



Technical Brief

Which Private AI Platform Is Built for Day 2?

Prowess Consulting compares Dell™ AI Factory and HPE® Private Cloud AI based on lifecycle coverage, toolchain continuity, and management flexibility to determine how effectively the two competing solutions meet businesses' needs.

Enterprise AI platforms rarely fail at deployment. They fail later, when growth forces teams to integrate new workloads and extend existing governance and operational tools into the AI environment.

In this report, Prowess Consulting compares Dell AI Factory and HPE Private Cloud AI to assess which platform is better suited to production AI based on lifecycle coverage, toolchain continuity, and management flexibility.

This isn't a benchmark comparison; it's an assessment of how platform design affects long-term operations, governance, and control. Our evaluation centers on three questions that enterprise IT teams consistently raise when moving AI into production:

- Does the platform cover the full AI operating path?
- Can teams keep the tools and workflows they already use?
- Does the management model scale cleanly while meeting governance, security, sovereignty, and integration-effort requirements?

Our comparison shows a clear divide: Dell AI Factory extends the enterprise operating model while HPE Private Cloud AI places more of the operating model inside a vendor-defined management environment. Dell AI Factory offers broader lifecycle coverage and an integration-first approach that allows teams to keep how they operate. HPE Private Cloud AI offers a streamlined path for defined deployments, with a model that centers more strongly on its control plane and validated configurations.

That difference becomes increasingly visible over time: after Day 1, the platform either adapts to the organization or begins to shape it. For enterprises moving AI into production, the ability to adapt becomes the deciding factor.

To illustrate what this all means, consider the following hypothetical. An enterprise that has deployed its first private AI workload and that now needs to connect that workload to existing security, monitoring, service-management, and governance processes. At that stage, the deciding factor is whether the platform can scale without forcing teams to create a parallel operating model.

Key Insights



Dell AI Factory is built for Day 2, when production readiness is truly decided.



Dell AI Factory extends the enterprise, whereas HPE Private Cloud AI asks the enterprise to adapt to it.



Dell AI Factory covers a wider scope of AI operations, which reduces future integration and governance work.

What Matters After Deployment: Evaluating Private Cloud AI for Day 2

Enterprise AI programs face their hardest test starting on Day 2, when the platform has to support governance, integration, and day-to-day operations at scale. Data often becomes the limiting factor as environments expand beyond Day 1, forcing teams to address data access, governance, quality, and movement across a growing set of systems.

A platform can look effortless in a proof of concept but still create friction when it has to fit security and governance requirements across day-to-day operations at scale. Gaps discovered after production deployment can be harder and more expensive to correct because identity, data, security, and operations processes have already formed around the original architecture.

Why Focus on Operations After Deployment?

Enterprise IT teams face a consistent pattern after initial AI deployment. Many organizations invest in AI infrastructure to accelerate early use cases, and they then hit a second phase where the environment grows more complex as data volumes increase and governance requirements tighten. Operations teams then need to connect AI systems to existing monitoring, ticketing, and automation workflows. These pressures force alignment between the business requirements and operations platform.

We examined platform fit under those conditions because that is where long-term value is determined, particularly as AI solutions proliferate. Deployment speed matters, but it doesn't predict whether an AI platform will support production operations after Day 1.

Why Compare Dell AI Factory and HPE Private Cloud AI?

To explore how different platforms perform under those conditions, this report compares Dell AI Factory and HPE Private Cloud AI. Both Dell Technologies and HPE are major enterprise infrastructure providers with broad AI portfolios, and both offer pre-integrated private cloud AI solutions designed to simplify deployment. These solutions are both positioned as end-to-end platforms rather than point products.

That overlap makes the comparison meaningful. It allows our analysis to focus on how each platform behaves as requirements expand, not just how quickly it can stand up an initial workload. In other words, the analysis reflects a shift from Day 1 to Day 2. Day 1 describes deployment, with teams measuring how quickly they can configure hardware, deploy software, and run an initial workload. Day 2 is about operations, with teams discovering whether the platform still works when data grows, governance requirements evolve, and AI workloads become part of the broader IT environment.

Dell reports that Dell AI Factory with NVIDIA has reached 5,000 customers, including 1,000 added during the first quarter of 2026. Although adoption alone does not establish architectural superiority, it increases the practical significance of how the platform handles governance, integration, and Day 2 operations across a growing installed base.

