signal65

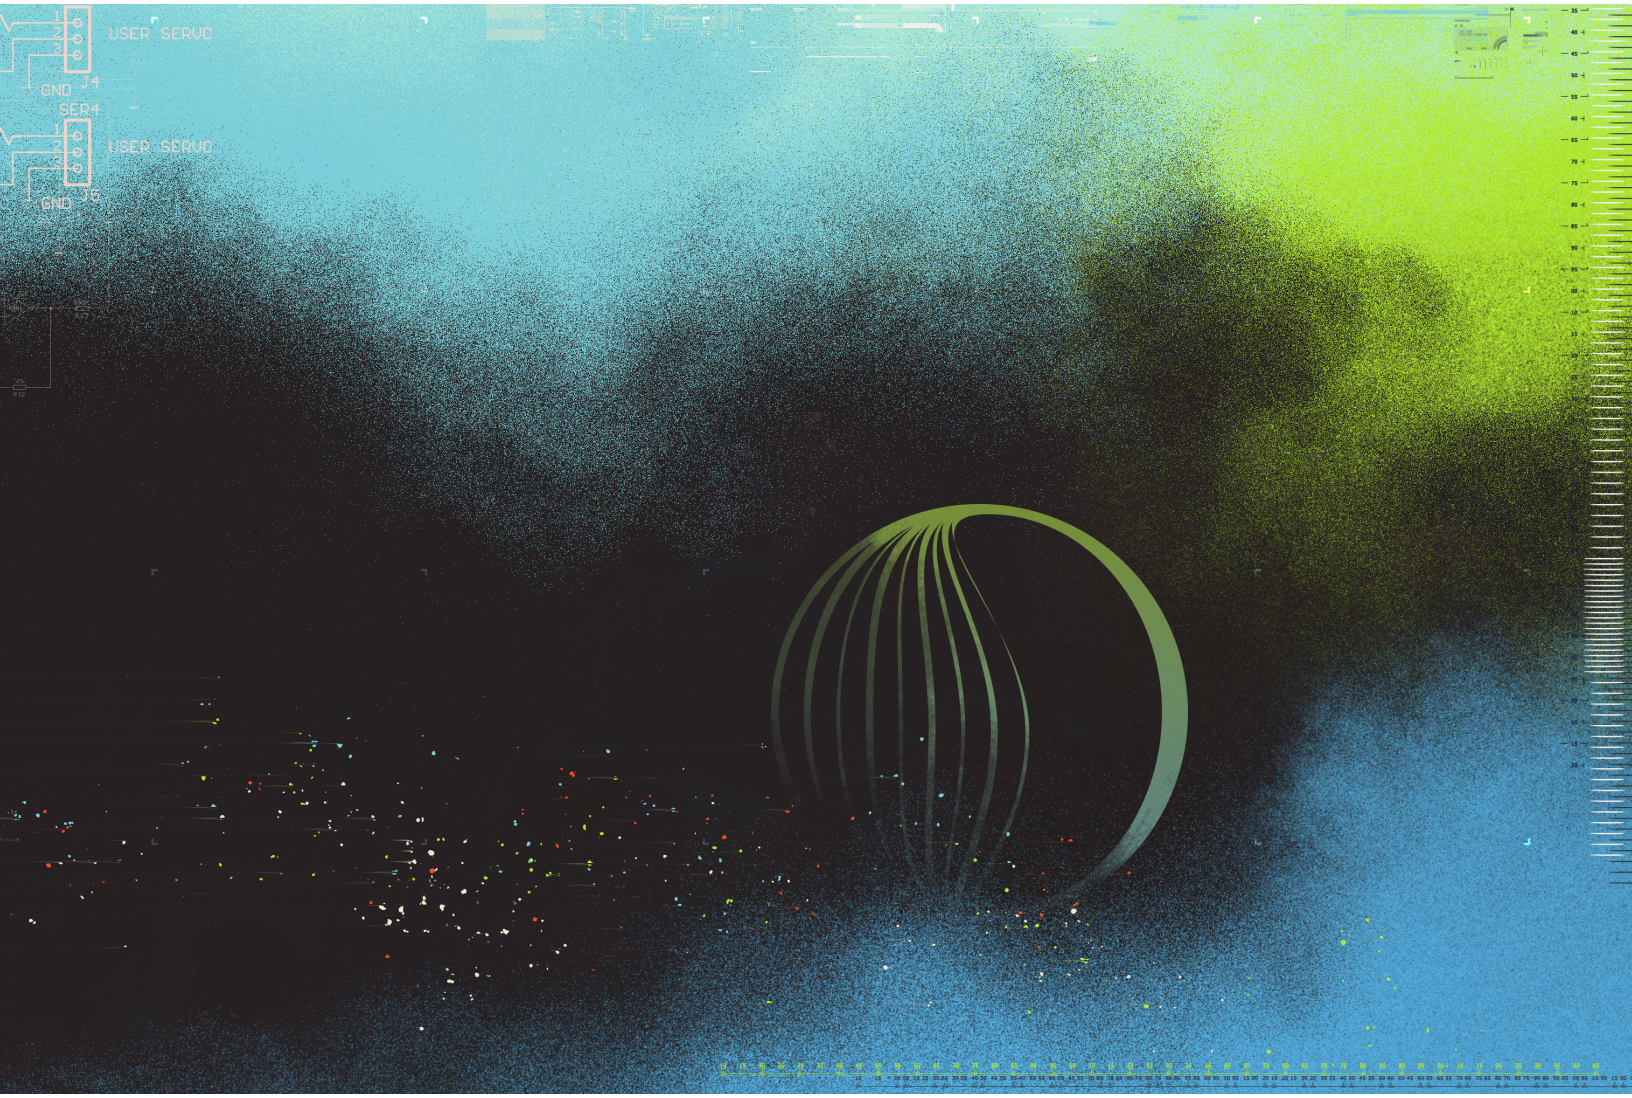

The Unified HPE ProLiant Compute Infrastructure Stack

Analyzing Platform-Level Efficiency, Architectural Security, and Edge Resilience

AUTHOR

Jonathan Fellows
Performance Validation Engineer | Signal65

IN PARTNERSHIP WITH

HPE

AMD

JANUARY 2026

Executive Summary

Organizations are increasingly operating distributed compute environments spanning datacenters, colocation sites, and edge deployments. In these environments where operational oversight can be more difficult, administrative time, security assurance, and consistent lifecycle operations directly impact operational costs and business continuity. At the edge, the management platform, not the hardware alone, determines security posture, lifecycle repeatability, and total operating cost.

Modern enterprise compute requires compute anywhere, including the edge. In this report, we cover the HPE ProLiant DL145 Gen11 server powered by the AMD EPYC™ 8124P processor, evaluated to determine its resilience to thermal stress in terms of sound levels and workload latencies, along with edge deployment options and capabilities. This report also compares the advanced architecture of HPE's platform, specifically the HPE Integrated Lights-Out 7 (iLO 7) and the cloud-native HPE Compute Ops Management, against competitive offerings from Dell, detailing the resulting strategic advantages in security, automation, and predictive intelligence.

Key Findings

HPE ProLiant DL145 Gen11 server performs well under heat stress (105°F ambient temperature) with only small increases in sound levels and workload latency.

HPE iLO 7 delivers significant features and management capabilities compared to Dell iDRAC10. HPE iLO 7 executes most tasks quickly, and outperforms Dell's iDRAC10 on several key tasks.

HPE Compute Ops Management provides cloud-native, enterprise class fleet management capabilities enabling organizations to quickly and intelligently manage their infrastructure.

The HPE ProLiant DL145 Gen11 server with an AMD EPYC 8124P processor demonstrated high resistance to heat stress, delivering consistent latency on AI image inference workloads with minimal performance degradation (less than 2%) even when the ambient temperature increased by 30°F (from 75°F to 105°F). The server operated below 50 dB (measured at 1 meter outside the enclosure) while executing the AI image inference workload at an ambient temperature of 105°F inside a small enclosure. The HPE ProLiant DL145 Gen11's compact size and flexible mounting options enable deployment across diverse environments, including at the edge, while sustaining quiet operation under heavy loads.

A recent Futurum Research Enterprise IT Decision Maker Survey found that remote management is cited by over 60% of enterprises as a top requirement for deploying hybrid and edge workloads, enabling secure, out-of-band management in distributed infrastructures. HPE's Integrated Lights Out 7 and HPE Compute Ops Management extend beyond device-level management by providing a unified operational stack designed for secure, scalable, and repeatable lifecycle management. This enables organizations to standardize workflows, enforce