

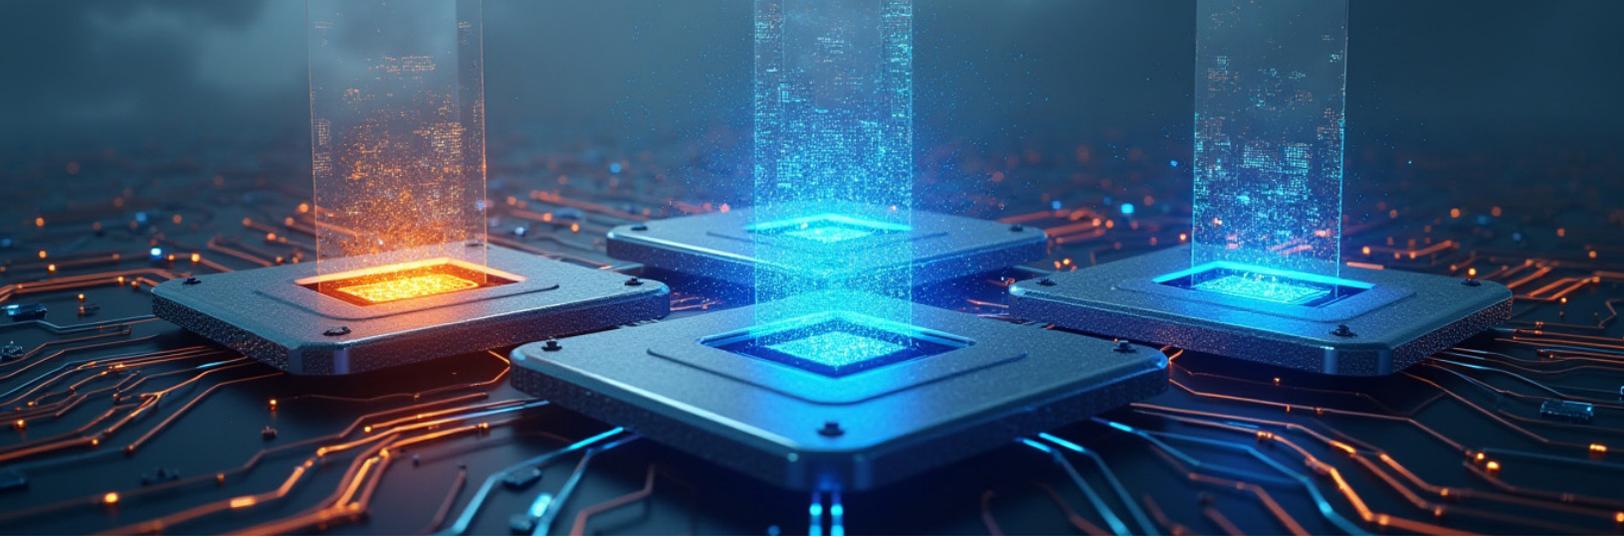
WHITE PAPER

On-Premises AI Deployment Alternatives: A Comprehensive and Strategic Analysis of Lenovo, Dell, and HPE

APRIL 2025

Authors:

Steven Dickens
CEO and Principal Analyst



Executive Summary: The On-Premises AI Imperative

The meteoric rise of AI in the last couple of years as a transformative force in enterprise operations has thrust infrastructure decisions into the spotlight. While public cloud solutions dominate the zeitgeist relating to AI, a tectonic shift is underway as enterprises increasingly turn to on-premises deployments to meet stringent demands for data sovereignty, latency, security, and cost control. This report delivers an exhaustive comparative analysis of the AI deployment offerings from three titans of enterprise infrastructure; Lenovo, Dell, and Hewlett Packard Enterprise (HPE), tailored explicitly for on-premises data centers or co-location facilities. HyperFRAME's evaluation hinges on two pivotal axes: time to value, the speed at which enterprises can extract tangible benefits from their AI investments, and flexibility of consumption, the adaptability of scaling and payment models to align with dynamic business needs.