Validated deployment paths have become table stakes. HPE emphasizes a simplified setup experience for HPE Private Cloud AI, including guided deployment for defined configurations. That approach addresses a real concern around speed and ease of implementation, but prepackaged deployment doesn't remove the harder work of identity integration, governance, data pipelines, enterprise security, and Day 2 operations.

As environments expand into Day 2 and beyond, gaps in the initial configuration become harder to ignore. New workloads bring different data requirements. Security teams extend existing controls into AI infrastructure. Operations teams integrate AI systems into established processes. The platform has to absorb that complexity without forcing teams to rebuild how they work.

That distinction between simplicity of setup and robustness of production platform defines the rest of this report. The sections that follow examine how Dell AI Factory and HPE Private Cloud AI perform after Day 1. That's when lifecycle coverage, toolchain continuity, and management flexibility determine whether an AI platform supports production operations or becomes a burden on them.

WHY THIS MATTERS:

Enterprises need to know whether a private AI platform will keep supporting them as their workloads, governance requirements, and operating processes expand beyond the initial deployment.

Does the Platform Cover the Full Scope of Enterprise AI?

Enterprise AI extends beyond a single workload or infrastructure layer. It spans client systems, edge environments, data pipelines, storage, networking, data-center infrastructure, and the technologies that connect and manage them. Gaps in that coverage often remain hidden during deployment and emerge only later, when teams try to extend an initial AI workload into enterprise-wide production use.

Coverage from Client to Data Center

Dell AI Factory presents more complete coverage across the full operating scope, while HPE Private Cloud AI centers its solution in the data center. That broader scope includes Dell Deskside Agentic AI for client-side use cases and Dell PowerRack for data-center AI infrastructure. HPE's data center focus supports validated deployment for defined AI workloads, but it leaves more of the end-to-end environment outside the core platform. For example, HPE does not offer AI PCs as part of its portfolio, which limits its coverage across client-side AI use cases.

This difference shapes how organizations extend AI beyond initial deployments. A platform that covers more of the operating path reduces the number of integration points required as new use cases emerge.

The Data Layer Defines Operational Behavior

The data layer determines how AI workloads scale, move, and interact with the rest of the environment, and it shapes the effort required to operationalize models in production. As AI deployments mature, questions around data access, governance, metadata management, and movement often become more operationally significant than the underlying infrastructure itself. The Dell AI Data Platform combines Dell storage engines, including Dell PowerScale™, Dell ObjectScale™, and PowerFlex™, with open data services such as Starburst®, Elastic®, and Dataloop®. This modular approach lets organizations align storage, search, data access, and pipeline services with different AI workloads without treating the data layer as a single fixed stack.

HPE's storage approach is evolving. The company has expanded its data portfolio with offerings such as HPE Alletra® Storage MP X10000 and HPE Data Fabric Software. HPE positions these offerings as a unified AI data architecture that combines AI-optimized storage with a federated data layer to manage and access distributed data across environments. Earlier HPE Private Cloud AI deployments have relied more heavily on third-party storage integrations, while newer offerings aim to bring more of that capability into the HPE portfolio. That shift will likely strengthen the portfolio over time, but it can also raise questions about maturity, consistency, and migration across deployments, particularly for customers navigating a transition between older and newer approaches.

For customers, that transition can require additional migration planning as teams reconcile earlier and newer storage architectures, operational tooling, and support models. Data placement, migration paths, and operational tooling can change as the storage layer evolves. Existing customers should therefore ask how HPE will support earlier deployments based on VAST® solutions alongside its newer storage architecture, what migration paths will be available, and whether management and support practices will remain consistent across both approaches. In practice, this means more moving parts to align and potentially more effort to expand AI without disrupting existing operations.

Portfolio Consistency at High Scale

Differences in portfolio design become more visible at the high end of the performance range. HPE Private Cloud AI does not support HPE's highest-performance systems using NVIDIA® SXM™ modules. Organizations that require configurations using GPUs such as NVIDIA H200, B200, or B300 GPUs must move into the HPE Cray® portfolio, which uses different hardware and a separate management environment from HPE Private Cloud AI. That transition divides high-end training infrastructure from the platform used for other private AI workloads, creating another operational boundary as requirements grow.