consistent configuration and firmware policies, and reduce the time and risk associated with manual intervention, particularly across distributed environments. The result is a platform engineered for operational efficiency, resilience, and long-term infrastructure value, rather than incremental task-level improvements. While competitive management platforms like Dell's Integrated Dell Remote Access Controller (iDRAC) and Dell AIOps offer control, they can lack the automation, security attestation, and policy consistency required to more easily manage large-scale, multi-site infrastructure efficiently.

Our findings confirm that HPE iLO 7 delivers features that are valuable for IT professionals, increases administrative efficiency, while doing so with an easy-to-use interface. Administrative tasks such as updating firmware, creating users, and launching terminal windows are streamlined to provide benefits for IT admins.

HPE iLO 7 offers several security features including a hardened, zero-trust security foundation by implementing Security Protocol and Data Model (SPDM) device attestation for authentication, a critical feature. HPE custom iLO 7 technology provides in-market support for NIST and CNSA 2.0 quantum-resistant (PQC) algorithms for secure firmware signing. This provides a future-proof defense against "hack now, decrypt later" threats today. The HPE Secure Enclave is architected with physical tamper resistance in mind and is designed to meet the stringent FIPS 140-3 Level 3 requirements (currently pursuing certification). Dell's iDRAC10 is based on the Nuvoton Arbel chip that features a robust solution with its own silicon-based Root-of-Trust as well. However, its cryptographic implementation is built to a different, and less stringent, standard.

HPE Compute Ops Management demonstrates a strategic, cloud-native advantage by offering enterprise features such as the ability to add servers quickly, create and manage servers as a group, Redfish API telemetry, and more. HPE Compute Ops Management also allows organizations to monitor and predict their energy usage, cost and CO_2 emissions by delivering predictive AI sustainability insights. Other AI powered insights include server health and utilization insights to further aid organizations with management and planning.

Server Resilience Under Stress: Environmental and Benchmarking Analysis of HPE ProLiant DL145 Gen11

The Edge Computing Battlefield: Thermal and Acoustic Realities

The HPE ProLiant DL145 Gen11, powered by an AMD EPYC 8124P 16 core processor, is a server purpose-built for the unique challenges of the edge. The single greatest challenge for edge AI is the thermal bottleneck.

Unlike temperature-controlled data centers, edge environments are uncontrolled, non-data center locations such as factory floors, retail stores, or outdoor sites. In these environments, compute-intensive AI workloads generate significant heat. This heat can lead to thermal throttling, an emergency self-preservation mechanism where chips automatically reduce clock speeds to prevent damage. This throttling leads to unreliable performance and can be detrimental for mission-critical edge applications.