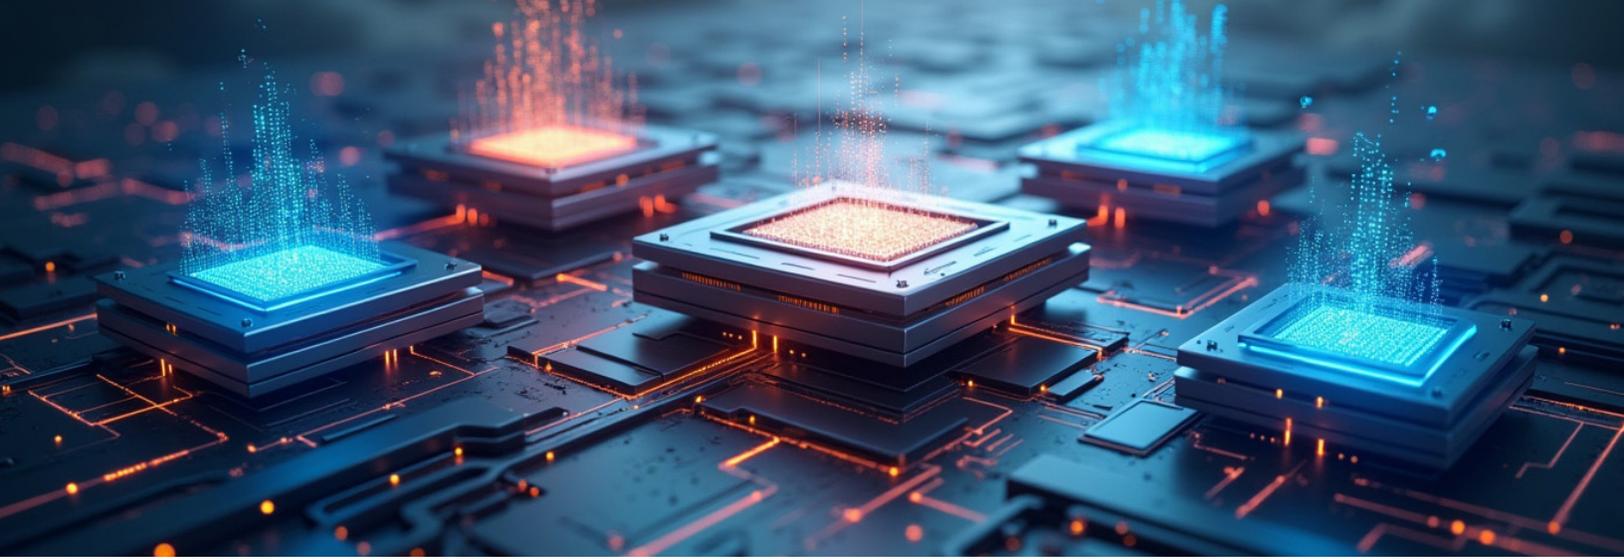
This analysis unveils distinct strategic profiles for each vendor. Lenovo emerges as a pacesetter in rapid deployment and consumption agility, leveraging its Hybrid AI Advantage and TruScale models. Dell carves a path of ecosystem breadth and hardware optimization, bolstered by its APEX as-a-service framework. HPE, meanwhile, stakes its claim with unparalleled deployment velocity through its Private Cloud AI and GreenLake offerings. This report goes beyond surface-level comparisons, dissecting the architectural underpinnings, service ecosystems, and strategic implications to empower enterprise architects with actionable intelligence for navigating the AI infrastructure landscape.

The Strategic Context: Why On-Premises AI Matters Now

Before delving into the vendor-specific analysis, it's worth stepping back to understand the broader forces propelling on-premises AI into the enterprise mainstream. The public cloud's allure, scalable compute, managed services, and rapid provisioning, remains undeniable. Yet, as AI workloads mature, enterprises are confronting a triad of challenges: regulatory compliance, escalating cloud costs, and the need for low-latency inference at the edge. On-premises deployments address these head-on, offering granular control over data, predictable economics, and the ability to tailor infrastructure to specific workloads.

Put another way HyperFRAME firmly believes that two truths can be true at the same time - Public cloud can be the right solution for certain workloads, and on-premises can be a good fit for other workloads - It depends. Non-functional requirements such as; performance, availability, scalability, security, data sovereignty, regulatory framework and environmental concerns all play a role in deciding where a workload should reside.

Moreover, the rise of agentic AI, systems that autonomously reason, plan, and act, underscores the need for robust, localized compute power. This isn't just about running models; it's about orchestrating complex, multimodal workflows that span text, video, and real-time sensor data. Vendors like Lenovo, Dell, and HPE are stepping into this breach, not merely as hardware providers but as strategic partners in the AI journey with templated, reference architectures that are often pre-tested. Their offerings reflect a convergence of hardware innovation, software integration, and consumption flexibility, a trifecta that defines the modern AI infrastructure playbook.



