 330	<ul style="list-style-type: none">Launched KNEO 330, its second-generation edge GPT server that can handle AI inference with 48 TOPs of AI computing power
	SmartRow2	<ul style="list-style-type: none">Launched SmartRow2, a new edge data center solution for streamlined edge deployments with increased scalability and flexibility
	ZEDED A Edge Sync	<ul style="list-style-type: none">Launched ZEDED A Edge Sync to address connectivity obstacles in edge deployments that lack consistent access to the cloud

Company	Product	Description
	Gcore Inference at the Edge	<ul style="list-style-type: none">Launched "Gcore Inference at the Edge" to deliver low-latency experiences for AI applications by allowing the distributed deployment of pre-trained machine-learning models to edge inference nodes
	MX Grid	<ul style="list-style-type: none">Introduced MX Grid to enhance operational technology (OT) responsiveness and decision-making in organizations
	Pinecone Serverless	<ul style="list-style-type: none">Released Pinecone Serverless, a serverless vector database for production deployments, enabling complex AI applications
	Surreal Cloud	<ul style="list-style-type: none">Released the beta version of Surreal Cloud, a fully managed version of its scalable database solution for serverless applications

Google, Microsoft, and Mirantis announce major Kubernetes additions alongside new Kubernetes version release

● Cloud-native Tech ● Cloud Optimization Tools

Company	Product	Description
	Google Distributed Cloud	<ul style="list-style-type: none">Integrated Google Kubernetes Engine (GKE) into its Google Distributed Cloud (GDC) platform, allowing the deployment of IT infrastructure setups on any cloud
	Automatic	<ul style="list-style-type: none">Previewed a new service called "Automatic," which is designed to streamline the setup of Kubernetes clusters on the Azure Cloud platform
	k0smotron	<ul style="list-style-type: none">Released k0smotron 1.0, an open-source tool for deploying and managing Kubernetes clusters across any infrastructure—private clouds, public clouds, bare metal, or the edge
	Kubernetes	<ul style="list-style-type: none">Introduced the 1.30 update with various feature updates and enhancements to improve resource utilization and scalabilities, fortify security measures, and provide better development support