The conventional solution, aggressive air-cooling, creates a secondary problem, high acoustic noise. This is unacceptable in human-proximate environments like hospitals, retail spaces, or offices. The HPE ProLiant DL145 Gen11 is engineered to solve both problems with a ruggedized, compact, and low-acoustic design.

The HPE ProLiant DL145 Gen11 demonstrates the thermal resilience, acoustic control, and form factor flexibility required to support AI, industrial, and latency-sensitive applications outside traditional datacenter environments. With the AMD EPYC 8124P processor, the ability to maintain stable inference performance with minimal latency impact, even when operating at elevated ambient temperatures, translates into predictable workload performance in real-world deployment conditions, not just controlled lab environments. This

performance reliability is central to edge AI deployment models, where physical access may be limited, and cooling infrastructure may be inconsistent. The HPE ProLiant DL145 Gen11 supports continuous, stable processing at the edge without requiring specialized environments or high-cost thermal adaptation.

This section analyzes the performance and acoustic data for the HPE ProLiant DL145 Gen11 server, evaluating its architectural resilience and stability when processing demanding AI inference workloads while under thermal stress. The benchmarking utilized two complex AI tasks, Imagenette (image classification) and Coco (image detection) with the yolo11m, yolo11s and yolo11n models from Ultralytics. The primary metrics measured were latency in milliseconds (ms) for inference times (the compute-intensive core task), and frames per second (FPS). Consistency in latency and high FPS are paramount for reliable production inferencing.

The server was placed on a shelf in a small enclosure approximately 2ft x 3ft, with a door that could open and close (see appendix for pictures of the server and enclosure). Sound measurements were taken from a distance of 1 meter from the server both with the door closed and open. An important variable in this testing was ambient temperature: standard 75°F versus extreme 105°F. The AI workloads utilized 100% of the AMD EPYC CPU cores, ensuring the server was running at full capacity and did not have any thermal issues.

Thermal Stability and Latency Consistency

The architectural design of the HPE ProLiant DL145 Gen11 with the 4th Gen AMD EPYC processor demonstrates exceptional thermal management. The minimal change in inference latency observed across a 30°F ambient temperature rise is an impressive finding regarding the HPE ProLiant DL145 Gen11 architectural design and is essential for demanding workloads.

When running the more intense yolo11m model, for the Imagenette Inference task, latency increased by only 1.88% from 21.3ms at 75°F to 21.7ms at 105°F. Similarly, the more complex Coco Inference task saw latency increase from 159.0ms to 161.4ms, a degradation of just 1.51%. This level of stability validates that the HPE ProLiant DL145 Gen11's

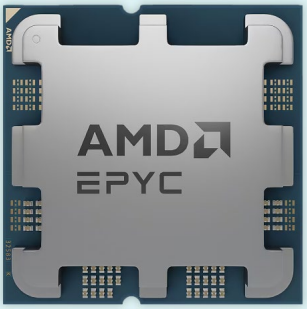

Efficiency engineered for single-socket deployments: The AMD EPYC 8004 Series is purpose-built for optimized single-socket environments, delivering up to 64 cores and 128 threads. This design aligns with the HPE ProLiant DL145 Gen11 to prioritize density and cost-effective scale-out in a smaller thermal envelope.

High performance-per-watt (Zen 4c): Built on 5nm "Zen 4c" architecture, these processors concentrate compute resources to maximize efficiency. This enables the HPE ProLiant DL145 Gen11 to offer strong throughput while minimizing cooling requirements, system acoustics, and long-term energy costs.

Low TDP options for diverse thermal needs: With TDPs as low as ~70W and broad thermal range support, the HPE ProLiant DL145 Gen11 runs reliably in power-constrained environments—such as retail, telecom, and edge sites—without compromising operational stability.

Lower infrastructure and operating costs at scale: The combination of high performance-per-watt and dense compute reduces server footprint and overall TCO. This makes the HPE ProLiant DL145 Gen11 a compelling choice for cloud, storage, and edge deployments where capital and operational efficiency are critical.

Security designed for distributed workloads: AMD Infinity Guard[1] delivers silicon-level security protections suited for environments with limited physical control, helping the HPE ProLiant DL145 Gen11 maintain strong data security across distributed and remote deployments.

thermal management system successfully prevented thermal throttling of the CPU cores. This ensures the server delivers predictable and reliable performance, even when deployed in thermally challenging environments, such as unoptimized data closets or other edge locations.

Metric	Measurement at 75°F	Measurement at 105°F	Delta
Acoustics (Door Closed, 1m)	41.31 dB	48.83 dB	+7.52 dB
Acoustics (Door Open, 1m)	52.81 dB	60.43 dB	+7.62 dB
Imagenette Inference Latency (Avg)	21.3ms	21.7ms	+0.4ms (+1.88%)
Coco Inference Latency (Avg)	159.0ms	161.4ms	+2.4ms (+1.51%)

Table 1: *HPE ProLiant DL145 Gen11's Server Resilience: Thermal Acoustics and AI Workload Latency*

Acoustic Profile and Edge Deployment Suitability

The exceptional thermal stability achieved by the HPE ProLiant DL145 Gen11 under high heat is directly correlated with a highly responsive and necessary Acoustic Profile. To maintain the CPU cores at an optimal operating temperature and prevent the latency spikes associated with throttling, the cooling system efficiently increased its operational speed.