In-Depth Analysis of Lenovo AI Solutions

Lenovo AI Factory: Hybrid AI Advantage with NVIDIA

Lenovo's foray into on-premises AI is anchored by its Hybrid AI Advantage with NVIDIA, a full-stack solution born from a deep partnership with the GPU giant. This isn't a mere hardware play; it's a strategic bid to redefine how enterprises operationalize AI. By marrying Lenovo's ThinkSystem and ThinkEdge server platforms and services with NVIDIA's H100/H200 NVL and L40S GPUs, alongside the NVIDIA AI Enterprise software suite and partner, ISV options, Lenovo delivers a cohesive ecosystem that straddles on-premises and hybrid paradigms. The "Hybrid" moniker isn't just marketing flair, it signals a deliberate design to accommodate enterprises juggling sensitive local workloads with cloud-augmented scalability. When coupled with Lenovo's storied history in liquid cooled systems, under the Neptune brand name, Lenovo has a comprehensive array of AI-ready systems that span from the edge to the datacenter.

At the heart of Lenovo's strategy lies a focus on agentic AI, a term that's rapidly becoming the litmus test for next-generation automation. These aren't your grandfather's predictive models; they're autonomous systems capable of reasoning through complex tasks, think supply chain optimization or real-time customer service orchestration. Lenovo's platforms are engineered to slash deployment timelines for such agents, a move that aligns with the growing enterprise appetite for AI that doesn't just analyze but acts. Performance metrics from Lenovo's own claims paint a compelling picture: content creation accelerated by up to 8x, customer service efficiency boosted by 50%, and knowledge assistants driving measurable productivity gains. These aren't abstract promises, they're quantifiable markers of time to value, a critical lens for evaluating ROI in the AI era. Enterprises need to validate these claims in their own environment but the benefits are tangible and real.

Lenovo's hardware innovation doesn't stop at GPUs. The ThinkSystem SC777 V4 Neptune server, with its liquid-cooling architecture, tackles the thermal challenges of dense AI deployments head-on. This isn't just about keeping servers cool; it's about enabling sustainable, high-performance data centers that can scale without breaking the bank on power bills. Add support for NVIDIA's Blackwell Ultra architecture, and Lenovo positions itself at the bleeding edge of AI compute, a strategic differentiator for enterprises eyeing long-term competitiveness.

Reference Architecture: Validated Design for AI Infrastructure

Lenovo's Validated Designs for AI Infrastructure on ThinkSystem Servers is less a product and more a philosophy, a meticulously crafted blueprint tested in Lenovo AI Labs that spans the AI stack lifecycle from compute, storage, network, OS, orchestration, software to data ingestion to and inference with AI ISV options. This isn't a one-size-fits-all template; it's a living framework, updated to incorporate big data storage and edge inference capabilities. Also a point of note is that Lenovo has a robust range of services that can be wrapped around these systems. They extend from the obvious support and break-fix level as a service option with Lenovo TruScale right up to highly skilled consultants who can help with every step of the AI journey from customizing NVIDIA NIMs and accelerating AI

agents leveraging Lenovo AI Library, design to deployment. For enterprises, this translates to reduced risk and accelerated deployment. The validated design isn't just about hardware specs; it's about providing a roadmap that mitigates the guesswork of AI adoption.

A standout feature is Lenovo's collaboration with Centific on the AI Data Foundry, a platform integrated into the validated design. This isn't a trivial add-on, it's a recognition that data is the lifeblood of AI, particularly for multimodal workloads. By streamlining ingestion, transformation, and deployment of diverse data types (video, text, speech), the AI Data Foundry lowers the barrier to leveraging unstructured data, a perennial pain point for enterprises. Partnerships with NetApp further amplify this, weaving Lenovo's solutions into broader storage ecosystems and enhancing interoperability. This open approach signals Lenovo's intent to play nicely with existing infrastructure, a boon for enterprises wary of vendor lock-in.

Services: Accelerating Adoption with Precision

Lenovo's service portfolio is where its on-premises vision comes into sharp focus. The AI Fast Start program is a standout, promising tangible business value within 90 days. This isn't a vague consultancy pitch, it's a structured engagement that identifies use cases, leverages enterprise data, and deploys production-ready solutions. Case studies bear this out: SAP's AI avatar for visitor engagement, ST Logistics' warehouse automation slashing order processing times, and AISHA's MRI analysis accelerated by over 99%. These aren't isolated wins, they're proof points of Lenovo's ability to deliver rapid, industry-agnostic value.