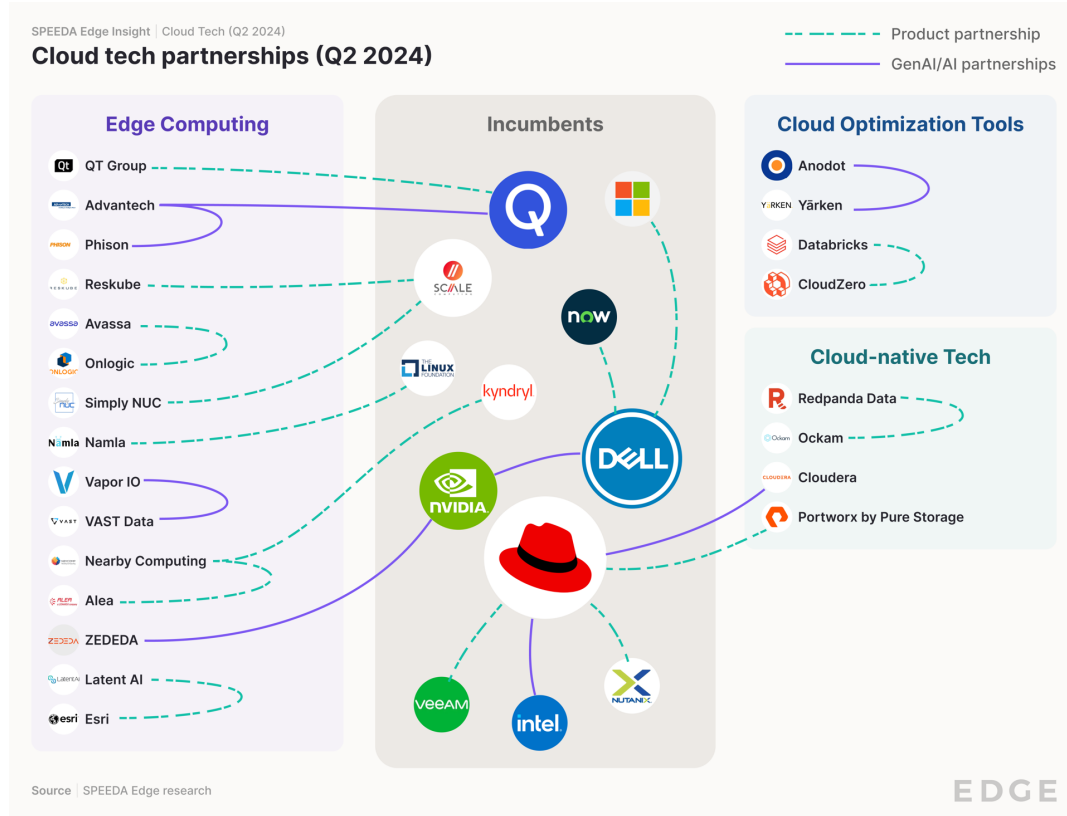
Company	Product	Description
	Kong Konnect	<ul style="list-style-type: none">Commercially launched its Kong Konnect dedicated cloud gateways on Amazon Web Services (AWS)
	Docker Desktop	<ul style="list-style-type: none">Released Docker Desktop 4.29 with Docker socket mount permissions in ECI, advanced error management, Moby 26, and new beta features
	AuthService	<ul style="list-style-type: none">Launched AuthService, an open-source project to simplify the deployment and management of the existing authorization service for the Istio service mesh
	FinOps Foundation (an initiative of the Linux Foundation)	<ul style="list-style-type: none">Launched version 1.0 of its open-source billing standard, FinOps Open Cost and Usage Specification (FOCUS)
	Crossplane	<ul style="list-style-type: none">Released an enterprise-ready Crossplane solution, providing a unified control plane experience with security, automation, and centralized management for cloud-native platforms

Partnerships


Analyst Take: Cloud Tech witnessed several notable partnerships, emphasizing advancements in AI, edge computing, and cloud-native workloads. A key trend is the collaborations aiming to simplify and enhance AI deployments and edge management. Noteworthy partnerships include [Vapor IO](#) and [VAST Data](#), integrating edge and core AI systems, and Red Hat with Intel, [Cloudera](#), [Veeam](#), and [Nutanix](#) to bolster AI, security, and virtualization.

Dell and NVIDIA's partnership aims to accelerate edge AI adoption through automated deployment and management platforms. Dell's NativeEdge now integrates with NVIDIA's AI Enterprise software stack, enabling developers to deploy AI applications at the edge using frameworks like Metropolis for video analytics and Riva for speech. The NativeEdge platform also integrates with [ServiceNow](#)'s Now Platform, providing closed-loop automation for orchestrating, managing, and securing Edge AI workloads. These solutions simplify the customer experience and streamline operations, making it easier for organizations to implement and scale AI capabilities at the edge.

Edge AI ecosystem strengthens through partnerships aimed at GenAI deployment



Red Hat enhances AI and cloud capabilities through partnerships with Intel, Cloudera, and others



Red Hat [partnered](#) with Intel to enhance enterprise AI applications through the product, Red Hat OpenShift AI, leveraging various Intel processor technologies.



Red Hat [partnered](#) with Cloudera to address growing demand for AI workloads by enhancing cloud infrastructure security and robustness. Additionally, Cloudera is transitioning its public cloud infrastructure to Red Hat's newly unveiled RHEL AI platform.

The [partnership](#) between Veeam and Red Hat aims to provide the necessary tools and support for organizations transitioning to OpenShift virtualization, which involves integrating virtual machines and containers.



Nutanix and Red Hat expanded their [collaboration](#) to use Red Hat Enterprise Linux as the foundation for Nutanix's cloud platform, providing enterprise-ready Linux capabilities.