The AMD EPYC 8124P processor has very low TDP down to ~70W, a wide thermal range, and broad operating temperature support lets HPE ProLiant DL145 Gen11 systems operate in constrained or edge locations where power and cooling budgets are limited, expanding deployment flexibility across retail, telecom, and remote sites. AMD has designed this processor to be optimized for 1 socket deployments, ideal for edge locations. Additionally, the AMD EPYC 8142P processor minimizes power usage with cores on a 5nm process, further supporting the edge use case.

The acoustic measurements were taken at a distance of 1 meter against a low ambient background level of approximately 30 dB. When ambient temperature rose from 75°F to 105°F, the sound output increased from 41.31 dB (door closed) to 48.83 dB (door closed), and from 52.81 dB (door open) to 60.43 dB (door open). For reference, ~40 dB is typical for a quiet home, ~60 dB for normal conversation or background music, and sustained exposure above ~85 dB can be harmful to hearing. This fan ramp, resulting in an increase of about 7.5 dB, confirms the server's architectural commitment to computational stability by utilizing aggressive cooling only when required to guarantee predictable performance in thermally demanding environments, such as unoptimized data closets or Edge locations.

Inferencing Workloads at the Edge

For edge environments where modest image detection workloads are common, such as retail or manufacturing, the yolo11n and yolo11s models were tested in addition to yolo11m to showcase the capabilities of the HPE ProLiant DL145 Gen11 with the AMD EPYC 8124P processor. Again, the workloads used the Imagenette (image classification) and Coco (image detection) datasets to simulate what would be used in an edge environment.

Workload	CPU Avg Latency (at 75°F)	Approximate Frames Per Second
yolo11n Imagenette (Classify) Inference	16.3ms	61.3
yolo11n Coco (Detect) Inference	89.3ms	11.2
yolo11s Imagenette (Classify) Inference	22.8ms	43.8
yolo11s Coco (Detect) Inference	116.0ms	8.6
yolo11m Imagenette (Classify) Inference	40.8ms	24.5
yolo11m Coco (Detect) Inference	185.8ms	5.4

Table 2: *Average Latencies and Frames Per Second for yolo11 models*

As seen in the table above, the FPS for the Imagenette dataset were near or above 30 FPS for all model sizes. This implies that the HPE ProLiant DL145 Gen11 can support image classification at the edge for most use cases. With the more complex detect task, the FPS were lower, but still useful for scenarios that don't require real-time detection. The performance on AI workloads can be scaled up, should consumers require, with a dedicated GPU acceleration device to achieve low latency targets. This would improve the latencies and FPS for the detect workloads, for example. The gains in performance provided by a GPU transforms the HPE ProLiant DL145 Gen11 from a general-purpose server into a viable, high-throughput AI inferencing platform.

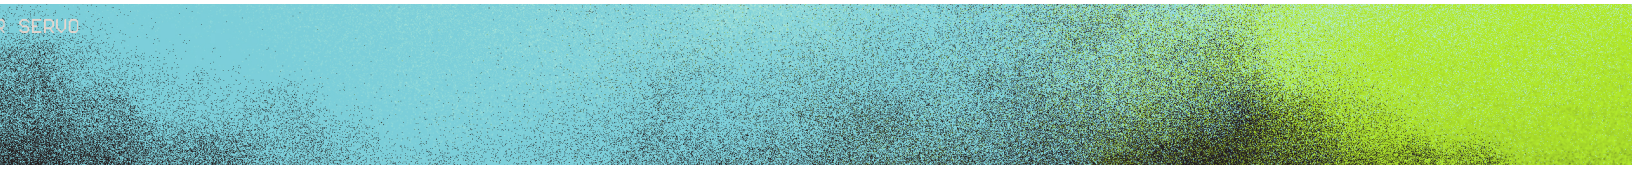

HPE iLO 7 Advanced vs Dell iDRAC 10 Datacenter

For decades, out-of-band (OOB) management was defined by basic, reactive functionalities, often relying on the dated and insecure Intelligent Platform Management Interface (IPMI) protocol. The modern IT landscape, however, demands a proactive, programmable, and deeply integrated management layer, a Baseboard Management Controller (BMC), to support automation, "infrastructure-as-code" (IaC) workflows, and a zero-trust security posture.

Quantified Operational Efficiency: HPE iLO 7 vs Dell iDRAC10

The primary benefit of HPE iLO 7 is converting critical, often manual administrative tasks into rapid, measurable workflows that are fully accessible via both the modern GUI and its Redfish API. This approach drastically reduces the operational friction and risk associated with tasks typically handled manually or via rudimentary interfaces. The tasks were performed 3 times, with the time/clicks measurements recorded on the 3rd execution.

Task	HPE iLO 7 Time (s)	HPE iLO 7 Clicks	Dell iDRAC10 Time (s)	Dell iDRAC10 Clicks	Analysis
Launch Remote Console	5	2	5	1	HPE Equal Time ~= Clicks
Check Security Events, Download File	9	4	5	2	Dell Faster/Fewer Clicks
Create Least-privilege User	19	11	16	8	Dell Faster/Fewer Clicks
Update Firmware from Repository	8	6	10	7	HPE Faster/Slightly Fewer Clicks
Change Boot Order	12	5–15	17	6–15	HPE Faster ~= Clicks
Set up SMTP Alerts	33	10	24	8	Dell Faster/ Slightly Fewer Clicks

Table 3: *iLO 7 Administrative Tasks HPE vs. Dell Comparison*

HPE iLO 7 showed an advantage vs iDRAC10 as it required only 8 seconds and 6 clicks to initiate a firmware update using a repository or local file, making it 2 seconds faster and requiring one fewer click than iDRAC10. With iLO 7, a minimal amount of time is spent on firmware updates during critical security patching. Similarly, the ability to change the boot order in 12 seconds provides 5 seconds of administrative velocity during critical recovery or maintenance tasks with HPE iLO 7, compared to iDRAC10's 17 seconds.