Complementing this is TruScale, Lenovo's as-a-service juggernaut. Offering Infrastructure-as-a-Service (IaaS) and GPU-as-a-Service (GPUaaS), TruScale flips the traditional CAPEX model on its head. Enterprises pay for what they use, scaling compute and GPU resources dynamically via Lenovo's Intelligent Computing Orchestration (LiCO) software. This isn't just about cost, it's about agility, enabling enterprises to match infrastructure to workload flux without overprovisioning. Whether deployed in a data center or co-lo facility, TruScale's metered approach redefines flexibility of consumption.

The AI Innovators Program rounds out Lenovo's ecosystem, delivering over 165 pre-validated solutions with partners like Intel, NVIDIA, and ISVs. Tailored for on-premises use cases, think manufacturing or smart cities, this program slashes deployment timelines by offering turnkey configurations. It's a pragmatic play, ensuring enterprises don't reinvent the wheel for every AI initiative.

Time to Value: Speed as a Competitive Edge

Lenovo's obsession with time to value is palpable and much needed as enterprises wrestle with quantifying value. AI Fast Start's 90-day ROI promise, backed by real-world successes, sets a high bar. The AI Innovators Program amplifies this by pre-validating solutions, while the Validated Design streamlines planning and execution. Together, these elements compress the journey from concept to production, a strategic advantage in an era where AI agility can make or break market leadership.

Flexibility of Consumption: Modular and Metered

TruScale for GPU is the latest evolution in Lenovo's consumption story, offering pay-as-you-go access to infrastructure hardware and the ability to optimize usage of NVIDIA GPUs with granular control. Paired with a modular hardware portfolio, from ThinkEdge for edge AI to SR675 V3 for large-scale training, Lenovo lets enterprises right-size their infrastructure. The Hybrid AI model adds another layer, blending on-premises control with cloud optionality. This isn't just flexibility, its strategic empowerment for enterprises navigating data sovereignty and workload variability.



In-Depth Analysis of Dell AI Solutions

Dell AI Factory with NVIDIA: Ecosystem Breadth Meets Hardware Might

Dell's AI Factory with NVIDIA is a tour de force of integration, blending Dell's infrastructure prowess with NVIDIA's accelerated computing stack. This isn't a bolted-on partnership, it's a two-decade collaboration distilled into a comprehensive AI platform. From PowerEdge servers to NVIDIA's GPUs, AI Enterprise software, and Inference Microservices (NIM), Dell's offering spans the AI lifecycle with surgical precision. The focus? Turning ideas into outcomes by weaving existing data and workflows into a scalable AI framework.

Dell's emphasis on data as the foundation sets it apart. Recognizing that AI success hinges on quality inputs, Dell and NVIDIA tackle the messy realities of data management, discovery, preparation, analytics, ensuring enterprises are GenAI-ready. This should come as no surprise given the EMC heritage that sits with Dell to this day. Hardware like the PowerEdge XE9680L, with liquid cooling and Blackwell GPU support, underscores Dell's commitment to cutting-edge compute. Cost efficiency claims, up to 75% savings over public cloud for LLM inferencing, further bolster the value proposition, hitting the "time to value" bullseye.

Reference Architecture: Leveraging NVIDIA's Playbook

Dell leans on NVIDIA Enterprise Reference Architectures (RAs) to deliver battle-tested blueprints for AI factories. These aren't generic guides, they're detailed prescriptions for server clusters, networking, and scalability, aligned with industry best practices. By standardizing on NVIDIA's RAs, Dell reduces deployment friction and enhances interoperability, a pragmatic move for enterprises integrating multivendor ecosystems.

Services: Expertise at Every Step

Dell's Professional Services are a force multiplier, spanning planning, design, deployment, and testing. Tailored to specific use cases, these services leverage Dell's AI Solutions Explorer to recommend configurations, a data-driven approach to optimization. The Dell Enterprise Hub adds another dimension, enabling on-premises deployment of open-source models with multicloud monitoring. It's a nod to enterprises craving control without sacrificing visibility.

Time to Value: Fast-Tracking ROI

Dell's PowerEdge AI servers promise to "turn ideas into action faster," with pre-configured systems minimizing setup time. NVIDIA NIM Blueprints accelerate use-case deployment, while cost savings on LLM inferencing turbocharge ROI. Professional Services further shorten the timeline, drawing on Dell's ecosystem to bridge planning and production.

Flexibility of Consumption: APEX and Beyond

The APEX portfolio is Dell's consumption cornerstone, offering subscription and as-a-service models for compute, storage, and networking. A sprawling hardware lineup, from AI PCs to high-density servers, ensures workload alignment, while modular cooling (air, liquid, immersion) enhances adaptability. An open ecosystem, integrating open-source tools and NVIDIA's stack, rounds out Dell's flexibility play, giving enterprises room to maneuver.