Dell AI Factory incorporates high-performance systems within a more unified management approach. Its modular architecture allows organizations to scale GPU infrastructure without shifting to a separate management framework.

Completeness Shows Up Over Time

Platform completeness is not a feature count. It shows up when organizations scale beyond the first workload.

As AI deployments expand, gaps in the operating path translate into integration work, process changes, or architectural tradeoffs. A more complete platform reduces those disruptions, allowing teams to extend their environment using the same tools and patterns they established at the start.

WHY THIS MATTERS:

A platform with broader operating-path coverage can help enterprises scale AI without turning each new workload into another integration project. As AI environments grow, the operational burden often shifts from infrastructure deployment to data management, governance, and integration across the broader enterprise.

Can Teams Keep the Tools and Operating Models They Already Use?

Most enterprises do not create an entirely new operating model for AI. They extend existing monitoring, ticketing, security operations, automation, and virtualization practices into the new environment. A platform that fits those practices preserves continuity; one that requires teams to rebuild them adds cost and operational friction.

How HPE Private Cloud AI Redefines the Operating Model

HPE Private Cloud AI centers its operating model on a unified control plane. This approach simplifies deployment and provides a clear management structure. It also defines how tools and workflows connect to the platform after Day 1.

The result is an operating model organized around HPE GreenLake®. Existing tools can connect to that model, but the platform establishes a new management center rather than simply extending the one the enterprise already uses. **The platform requires teams to use the GreenLake console for provisioning, orchestration, and lifecycle management**. Existing IT service management (ITSM), observability, and automation practices must fit within or alongside the GreenLake management model. Industry analysts describe **HPE's strategy as bringing compute, storage, networking, and hybrid-cloud operations together through GreenLake**, with operational capabilities integrated within a unified control plane. That approach can strengthen integration across the HPE ecosystem, but it can also reduce flexibility outside it. Initial simplicity can therefore come with a longer-term tradeoff: the enterprise adapts more of its operations to the platform.

How Dell AI Factory Integrates into Existing Operations

Dell AI Factory starts from a different assumption. Most organizations already have established tools and workflows, and the Dell platform is designed to integrate into that environment rather than redefine it.

Integration Versus Replacement

Dell supports integration across security, operations, and management systems. Supported integration points for Dell AI Factory include ServiceNow® for IT service management; Splunk®, IBM® QRadar®, and Elastic for security and observability; and Palo Alto® Cortex XSOAR® for security orchestration. **Dell™ OpenManage™ Enterprise (OME) and SmartFabric Manager provide a unified operational view across server, network, and GPU infrastructure, while Dell AI Ops and the Dell APEX™ Console extend monitoring and management into broader enterprise workflows**.

These integrations allow teams to extend existing processes into the AI environment. Monitoring, incident response, and automation workflows can remain consistent across traditional and AI infrastructure, resulting in greater continuity and saving teams from having to rebuild how they work to accommodate the platform. Extending integration further, Dell Automation Platform can expose partner platforms as modular blueprints so that customers change stacks without redoing infrastructure. This continuity can reduce retraining, limit process changes across security and service management, and lower the operational risk of scaling AI into production.

Choice in Virtualization and Orchestration

Openness also appears in how the platform handles virtualization and orchestration. While both platforms support Kubernetes® and kernel-based virtual machines (KVMs), **HPE Private Cloud AI is Kubernetes-native first and VMware® is not part of the packaged GreenLake solution**; customers with significant VMware investments must run VMware Cloud Foundation® 9.0 (VCF) as a separate stack. Dell AI Factory, by contrast, has deep integration with VCF. That range of integration allows organizations to align AI workloads with their existing infrastructure strategies.

Openness Is an Operational Decision

Openness is often framed as an architectural property. In practice, it shows up in day-to-day operations. A platform is open when teams can keep its tools and evolve its stack over time. After Day 1, that flexibility can narrow as the environment grows. As more systems connect to the platform, the cost of changing direction increases. The operating model either expands with the organization or constrains it.

WHY THIS MATTERS:

Teams can move AI into production faster when the platform extends the tools and workflows they already trust, instead of making them rebuild operations around a new control plane.

Does the Management Model Fit Governance and Security Requirements?

