



Strategic Partnership Announcement

June 16th, 2026

Disclaimer

Forward-Looking Statements

Rackspace Technology has made statements in this presentation and other reports, filings, and other public written and verbal announcements that are forward-looking and therefore subject to risks and uncertainties. All statements, other than statements of historical fact, included in this presentation are, or could be, “forward-looking statements” within the meaning of the Private Securities Litigation Reform Act of 1995 and are made in reliance on the safe harbor protections provided thereunder. These forward-looking statements relate to anticipated financial performance, including our financial outlook and expected future results, management’s plans and objectives for future operations, business prospects, outcome of regulatory proceedings, market conditions, and other matters.

In particular, our discussion today will include forward-looking statements regarding our recently announced definitive agreement with AMD, including without limitation, the ability to dedicate, maintain and make available an aggregate of 30 megawatts of AMD Products contemplated by the GPUaaS Agreement, which may not be achieved in full or at all, or may be achieved on a materially different timeline; the anticipated benefits and performance of GPU and CPU compute deployments; the expected delivery of Enterprise AI Cloud, Enterprise Inference Engine, Inference as a Service, and Bare Metal AMD Instinct capabilities; anticipated end-customer demand; the expected commercial and financial benefits of the collaboration to each company; and the parties’ respective outlooks on the AI industry. While the parties have executed a definitive agreement establishing a commercial framework for the collaboration, individual deployment authorizations are subject to separate execution and certain commercial terms, including pricing and financial parameters, remain subject to further agreement between the parties. AMD has no obligation to agree to any particular deployment as being within the scope of the framework. Any third-party financing required to implement planned deployments is subject to availability on terms acceptable to the Company. The GPUaaS Agreement is subject to certain financing, operational, and legal conditions and provides AMD with a right of first refusal that may affect the Company’s flexibility in selling capacity to third parties. There can be no assurance that deployments will occur on the anticipated timeline, that financing will be obtained, that AMD will agree to future deployments, or that the anticipated benefits of the collaboration will be realized. Deployments are subject to availability of and lead times for AMD Products from third party original equipment manufacturers.

Our discussion will also include forward-looking statements relating to the Company’s workforce realignment plan, including without limitation, the expected number of employees affected, the anticipated timing and implementation of the reduction in force across jurisdictions, the estimated one-time expenses associated with the workforce realignment plan, and the anticipated gross annualized savings and reinvestment plans. Actual expenses, savings and reinvestments may differ materially from these estimates as a result of changes in the scope, timing or implementation of the workforce realignment plan, variations in severance obligations across jurisdictions, the timing of employee exits, regulatory or legal requirements applicable in certain jurisdictions, threatened or actual litigation, and other factors. There can be no assurance that the Company will realize the anticipated savings from the workforce realignment plan within the expected timeframe, or at all. The Company undertakes no obligation to update or revise these forward-looking statements except as required by law.

Market Overview

Infrastructure without an operating model is capacity, an operating model without committed infrastructure is an aspiration

Addressing the Market Gap

- Enterprise AI today puts the full operational burden on the customer — companies procure GPU hours and then independently manage integration, security, governance, and accountability
- For regulated industries — financial services, healthcare, government — that model fails. Governance, data sovereignty, and operational continuity are not selling points — they are the price of entry
- No provider has delivered a full governed stack from silicon through business outcomes under a single accountable operator. That is the gap Rackspace is built to fill

This is the right time

Production at Scale

Enterprise AI is moving beyond pilots into real workloads requiring dedicated, performant, governed infrastructure at scale

Regulated Demand

Financial services, healthcare, and government represent massive underserved demand for managed, governed AI environments

Structural Gap

Hyperscalers deliver compute. System integrators deliver services. Neither is accountable for governed AI in production, end to end

Rackspace x AMD: The Governed Enterprise AI Operator

The Solution

- Rackspace and AMD define a new category: governed enterprise AI infrastructure purpose-built for regulated industries where sovereignty, auditability, and deterministic performance are non-negotiable
- Rackspace operates the full stack from silicon to outcomes as the single accountable operator — for availability, scaling, performance, and compliance — with AMD supplying the accelerated compute foundation

What Each Partner Brings

rackspace
technology.

AMD 

Core Asset and Differentiation	Governed, operator-led, model-agnostic; customers choose the model, Rackspace keeps it running and governed	Next-gen silicon roadmap and portfolio (MI355X, MI350P)
Channel	Deep enterprise relationships across regulated verticals and key AI partnerships	Strategic partner ecosystem and co-marketing reach
Infrastructure	Existing global DC footprint with capacity available for near-term deployment, scaling to a total committed footprint of 30 MW	Access to next generation Instinct, EPYC and AMD Enterprise AI roadmap

The Rackspace Value Proposition

Four structural advantages that made Rackspace the right operator partner

Global Data Center Footprint

Capacity available for near-term deployment.