In-Depth Analysis of HPE AI Solutions

HPE Private Cloud AI: Turnkey Velocity

HPE Private Cloud AI, developed in close collaboration with NVIDIA, exemplifies a streamlined approach to enterprise-grade AI deployment, blending simplicity with exceptional speed. Marketed as a turnkey solution tailored for on-premises or co-location environments, this offering integrates HPE's robust ProLiant servers and high-performance Cray supercomputers with NVIDIA's cutting-edge AI Data Platform. The result is a cohesive, pre-configured system designed to slash deployment timelines to "just a few hours", a bold assertion that challenges traditional expectations and redefines the concept of time to value in AI infrastructure. HPE's private cloud philosophy emphasizes ease of management, reducing operational complexity for IT teams, while powerful integrations such as NVIDIA AI-Q Blueprints and NIM microservices provide pre-built frameworks and tools to accelerate adoption across diverse AI workloads. Additionally, the HPE Data Fabric Software plays a pivotal role by enabling seamless, unified data access across edge, on-premises, and cloud environments, an essential capability for organizations navigating the complexities of hybrid AI strategies and seeking to maximize data utility.

Reference Architecture: NVIDIA Alignment

Much like its competitor Dell, HPE leverages NVIDIA Enterprise Reference Architectures (RAs) to deliver meticulously pre-tested configurations that are fine-tuned for critical AI functions such as model training, inference, and robust security. This strategic alignment with NVIDIA's ecosystem ensures not only scalability to meet growing enterprise demands but also broad compatibility with a wide array of AI tools and frameworks, an indispensable feature for businesses aiming to establish resilient, future-proof AI foundations. By adhering to these validated designs, HPE minimizes deployment risks and enhances performance optimization, offering enterprises a reliable blueprint to build and expand their AI capabilities with confidence over time.

Services: Data-Driven Acceleration

HPE's Data Services extend far beyond hardware, providing a comprehensive suite of offerings that includes hands-on workshops, detailed readiness assessments, and expert-led implementation support. These services are designed to transform fragmented or outdated data environments into AI-ready ecosystems capable of supporting advanced analytics and machine learning initiatives. The HPE AI Factory, a cornerstone of this portfolio, reinforces HPE's claim of 24-hour deployment by pairing rapid setup with access to seasoned professionals who guide customers through every stage of the process, from planning to production. This combination of speed and expertise positions HPE as a leader in delivering accelerated AI solutions, particularly for organizations eager to transition quickly from concept to operational value.

Time to Value: Speed Meets Scale

HPE's promise of a 24-hour deployment window stands out as a transformative advantage, dramatically compressing the timeline from initial setup to production-ready AI systems. This rapid rollout is complemented by an iterative scaling model that allows enterprises

to start with modest deployments and expand incrementally as needs evolve, effectively minimizing upfront risk while maximizing flexibility. Coupled with HPE's Data Services, which help prioritize high-impact use cases and streamline workflows, this approach ensures a swift return on investment, catering to businesses that need both immediate results and long-term scalability to stay competitive in a fast-moving AI landscape.