Security in enterprise AI starts with control. Teams need to know where management resides, how the platform connects to the outside world, and whether the operating model fits the rules that govern the environment. Those questions matter most in regulated, sovereign, and air-gapped settings, where management architecture can become a compliance issue instead of a design preference.

The Control Plane Affects Governance

HPE Private Cloud AI uses the HPE GreenLake control plane as its standard management model. In that model, core administration is tied to GreenLake console services, which changes how organizations evaluate connectivity, administrative control, and governance fit. HPE does offer a standalone or air-gapped GreenLake option for restricted environments. However, that option preserves the GreenLake operating model inside the isolated environment and can leave organizations with systems, workflows, policies, and skills that must be added or duplicated.

Dell AI Factory gives customers on-premises management options that can keep more administrative control inside the customer environment; it integrates directly into the existing management infrastructure and extends the customers' established management models. Such an approach can simplify governance for organizations that need local control.

For organizations in regulated, sovereign, or air-gapped environments, the management model affects more than architecture. It can increase compliance review, internal governance effort, and the operational burden of maintaining an isolated environment.

Security Teams Need Continuity, Not a Parallel Workflow

Security operations work best when AI infrastructure fits into the processes teams already use. Most enterprises already run established workflows for incident response, log analysis, change control, and service management. A platform adds operational friction when it forces those teams into a separate model.

Dell AI Factory supports integration with existing security information and event management (SIEM); security orchestration, automation, and response (SOAR); and information technology service management (ITSM) workflows through tools such as OME, AIOps, and related integrations. **That continuity allows security and operations teams to extend current practices into AI infrastructure instead of building a separate set of procedures around it.**

The benefit is practical. Teams can apply familiar controls and route events into existing monitoring paths. Infrastructure changes can follow the same operational mechanisms used elsewhere in the environment. That continuity reduces training overhead and shortens the distance between AI operations and core IT operations.

WHY THIS MATTERS:

Management architecture can determine whether AI infrastructure fits existing governance rules or creates new control-plane risk for security teams to manage.

What Determines Success After the First Workload

Private cloud AI platforms prove their value on Day 2, when the operating model has to scale with the enterprise. In this report, the evidence shows a clear pattern in which platforms either extend the existing enterprise operating model or define the operating model inside a vendor-defined management environment.

Platforms that cover more of the operating path, such as Dell AI Factory, fit into existing tools and create fewer disruptions after Day 1. Conversely, we found that platforms like HPE Private Cloud AI that concentrate value in validated configurations or a centralized control model can introduce friction as new workloads and governance demands enter the environment.

Dell AI Factory offers broader lifecycle coverage, works more naturally with existing enterprise tools and processes, and keeps more management control in the customer environment than HPE Private Cloud AI. HPE Private Cloud AI delivers a streamlined path for defined deployments, but its design places more weight on a central control plane and a narrower operating scope. For enterprises moving AI into sustained production, the deciding factor is whether the platform adapts to the environment over time.

[See the latest research from Prowess Consulting.](#)

Dell AI Factory and HPE Private Cloud AI at a Glance

Differentiator	Dell AI Factory	HPE Private Cloud AI
Enterprise AI scope	Covers a broader operating scope spanning client systems, data, infrastructure, and data-center environments.	Focuses primarily on private AI infrastructure and operations within the data center.
Client-side AI	Extends to client-side AI use cases through Dell Deskside Agentic AI and Dell AI PCs.	Does not offer AI PCs as part of its portfolio.
Data architecture	Dell AI Data Platform combines multiple Dell storage engines with open data services, allowing organizations to select components by workload.	Combines HPE storage and data-fabric technologies in an evolving unified AI data architecture.
Storage portfolio continuity	Uses established Dell storage platforms, including PowerScale, ObjectScale, and PowerFlex, within a modular architecture.	Is transitioning from earlier third-party storage integrations toward newer HPE storage offerings, which can introduce migration and consistency considerations.
High-performance AI infrastructure	Supports scaling to high-performance GPU systems, including Dell PowerRack, within a more unified platform and management approach.	Customers requiring the highest-performance NVIDIA SXM systems must move from HPE Private Cloud AI to the separately managed HPE Cray portfolio.
Existing tools and workflows	Integrates AI infrastructure into existing enterprise monitoring, service-management, security, and automation processes.	Centers operations on the HPE GreenLake management model, requiring existing tools and workflows to integrate with that control plane.