Rackspace's global data center footprint includes 30 megawatts of capacity available for phased deployments of AMD AI compute beginning in late 2026 through 2028. No material new construction contemplated in the near-term.

Regulated Industry Presence

One operator, model-agnostic and accountable for the full stack

More than two decades operating regulated, mission-critical workloads – proven with Epic at scale, sovereign cloud in the UK & US, and a deep installed base of regulated enterprise customers. Stack is model-agnostic: with a single orchestration layer & context-aware inference keeping knowledge intact.

Managed Infrastructure Expertise

Deep operational expertise at enterprise scale.

Deep operational expertise in managed infrastructure at enterprise scale where failure is more than just a performance issue. That discipline is not assembled quickly and it is not replicated by operators whose accountability ends at the infrastructure perimeter.

Operator Model, Not a Reseller

We operate compute. We own the outcome.

We do not simply resell compute — we operate it and remain accountable for the outcome. That governed, operator-led model is what makes Rackspace the right partner to bring AMD Instinct into regulated enterprise production and what differentiates us from every hyperscaler and GPU cloud provider.

GTM Strategy and Product Evolution

Go-to-Market Motions

Activating the Installed Base and Building the Ecosystem

1 Installed Base Activation

Near-term demand from Rackspace's existing enterprise customers

2 Co-Sell with Partners

RXT account teams co-sell AMD GPUs with AMD, Palantir and Uniphore into installed base

3 New-Logo Acquisition

Bring new logos via FDE-led sales motions as well as traditional sales channel

4 AMD Demand Funnel

Enterprise GPU demand redirected to RXT with co-branded campaigns

From Silicon to Outcomes

More Value at Every Layer

Tier 4: Enterprise AI Cloud

Full-stack margin capture

Tier 3: Enterprise Inference Engine

Higher-margin managed service

Tier 2: Inference as a Service

Software + ops margin layer

Tier 1: Bare Metal

Foundation: infrastructure-as-a-service

Launch footprint: 30 MW of dedicated capacity for AMD Instinct GPUs and AMD EPYC CPUs

Partnership Ecosystem to Deliver Enterprise AI



20+ Year Partnership with Dell Infrastructure operating within Rackspace Data Centers



Q1-2025 Partnership with Rubrik to redefine the Cyber Resiliency layer for regulated and sovereign cloud environments



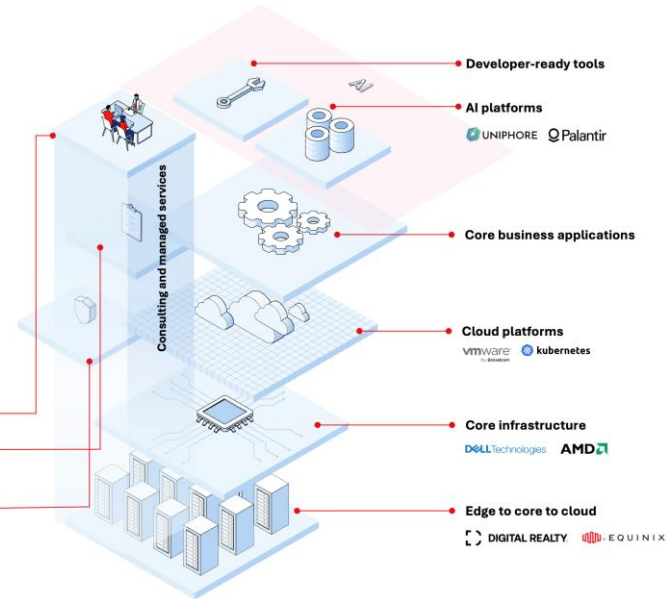
AMD GPU and CPU silicon underpins Rackspace's AI infrastructure, enabling bare metal GPUaaS, managed Inferencing & Fine-Tuning, and AI Platforms



VMware as a core workload and application control plane with Rackspace as one of the earliest and largest CSPs for more than 15 years



Q1-2026 Palantir and Uniphore establish the AI Platform layer running at Rackspace to help increase time to value for Enterprise AI adoption



Conclusion

rackspace
technology™

AMD 

Category of one: the governed enterprise AI operator
— accountable from silicon to outcomes

Proven by what we already run: Epic at scale, sovereign cloud, strategic partnerships with Palantir and Uniphore

AMD underneath: the first committed infrastructure anchoring the governed AI stack

One operator. One accountable stack. Built for the industries that cannot afford to get it wrong

The category is defined. The infrastructure is committed. The operator is Rackspace.