Flexibility of Consumption: GreenLake's Edge

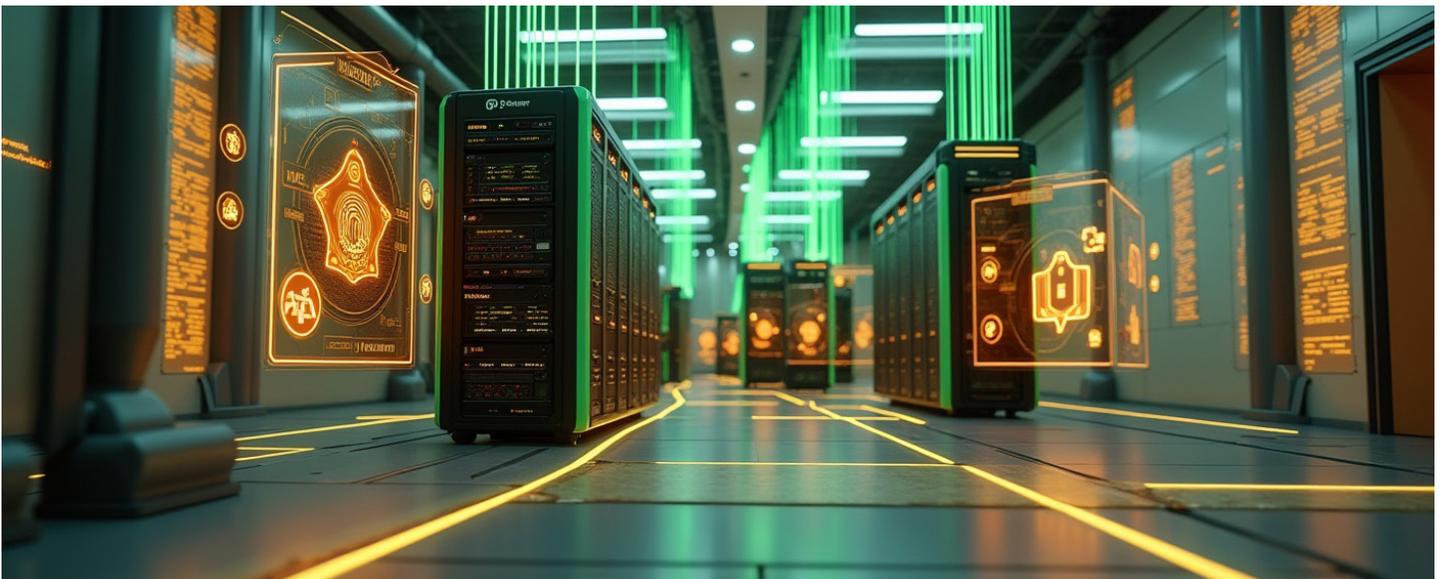
HPE pioneered consumption-based infrastructure with its GreenLake platform, and this legacy continues to provide a competitive edge in the AI space. GreenLake offers pay-per-use flexibility, empowering enterprises to dynamically adjust compute, storage, and networking resources in response to fluctuating workloads or budget constraints. The modular design of HPE's Private Cloud AI supports a hybrid economic model, enabling cost-efficient operations across on-premises and cloud environments. Further enhancing this adaptability is a rich software ecosystem, featuring NVIDIA's AI tools alongside integrations with independent software vendors (ISVs), which ensures that the platform can handle a diverse range of workloads, from generative AI to data-intensive analytics. This combination of economic agility and workload versatility makes HPE's solution particularly appealing to organizations seeking to balance innovation with operational efficiency in an increasingly AI-driven world.

Rapid Start Services

HPE offers a comprehensive 90-day onboarding services engagement as part of its Private Cloud AI solution, designed to accelerate the adoption of generative AI for enterprises. This service ensures a smooth transition by providing expert guidance on deployment, configuration, and optimization of AI workloads within a secure private cloud environment. Over the 90 days, HPE's team collaborates closely with organizations to tailor the platform to their specific needs, integrating it with existing infrastructure while addressing security, compliance, and scalability requirements. The engagement includes hands-on support for setting up AI models, fine-tuning workflows, and training staff to manage the system effectively. By leveraging HPE's expertise and NVIDIA's AI technology, businesses can quickly operationalize AI solutions, minimize downtime, and achieve faster time-to-value.

Beyond NVIDIA Software Value Add

HPE further enhances its AI portfolio with AI Essentials, a software suite that complements NVIDIA AI Software. AI Essentials provides operationalization and AI workflow management, offering a unified control plane for adaptable, compliant, and reproducible AI pipelines. This approach ensures enterprises can efficiently manage AI workloads, from model training to deployment, across hybrid environments. By combining NVIDIA's advanced AI capabilities with HPE's infrastructure and software, HPE empowers organizations to accelerate AI-driven innovation while addressing operational complexity and compliance needs.



Comparative Analysis: Strategic Trade-Offs

The table below lays out the key points of differentiation for the the three core vendors.

Feature			
AI Factory / Services	Hybrid AI Advantage, AI Fast Start, TruScale, AI Innovators Program	Dell AI Factory with NVIDIA, Professional Services, APEX, Enterprise Hub	HPE Private Cloud AI, Data Services, GreenLake
Reference Architecture	Validated Design for ThinkSystem Servers	NVIDIA Enterprise Reference Architectures	NVIDIA Enterprise Reference Architectures
Time to Value	90-day ROI via AI Fast Start, pre-validated solutions	Fast deployment, cost efficiency (LLM inferencing)	24-hour deployment promise
Flexibility	TruScale GPUaaS, modular portfolio, hybrid support	APEX as-a-service, broad hardware, open ecosystem	GreenLake pay-per-use, modular design, ecosystem integration