Redpanda [partnered](#) with Ockam to launch Redpanda Connect with Ockam, a zero-trust streaming data platform developed for large-scale enterprise use.



The [partnership](#) between Anodot and YärKEN aims to provide a cloud cost management solution that integrates on-premises and cloud environments, leveraging Anodot's AI-based optimization and YärKEN's platform for aligning tech spend with business objectives.

Dell and partners collaborate to streamline edge computing and AI deployment

- Product partnership
- GenAI/AI partnership

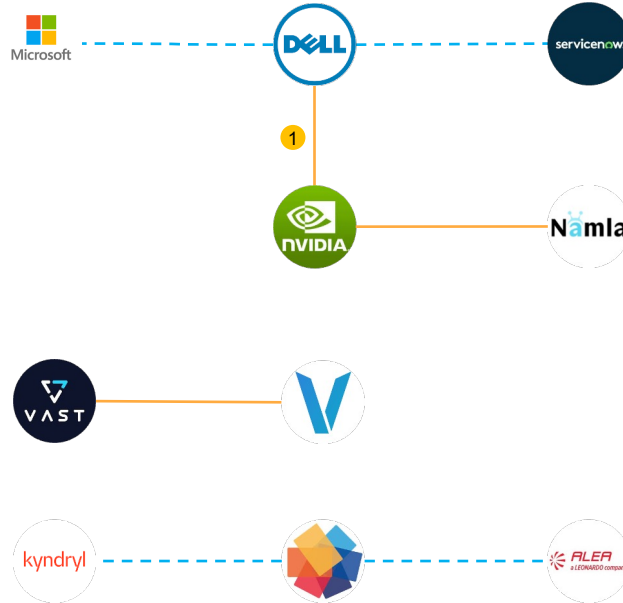
Dell NativeEdge introduced Azure Arc [enablement automation](#) to improve the Azure customer experience and enhance security at the edge, simplify edge operations, and advance AI capabilities.

1

Dell NativeEdge became the first edge orchestration platform that [automates](#) the delivery of NVIDIA AI Enterprise, enabling the deployment of NVIDIA AI frameworks on Dell NativeEdge Endpoints powered by NVIDIA-accelerated computing.

VaporIO [partnered](#) with VAST Data to simplify AI deployments for organizations, merging Vapor IO's Zero Gap AI and the VAST Data platform into an integrated edge-to-core AI system.

Nearby Computing [partnered](#) with Kyndryl to provide combined edge computing services to customers. Kyndryl will use Nearby Computing's NearbyOne platform to simplify the deployment of industry-specific business applications at the edge.



Dell NativeEdge [integrates](#) with ServiceNow's Now Platform to enable businesses to extend IT operations from the data center to the edge, simplifying the orchestration, management, and workflow of edge computing resources.

The [partnership](#) between Namla and NVIDIA aims to improve the scalability and management of edge AI deployments across various sectors by integrating NVIDIA's technology into Namla's edge AI platform.

Nearby Computing and Alea [partnered](#) to automate and orchestrate Alea's mission-critical communications applications through Nearby Computing's edge computing platform.

M&As

Analyst Take: The number of deals during the quarter remained broadly flat (10 vs. 11 in Q1 2024). This was amidst a broader trend where the surge in US M&A activity observed in Q1 2024 experienced a [slowdown](#) during Q2.

IBM's USD 6.4 billion acquisition of HashiCorp strengthens its hybrid cloud capabilities by potentially integrating HashiCorp's multi-cloud infrastructure automation tools with Red Hat offerings.

Harness continued its M&A growth strategy, acquiring Split Software to enhance its products. This move comes after Harness' competitor, [Octopus Deploy](#), acquired the struggling Kubernetes continuous delivery expert [Codefresh](#) in Q1 2024. Moreover, Suse, [Timescale](#), [Cloudflare](#), [Finout](#), and Vercel enhanced their tech stacks and offerings through acquisitions.