While Dell iDRAC10 demonstrated quicker completion times for setup tasks like creating a least-privilege user (16 seconds vs 19 seconds) and setting up SMTP alerts (24 seconds vs 33 seconds), HPE iLO 7 provides a highly streamlined process for Role-Based Access Control (RBAC) creation and proactive monitoring that is foundational for fleet management, ensuring high value despite the slight time difference. With basic task speed largely at parity, the architectural differences between HPE and Dell in security and fleet management become the primary, strategic differentiators.

HPE iLO 7 proves to be a highly responsive and modern management platform. While both solutions are capable, iLO 7 distinguishes itself by streamlining common maintenance tasks—specifically firmware updates and boot configuration—reducing the time and steps required for these frequent operations. This efficiency, paired with a user-friendly interface, ensures that administrators can execute critical workflows with confidence and minimal friction.

Strategic Impact of HPE iLO 7 in Daily Operations

The value of HPE iLO 7 is not simply in the number of clicks or seconds needed to complete administrative tasks, it is in the repeatability and reliability of those tasks across hundreds or thousands of systems.

When firmware updates, access control provisioning, or security verification workflows can be executed consistently through a hardened, policy-driven interface, organizations gain:

- Reduced administrative overhead
- Faster security patch response
- Fewer configuration drift incidents
- Higher confidence in system integrity

This consistency is especially critical in hybrid operations models where infrastructure teams must support many small sites, not just centralized datacenters. HPE iLO 7 enables those teams to maintain the same security posture and management standards from core to edge without additional tooling layers or integration overhead. HPE also introduced a search function in iLO 7 that allows users to find a task more easily from anywhere in the workflow, similar to Dell iDRAC10.

Architectural Hardening and Zero-Trust Security

HPE iLO 7 provides a management stack that is integral to a comprehensive zero-trust model:

- **Security Protocol and Data Model (SPDM) Device Attestation:** HPE iLO 7 supports DMTF SPDM to authenticate “applicable components” and establish trust at the device level. When enabled, HPE iLO 7 logs verified/unverified results for option cards such as storage controllers and Network Interface Cards (NICs), and where supported, NVMe devices and accelerators, along with individual server chassis components. This ability to audit and verify hardware integrity is a fundamental defense mechanism against hardware tampering.
- **Encryption Key Management:** HPE currently supports ESKM solutions utilizing the NAE-XML protocol established by Thales.
- **Secure Enclave:** The HPE Secure Enclave, embedded in the custom-designed iLO 7 ASIC, is a hardware-based vault that securely stores server encryption keys and provides multi-layered protection against tampering and intrusion. It is the first industry-standard server security solution to meet **FIPS 140-3 Level 3** requirements (currently pursuing certification) and support NIST and CNSA 2.0 quantum resistance standards, enabling secure firmware signing to defend against future quantum computing threats.
- **FIPS 140-3 Level 3 Architecture:** The HPE Secure Enclave is architected with physical tamper resistance designed to meet the stringent FIPS 140-3 Level 3 requirements. This is a crucial distinction. FIPS 140-3 Level 1 is a basic software-level cryptographic validation. Level 3 is a far higher standard, requiring physical tamper resistance and identity-based authentication. This provides a robust defense against physical access, side-channel, and tampering attacks, which software-level validation does not address.
- **Quantum-Resistant Cryptography (PQC):** HPE iLO 7 has an enabled "CNSA" security state that "uses CNSA 2.0 signing". This aligns with NIST and the NSA's CNSA 2.0 standards for post-quantum cryptography, ensuring the integrity of firmware signing. This is an in-market capability today. It provides a viable defense against "hack now, decrypt later" attacks, where an adversary could capture firmware-signing data today and use a future quantum computer to break its encryption, allowing them to create malicious, validly-signed firmware.

Dell's Approach: A Standard, Present-Day Implementation

Dell's iDRAC10 is a robust solution with its own silicon-based Root-of-Trust. However, its cryptographic implementation is built to a different, and less stringent, standard.

1. **FIPS 140-3 Level 1 Validation:** Dell's BSAFE cryptographic modules, which are leveraged by iDRAC, have an active, official validation from NIST for **FIPS 140-3 Level 1**. As noted, this is a software-level validation that does not require the physical tamper-resistance mechanisms mandated by the Level 3 standard.
2. **PQC Roadmap:** Dell's public-facing posture is that PQC is a "strategic imperative" and they are "actively working to implement" it in line with 2025 government guidance. While some specific 17th-generation server components, such as certain HBAs, are beginning to mention CNSA 1.0 capabilities, this is not the same as a systemic, platform-level PQC signing standard for the BMC firmware itself. This contrasts with HPE's in-market CNSA 2.0-aligned signing.

Security Vector	HPE iLO 7 (Gen12 Servers)	Dell iDRAC 10 (Gen17 Servers)	Strategic Implication
Cryptographic Module Validation	Designed to meet FIPS 140-3 Level 3	Validated at FIPS 140-3 Level 1	HPE architecture is designed for physical/tamper resistance, a higher standard required by regulated industries.
Post-Quantum Cryptography (PQC) for Firmware Signing	In-Market Support (CNSA 2.0-aligned signing)	Roadmap Commitment ("Actively working to implement")	HPE provides protection against long-term "hack now, decrypt later" threats.
Device Attestation	DMTF SPDM supported	DMTF SPDM supported	Both platforms use the modern SPDM standard, but HPE PQC-enabled Secure Enclave provides a more secure foundation.