The comparison of Lenovo, Dell, and HPE’s AI solutions reveals distinct strengths tailored to diverse enterprise needs. Lenovo stands out with its Hybrid AI Advantage and AI Fast Start, delivering a 90-day ROI through pre-validated solutions, alongside the flexible TruScale GPU-as-a-Service model that supports modular and hybrid deployments. Dell differentiates itself with the Dell AI Factory, leveraging NVIDIA Enterprise Reference Architectures, a broad hardware portfolio, and the APEX as-a-service model, emphasizing ecosystem depth and cost-efficient LLM inferencing for rapid deployment. Meanwhile, HPE’s Private Cloud AI, paired with Data Services and GreenLake’s pay-per-use flexibility, shines with its 24-hour deployment promise and turnkey simplicity, bolstered by seamless NVIDIA integration and a modular design. Ultimately, Lenovo excels in rapid ROI and consumption agility, Dell in ecosystem depth and hardware variety, and HPE in deployment speed and streamlined, ready-to-use solutions, offering enterprises a spectrum of choices to align with their AI strategies and operational priorities.

Recommendations: Strategic Considerations for Enterprise Architects and OT Decision-Makers

As enterprise architects and IT decision makers weigh the merits of on-premises AI deployment alternatives from Lenovo, Dell, and HPE, the decision transcends mere feature comparisons, it's a strategic alignment exercise rooted in organizational realities. To navigate this landscape effectively, four critical evaluation lenses emerge: the inertia of current vendors, the hybrid vs. cloud-only deployment paradigm, security framework alignment, and performance demands. These considerations aren't just checkboxes; they're the guardrails for ensuring AI investments deliver sustainable value without destabilizing existing operations.

Current Vendor Ecosystem: The Cost of Change vs. the Risk of Stagnation

Switching vendors or overhauling tools and processes is a perennial challenge, a reality enterprise architects know all too well. If your organization is entrenched with Lenovo ThinkSystem servers, Dell PowerEdge infrastructure, or HPE ProLiant and Cray systems, the path of least resistance may lean toward sticking with the incumbent. Change introduces friction: retraining staff, rewriting workflows, and potentially disrupting environments where uptime is non-negotiable. Yet, clinging to the status quo carries its own risk, root rot. If your current vendor's AI offerings lag in innovation (e.g., lacking agentic AI support or flexible consumption models), you're betting on yesterday's technology to solve tomorrow's problems. Evaluate whether Lenovo's TruScale, Dell's APEX, or HPE's GreenLake can integrate with your existing stack without a full teardown. Ask: Does the vendor's roadmap align with your AI ambitions, or are you papering over gaps with outdated tools?

Hybrid vs. Cloud-Only: Where Will Your AI Live?

The deployment model, hybrid or cloud-only, isn't a philosophical debate; it's a pragmatic choice driven by workload distribution. AI environments often demand on-premises control for latency-sensitive tasks (e.g., real-time manufacturing analytics), while strategic AI initiatives may tap cloud scalability for training or burst compute. Lenovo's Hybrid AI Advantage shines for organizations straddling both worlds, offering edge-to-core flexibility. Dell's APEX and HPE's GreenLake similarly support hybrid economics, but their strengths skew toward data center-centric deployments. Assess your portfolio: Where will the majority of your AI workloads reside in the next 3-5 years? Cloud-only leanings favor none of these explicitly, look elsewhere if that's your north star.

Security: Aligning AI with Your Framework

Security isn't an afterthought, it's the bedrock of trust in AI deployments, especially in regulated industries like healthcare or finance. Each vendor's solution must slot into your existing security framework, not force a rewrite. Lenovo's Validated Design emphasizes data sovereignty, Dell's Enterprise Hub offers multicloud monitoring, and HPE's Data Fabric ensures unified access with compliance in mind. Scrutinize how these integrate with your identity management, encryption standards, and audit trails. Does TruScale's GPUaaS expose new attack surfaces? Can APEX enforce your zero-trust policies? Map their capabilities against your security posture, misalignment here is a dealbreaker.

Performance: Matching Muscle to Mission

AI performance isn't a monolith, it's workload-specific. Agentic AI or LLM inferencing demands GPU density (e.g., NVIDIA H100s or Blackwell), while edge inference may prioritize compact efficiency. Lenovo's ThinkSystem SC777 V4, Dell's XE9680L, and HPE's Cray supercomputers each cater to distinct needs. Benchmark your requirements: What throughput, latency, and scalability do your use cases demand? Lenovo's 8x content creation boost or HPE's 24-hour deployment won't matter if they choke on your data volumes. Test vendor claims against your AI KPIs, performance mismatches erode ROI fast.