Differentiator	Dell AI Factory	HPE Private Cloud AI
Management model	Provides integration-based and on-premises management options that extend established enterprise operating models.	Uses HPE GreenLake as the primary provisioning, orchestration, and lifecycle-management environment.
Third-party integrations	Supports integrations with tools including ServiceNow, Splunk, IBM QRadar, Elastic, and Palo Alto Cortex XSOAR.	Supports external integrations through the GreenLake-centered operating model.
Virtualization and orchestration	Supports Kubernetes and KVM and offers deep integration with VMware Cloud Foundation.	Supports Kubernetes and KVM; VMware Cloud Foundation operates as a separate stack outside the packaged GreenLake solution.
Governance and local control	Can keep more management and administrative control within the customer environment.	Standard management is tied to GreenLake services; standalone and air-gapped options retain the GreenLake operating model.
Security-operations continuity	Extends existing SIEM, SOAR, ITSM, observability, and incident-response workflows into the AI environment.	Might require security and operations teams to align existing processes with a separate GreenLake-centered workflow.
Day 2 operating approach	Designed to adapt to existing enterprise tools, governance requirements, and infrastructure as AI deployments expand.	Provides streamlined deployment for validated configurations but can require the enterprise operating model to adapt as requirements expand.

FAQ

Q: Who is this paper for?

A: This paper is written for enterprise IT decision-makers (ITDMs) who are evaluating private cloud AI platforms for production use. That audience includes infrastructure architects, platform engineering leads, and operations teams responsible for turning AI projects into stable, governed, and scalable environments. It is also relevant to business stakeholders who need to understand the operational risks behind AI platform choices. The analysis focuses on how platform design decisions affect long-term fit, especially as AI workloads move beyond initial deployment into everyday enterprise operations.

Q: Why do enterprise AI projects often fail after deployment?

A: They fail because operational complexity increases once systems scale. Data pipelines expand, governance requirements tighten, and integration with existing tools becomes critical. Platforms that looked simple during proof of concept can introduce friction in production. This report emphasizes the shift from Day 1 (deployment) to Day 2 (operations) as the primary failure point.

Q: What differentiates Dell AI Factory from HPE Private Cloud AI?

A: Dell AI Factory focuses on lifecycle coverage, integration, and flexibility, whereas HPE Private Cloud AI emphasizes streamlined deployment with a centralized control model. Dell's platform spans from client to data center and integrates with existing tools, whereas HPE's model often depends on GreenLake as a central operational layer, shaping how environments evolve.

Q: Why is lifecycle coverage important in AI infrastructure?

A: Lifecycle coverage ensures the platform supports the full AI operating path, including client devices, edge, data pipelines, and core infrastructure. Gaps in this path might not appear initially but can create integration work later. This report shows that broader coverage reduces long-term complexity and operational disruption.

Q: How does management architecture affect governance and security?

A: Management architecture determines how systems connect. HPE's control plane defines the governance structure around which existing systems and workflows must adapt. Dell's integrated management option keeps control local and can better align with management systems and workflows already in place. This can be particularly important in regulated, sovereign, or air-gapped environments.

Q: What role does toolchain integration play in long-term success?

A: Toolchain integration allows teams to continue using existing monitoring, security, and IT service processes. Dell AI Factory supports integration with SIEM, SOAR, ITSM, and observability tools, enabling continuity. Platforms that require new workflows increase operational cost and complexity over time.

Q: Why does the control plane model matter for scalability?

A: The control plane defines how infrastructure is managed as it grows. A centralized model like HPE's can simplify setup but create dependency over time, making changes harder. A more flexible, integration-based model allows organizations to scale without locking into a single management framework.



The analysis in this document was done by Prowess Consulting and commissioned by Dell Technologies.
Results have been simulated and are provided for informational purposes only.
Any difference in system hardware or software design or configuration may affect actual performance.
Prowess Consulting and the Prowess logo are trademarks of Prowess Consulting, LLC.
Copyright © 2026 Prowess Consulting, LLC. All rights reserved.
Other trademarks are the property of their respective owners.
0826/260027