Table 4: *Comparative Security Architecture (HPE iLO 7 vs. Dell iDRAC 10)*

For any organization in regulated industries (government, finance, healthcare) or those with long (10+ year) infrastructure lifecycles, PQC-readiness and a physical (Level 3) security architecture are foundational requirements. The evidence shows that HPE iLO 7 architecture is demonstrably built for tomorrow's threats, not just today's.

Modern Observability and Programmability (Redfish)

In contrast to older management standards, HPE iLO 7 provides a rich, programmatic interface required for modern telemetry and infrastructure-as-code:

- **Full Redfish Support:** HPE iLO 7 implements a hardened management stack with full support for Redfish, the modern RESTful API standard for OOB control, replacing the outdated IPMI protocol. This standardization allows all core workflows, remote console/KVM, user/RBAC management, firmware views, and log export to be straightforward, modern, and fully automatable through scripting and infrastructure tools.
- **Enhanced Telemetry & Observability:** HPE iLO 7 surfaces detailed power, thermal, and sensor data in the Graphical User Interface (GUI) (under Power & Thermal/Sensors) and exposes the same data via Redfish for programmatic collection. This rich data stream enables streaming and broader observability use cases by polling Redfish into a centralized telemetry stack.
- **iLO Advanced Licensing Value:** Upgrading to HPE iLO Advanced unlocks features essential for large teams and regulated environments, including remote console, session recording/playback, virtual media/folders, enhanced power/thermal controls, one-button secure erase, and broad enterprise authentication options (e.g., CAC/PIV, Kerberos).

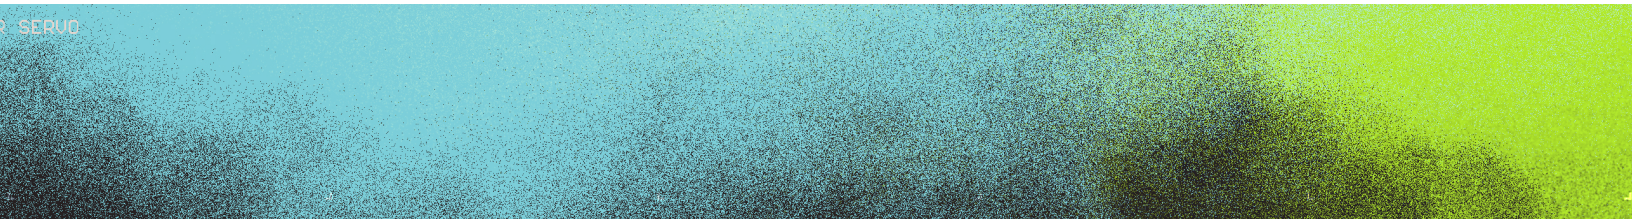

HPE Compute Ops Management vs. Dell AIOps

Moving from individual server management to fleet-level operations requires a cloud-native, policy-driven platform. HPE Compute Ops Management provides the necessary automation and architectural differentiation for scaling management complexity compared to Dell AIOps.

Another Futurum Research finding shows that Cloud Operations Management platforms that integrate with hardware management APIs (e.g., HPE iLO 7) are now viewed as essential for orchestrating multi-site enterprise environments, particularly as edge and remote site deployments scale. Similarly, more Futurum Research shows that among enterprises with edge deployments, 71% report that seamless integration between server hardware management (such as HPE iLO 7) and cloud-native management platforms is a major driver for vendor selection and operational success.

Quantified Fleet Management Workflows with HPE Compute Ops Management

HPE Compute Ops Management moves management beyond individual server tasks by focusing on group operations and rapid, cloud-native control, offering measured efficiency gains for administrators handling large fleets. Compute Ops Management provides a centralized, cloud-based platform for grouping servers, applying lifecycle policies, monitoring sustainability impacts, and generating actionable health insights. HPE Compute Ops Management workflows are designed for fleet-level operations, not just device-by-device administration.

Hands-on testing demonstrates the speed of these cloud-native, fleet-level tasks and highlights a key architectural differentiator for HPE.

Again, the time measurements were recorded on the **3rd task attempt**.

HPE Compute Ops Management now supports advanced status monitoring, including air filter health for servers equipped with the FIO front bezel and filter kit (P72582-B21). For the HPE ProLiant DL145 Gen11, the system automatically alerts administrators when the filter reaches maximum usage in high-particulate environments, transitioning status from OK to Warning or Critical.

Fleet Task	HPE COM Time (s)	HPE COM Clicks	Dell AIOps Time (s)	Dell AIOps Clicks	Analysis
View Sustainability Data	5	4	7	3	HPE Faster Time
Configure alerts/ notifications	8	5–7	13	4–15	HPE Faster Time
Check Security Events on a server	7	3	10	3	HPE Faster Time
Create a group of servers and save it	15	11	N/A (can only tag servers)	N/A	HPE Unique Persistent Grouping
Add individual servers (generate access key)	15	8	N/A (requires collector)	N/A	HPE Easier Onboarding
Create firmware baseline/policy	19	14	18	9	Dell Slightly Fewer Clicks
Check firmware policy, update server	24	10	17	9	Dell Faster/ ~= Clicks

Table 5: *Cloud Administrative Tasks: HPE Compute Ops Management vs. Dell AIOps*

A major differentiator for HPE is Persistent Group Segmentation. HPE Compute Ops Management enables administrators to create a group of servers using complex filters and save it for continuous, policy-driven use (15 seconds, 11 clicks). By contrast Dell AIOps is unable to create a saved group, but can tag servers forcing administrators to rely on manual tagging or recreate complex filters for recurrent tasks. Similarly, HPE's cloud-native architecture allows for direct server addition via an access key (15 seconds, 8 clicks), removing the reliance on an intermediary collector required by Dell AIOps, simplifying deployment across distributed environments. HPE Compute Ops Management also has the ability to integrate servers via traditional collectors should consumers choose that method.