Conclusion: Navigating a Dynamic AI Frontier in a Fast-Moving Market

The on-premises AI deployment landscape, as illuminated through this comparative analysis of Lenovo, Dell, and HPE, reveals a market in the throes of rapid evolution, a crucible of innovation where enterprise needs and vendor ambition collide. Lenovo's Hybrid AI Advantage and TruScale offerings showcase a relentless focus on rapid deployment and consumption flexibility, delivering agentic AI capabilities and a modular edge-to-core portfolio that compresses time to value while empowering scalability. Dell's AI Factory with NVIDIA counters with ecosystem depth and hardware breadth, leveraging APEX as-a-service and a data-centric approach to fast-track ROI and adaptability across diverse workloads. HPE, with its Private Cloud AI and GreenLake model, stakes its claim on unparalleled deployment speed, 24 hours from zero to production, paired with a turnkey simplicity that resonates with enterprises craving immediate impact. Each vendor brings a distinct flavor to the table, their strengths mapped against the twin axes of time to value and flexibility of consumption, offering enterprise architects and decision-makers a rich palette of options for aligning infrastructure with strategic imperatives.

Yet, this report is but a snapshot in a fast-moving market, a dynamic arena where choice and competition are catalysts for relentless innovation. Lenovo's push into agentic AI, Dell's integration of cutting-edge GPUs like Blackwell, and HPE's fusion of Cray supercomputing with private cloud economics aren't isolated advancements; they're responses to a competitive pressure cooker that's driving vendors to outpace one another. This is unequivocally good news for clients. A rising tide raises all boats, each leap forward, whether in deployment velocity, cost efficiency, or workload optimization, elevates the baseline of what enterprises can expect from on-premises AI solutions. Competition isn't just a race for market share; it's a forcing function that ensures vendors like Lenovo, Dell, and HPE continually refine their offerings, delivering tools that are faster, smarter, and more attuned to the real-world complexities of enterprise operations.

But here's the kicker: while on-premises architectures are evolving at breakneck speed, so too are public cloud deployments. The lines between these paradigms are blurring, hybrid models are no longer a niche but a mainstream reality, and cloud providers are doubling down on AI-optimized services that rival the control and economics of on-premises setups. This parallel evolution demands vigilance from ITDMs and architects. The optimal choice today, be it Lenovo's 90-day ROI promise, Dell's ecosystem flexibility, or HPE's turnkey velocity, may shift tomorrow as cloud innovations like serverless AI or enhanced edge compute reshape the calculus. Monitoring this fast-moving market isn't optional; it's a strategic imperative. Enterprises must continually reassess vendor roadmaps, benchmark performance against emerging standards, and weigh hybrid economics against cloud-only advancements.

In this fluid landscape, Lenovo, Dell, and HPE aren't just vendors, they're partners in a journey where adaptability is king. The on-premises AI frontier is vibrant and contested, but it's the relentless pace of change, and the competition it fuels, that ensures enterprises will reap the rewards of innovation. Stay sharp, stay curious, and keep your eyes on the horizon.



ABOUT HYPERFRAME RESEARCH:

HyperFRAME Research delivers indepth research and insights across the global technology landscape, spanning everything from hyperscale public cloud to the mainframe and everything in between. We offer strategic advisory services, custom research reports, tailored consulting engagements, digital events, go to market planning, message testing, and lead generation programs.

Our industry analysts specialize in rigorous qualitative and quantitative assessments of technology solutions, business challenges, market forces, and end user demands across industry sectors. HyperFRAME Research collaborates closely with your Analyst Relations, Product, and Marketing teams to build and amplify your thought leadership, positioning your expertise to enhance brand and product recognition. Through content that engages readers, viewers, and listeners alike, we ensure your voice resonates across channels.

CONTACT HYPERFRAME RESEARCH:

Steven Dickens

CEO & Principal Analyst | HyperFRAME Research

Email Address:

steven.dickens@hyperframeresearch.com

Telephone Number:

+1 845 505 1678

X: - [@StevenDickens3](#)

LinkedIn: [Steven Dickens](#)

BlueSky: [Steven Dickens](#)

CONTRIBUTORS

Steven Dickens CEO & Principal Analyst
HyperFRAME Research

INQUIRIES

Contact us if you would like to discuss this report and HyperFRAME Research will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "HyperFRAME Research." Non-press and non-analysts must receive prior written permission by HyperFRAME Research for any citations.

LICENSING

This document, including any supporting materials, is owned by HyperFRAME Research. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of HyperFRAME Research.

DISCLOSURES

HyperFRAME Research provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