IBM bolsters cloud capabilities with HashiCorp acquisition during steady M&A period

Deal spotlight



IBM acquired **HashiCorp** for USD 6.4 billion.

Announced date: April 24, 2024

Industry hubs: Cloud-native Tech*, Cloud Optimization Tools



Harness acquired **Split Software** to integrate its capabilities and strengthen Harness's offerings.

Announced date: May 29, 2024

Industry hubs: Cloud Optimization Tools*, Cloud-native Tech

- IBM acquired HashiCorp for USD 35 per share, a significant premium over HashiCorp's pre-announcement trading price of around USD 25.
- HashiCorp's product pillars (Terraform, Vault) align well with Red Hat's (IBM subsidiary) business and IBM's focus on the hybrid/multi-cloud market.
- In 2023, HashiCorp changed its business model, including a [shift](#) away from open-source licensing, which drew [attention](#) and new competitors.
- The transition to a SaaS-style cloud subscription business was taking longer than planned. The company's gross margins on cloud services were only [around 60%](#) (FY 2024), while its traditional support revenue had higher margins at 86%.
- Harness, a cloud-based software delivery platform, entered an agreement to acquire Split Software Inc., a company that specializes in feature flag management and experimentation.
- The acquisition will allow Harness to provide a more robust software delivery platform, including the ability to run A/B tests and measure the performance of new features.
- This is Harness' second acquisition in 2024 and its seventh overall. In [January 2024](#), the company acquired Armory, a continuous delivery platform company, for USD 7 million. This acquisition announcement comes just weeks after Harness secured a [USD 150 million](#) debt financing round, backed by Silicon Valley Bank and Hercules Capital.

Activity centered around Cloud-native Tech and Serverless Computing

Cloud-native Tech

IBM [acquired](#) HashiCorp, a leading multi-cloud infrastructure automation software provider, for USD 6.4 billion. IBM intends to bolster its cloud offerings with the acquisition of HashiCorp



Swedish investment giant EQT [acquired](#) API management and identity access management company WSO2 in a deal valued at more than USD 600 million.



SUSE, an open-source enterprise provider, [acquired](#) StackState's full-stack observability platform to integrate its capabilities into Rancher Prime.




Serverless Computing

Cloudflare [acquired](#) Baseline, a serverless application observability platform, to expand its capabilities and further enhance the developer experience for building, shipping, and troubleshooting applications on the Cloudflare platform.



Vercel a provider of a serverless web and app development platform [acquired](#) ModelFusion.dev, expanding the capabilities of its Vercel AI SDK 3.1 to further position the company as a leader in the enterprise AI development space.



Cloud-optimization Tools

Finout, a cloud cost observability platform provider, [acquired](#) Cloudthread, a cloud analytics platform, to enhance Finout's cloud cost optimization and FinOps capabilities.



Timescale, a cloud database company, has [acquired](#) PopSQL, an SQL editor and collaboration tool, to offer Timescale customers an intuitive interface for database querying, dashboard creation, and collaborative data analysis.



Redpanda, a streaming data platform, has [acquired](#) Benthos, an open-source stream-processing platform, to bolster its end-to-end streaming capabilities for data-intensive applications.



OpenAI has [acquired](#) Rockset, a real-time analytics database, to enhance its retrieval infrastructure and enable users, developers, and enterprises to better leverage their own data for more helpful AI applications.



Harness, a cloud-based software delivery platform, [acquired](#) Split Software, a feature management and experimentation company, to integrate its capabilities and strengthen Harness' offerings.



©2025 Uzabase, Inc. All Rights Reserved. The information contained herein (1) is proprietary to Uzabase Inc. and/or its content providers; (2) may not be copied or distributed; and (3) is not warranted to be accurate, complete, or timely. Neither Uzabase Inc. nor its content providers are responsible for any damages or losses arising from any use of this information.