HPE Compute Ops Management demonstrates a speed advantage in monitoring workflows. Configuring alerts is faster in HPE Compute Ops Management (8 seconds vs 13 seconds for Dell), prioritizing rapid operational protection. HPE Compute Ops Management is also faster at providing a top-level security status check in 7 seconds (vs 10 seconds for Dell). Furthermore, accessing the Sustainability Data card is faster in HPE Compute Ops Management (5 seconds vs 7 seconds for Dell).

Dell AIOps demonstrated quicker execution in traditional lifecycle management tasks, such as creating a firmware baseline (18 seconds, 9 clicks vs 19 seconds, 14 clicks for HPE Compute Ops Management) and checking/updating a server (17 seconds, 9 clicks vs 24 seconds, 10 clicks for HPE Compute Ops Management). Additionally, Dell AIOps maintains a maturity lead in traditional reporting, offering the capability to schedule reports. However, this is balanced by HPE Compute Ops Management's unique architectural strengths in policy segmentation and predictive insights, which are more critical for automated, future-focused management at scale.

Beyond the comparable task timings, HPE Compute Ops Management offers a distinct architectural benefit for scaling operations: the ability to create and save persistent server groups. Unlike traditional tagging methods, this feature allows for continuous policy enforcement across dynamic fleets. Combined with a simplified onboarding process that removes the need for intermediary collectors, the platform is particularly well-suited for organizations looking to manage distributed environments with greater agility.

A Strategic Differentiator: Policy-Driven Grouping

Traditional fleet management often relies on manual filtering or basic tagging for group operations. HPE Compute Ops Management provides a necessary architectural advantage for true cloud management and scaling:

- **Persistent Group Segmentation:** HPE Compute Ops Management offers the critical capability to create a group of servers using tags/filters (e.g., location + model + firmware level) and save it for persistent, policy-driven use. This ability to save and re-use complex, multi-criteria groups is essential for scaling policy-driven infrastructure and significantly reduces the risk of human error associated with manual, ad-hoc filtering. Dell AIOps is unable to create a saved group, forcing administrators to manually recreate complex filtering rules for recurrent tasks.
- **Onboarding:** HPE Compute Ops Management allows direct server addition via an access key, removing the requirement for an external collector. Generating this key takes only seconds, and the key can then be used for adding additional servers. Dell AIOps is unable to directly add servers, requiring the installation of an intermediary collector (Open Manage Enterprise) before management can begin, adding setup complexity. This direct, cloud-native approach simplifies the management architecture, lowers the setup barrier, and suggests a design philosophy optimized for rapid, distributed deployment across various environments.

Predictive AI: ESG Forecasting and Strategic Intelligence

The transition from basic monitoring to predictive AIOps requires leveraging AI to generate strategic business insights that drive Total Cost of Ownership (TCO) reduction and environmental governance. While Dell AIOps focuses on historical energy use data, HPE Compute Ops Management offers the ability of predictive energy use, adding substantial value.

HPE Compute Ops Management: Predictive Sustainability Insights

HPE Compute Ops Management distinguishes itself by leveraging AI for prediction and forecasting in the realm of sustainability.

- **Projected TCO and Environmental Governance:** HPE Compute Ops Management provides sustainability data on the main page, along with historical and projected CO2, energy, and cost data for servers with their AI Insights. This future-looking capability is essential for aligning infrastructure operations with corporate Environmental, Social, and Governance (ESG) goals and long-term TCO modeling.
- **Prescriptive AI Insights:** HPE Compute Ops Management favors a prescriptive, dashboard-driven AI model, with "AI Insights" structured into three specific, high-value outputs: Sustainability insights, Server utilization insights (limited to Intel CPU based ProLiant servers for now), and a Server hardware inventory Report. This model emphasizes delivering specific, pre-calculated, actionable intelligence directly to the user.

Cloud-Native Lifecycle Management for Distributed Deployments

HPE Compute Ops Management provides organizations with a centralized platform for grouping servers, applying lifecycle policies, monitoring sustainability impacts, and generating actionable health insights.

As opposed to fleet management approaches that rely on manual device-by-device administration, HPE Compute Ops Management treats infrastructure as a unified, policy-driven system that can be managed programmatically and at scale. This is especially valuable for organizations expanding AI inference, retail/branch compute, and industrial edge environments where the operational model shifts from “manage servers where they are” to “manage servers as a coordinated fleet, regardless of where they are.”

The ability to onboard devices quickly, enforce consistent firmware baselines, and surface sustainability and utilization insights through structured AI assistance positions HPE Compute Ops Management as a forward-aligned operational architecture, not just a management tool.

Summary

The HPE infrastructure stack demonstrates critical architectural and operational advantages compared to Dell's platforms, focusing on security, efficiency, and future-ready AI capabilities. Through the combination of HPE iLO 7, HPE Compute Ops Management, a cohesive and forward-ready operational foundation for modern compute environments is provided. Together, these capabilities enable organizations to reduce operational complexity, enforce a stronger and more measurable security posture, and deploy reliable compute resources across datacenter, colocation, and diverse edge environments highlighted by the HPE ProLiant DL145 Gen111, powered by AMD EPYC processors. The result is a platform engineered for lifecycle efficiency, predictable performance, and long-term infrastructure value, aligned with the growing shift toward distributed and AI-enabled workloads.

Server Resilience (HPE ProLiant DL145 Gen11): The HPE ProLiant DL145 Gen11 server powered by AMD EPYC processors validated its architectural commitment to workload stability, maintaining AI inference latency degradation below 2% even when subjected to an ambient temperature increase of 30°F (from 75°F to 105°F). This thermal resilience confirms its suitability for deployment in uncontrolled edge environments. A GPU compatible with the HPE ProLiant DL145 Gen11 can further boost performance for demanding workloads.

Granular Management (HPE iLO 7 vs. Dell iDRAC10): HPE iLO 7 provides superior velocity in high-impact lifecycle tasks, completing firmware updates and boot order changes faster than Dell iDRAC10 (8 seconds vs. 10 seconds for firmware, 12 seconds vs. 17 seconds for boot order). Architecturally and in alignment with NIST and the NSA's CNSA 2.0 standards, HPE offers a higher security standard with its Secure Enclave designed to meet FIPS 140-3 Level 3 requirements and its in-market readiness for Post-Quantum Cryptography (PQC), positioning it for long-term security against advanced threats.

Centralized Fleet Management (HPE Compute Ops Management vs. Dell AIOps): HPE Compute Ops Management offers unique, scale-critical features that Dell AIOps lacks, notably the ability to create and save persistent server groups for policy enforcement and deliver predictive sustainability data (forecasting future CO2, energy, and cost), essential for ESG planning. HPE Compute Ops Management also demonstrated faster time-to-value in crucial monitoring tasks, including configuring alerts (8 seconds vs. 13 seconds for Dell) and checking security status (7 seconds vs. 10 seconds for Dell). While Dell retains a maturity lead in traditional reporting (ability to schedule reports), HPE's focus is on predictive, cloud-native policy orchestration.

Appendix

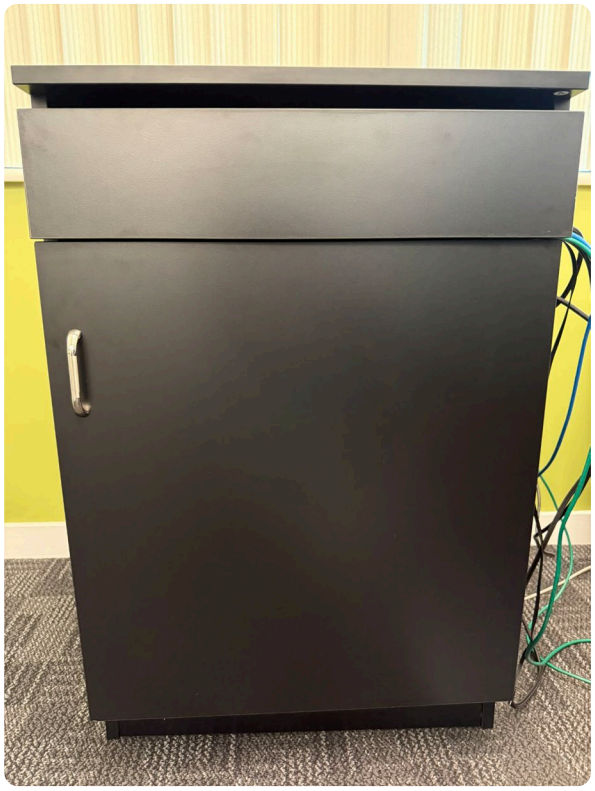

HPE ProLiant DL145 Gen11 server and its enclosure.

All task time/click measurements were started from the home screen in all scenarios for both iLO7 vs iDRAC 10 and AIOps vs COM.

Tasks with a range of clicks included using the minimal and maximum amount of options available for that task. Time was recorded for the minimal amount of options in these cases.

Frames Per Second values for the yolo benchmarks were calculated using the average, and p50/p95 latencies.

[1] AMD Infinity Guard features vary by AMD EPYC Processor generations and/or series. Infinity Guard security features must be enabled by server OEMs and/or Cloud Service Providers to operate. Check with your OEM or provider to confirm support of these features. Learn more about Infinity Guard at http://www.amd.com/en/products/processors/server/epyc/infinity-guard.html.

Important Information About this Report

CONTRIBUTORS

Jonathan Fellows
Performance Validation Engineer | Signal65

Cameron Moccari
Operations and Project Management Director | Signal65

PUBLISHER

Ryan Shrout
President and GM | Signal65

INQUIRIES

Contact us if you would like to discuss this report and Signal65 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 "Signal65." Non-press and non-analysts must receive prior written permission by Signal65 for any citations.

LICENSING

DISCLOSURES

Signal65 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.

ABOUT SIGNAL65

Signal65 is a leading research organization specializing in enterprise AI infrastructure optimization and deployment strategies. Our lab focuses on evaluating and optimizing AI hardware and software solutions for real-world enterprise applications, with particular expertise in large language models, retrieval-augmented generation systems, and distributed AI architectures.

For more information, visit signal65.com
or contact research@signal65.com

IN PARTNERSHIP WITH

signal65

CONTACT INFORMATION
Signal65 I signal65.com