

Lenovo Hybrid AI Advantage™ with NVIDIA:

Driving hospitality success

A scalable, validated AI solution combining Lenovo's ThinkSystem infrastructure, NVIDIA's cutting-edge AI acceleration, Centific's AI Data Foundry, delivering real-time guest personalization and streamlined operations.



Hospitality leaders face high expectations for personalized service, cost efficiency, and sustainable growth. Today's guests demand seamless check-ins, tailored experiences, and smart amenities—while operators must optimize back-of-house processes and achieve strong ROI.

Challenges:

In today's hospitality environment, fragmented systems and manual processes hinder efficiency and reduce guest satisfaction.

Key challenges include:

- **Elevated guest expectations:** Travelers seek hyper-personalized experiences.
- **Operational complexity:** Manual workflows in housekeeping, maintenance, and security slow responses, especially during peak periods.
- **Data silos:** Disconnected data from check-ins, occupancy, and F&B sales limits real-time decision-making.
- **Labor & cost pressures:** Rising costs and inefficiencies demand data-driven automation to reduce waste.
- **Complex system integration:** Integrating diverse systems and technologies seamlessly is critical, and traditional solutions often fall short in delivering a unified, scalable AI environment.

Integrated AI infrastructure

Powered by Lenovo's ThinkSystem portfolio, Accelerated by NVIDIA AI, and Centific AI Data Foundry, this solution delivers a pre-validated, enterprise-grade platform for hospitality. Built on the Lenovo Validated Design (LVD) for Multimodal Agentic AI, this solution empowers hospitality leaders to overcome challenges.

Solutions:

- **Real-time Personalization:** Anticipate guest preferences instantly, enabling multilingual AI check-ins and in-room customizations.
- **Streamlined Operations:** Automate tasks like housekeeping scheduling and early maintenance detection, freeing staff for guest-focused roles.
- **Data-Driven Insights:** Consolidate property-wide metrics into a single dashboard, optimizing staffing, inventory, and marketing for maximum ROI.
- **Scalable Deployment:** Start with a pilot system and expand to enterprise-grade implementations without infrastructure redesign.

30%



Increase in guest satisfaction with AI-driven personalization

Lenovo





Achieving tangible improvements for hospitality

Whether managing a boutique hotel or a large resort chain, modern operations demand real-time insights, faster service delivery, and proactive guest engagement. By leveraging data from sources such as room occupancy, on-site sensors, and guest feedback, Centific AI Data Foundry delivers immediate, measurable outcomes:

- ✓ **Real-time personalization:** An AI concierge transforms every guest into a VIP—delivering authentic, tailored recommendations, suggesting relevant upgrades, and creating a seamless, unforgettable experience.
- ✓ **Proactive issue resolution:** Identify potential disruptions (maintenance, safety concerns) early, significantly reducing downtime and ensuring a consistent guest experience.
- ✓ **Efficiency gains:** Automate labor-intensive tasks like housekeeping and supply tracking, allowing your team to focus on higher-value interactions.
- ✓ **Scalable growth:** Expand seamlessly from small pilots to enterprise-grade deployments, applying consistent data standards and AI frameworks across your portfolio.

Key takeaways

Centific AI Data Foundry transforms hospitality into a proactive, data-driven operation: **Key advantages include:**



Unified architecture

Deploy a complete, end-to-end AI environment quickly by leveraging a pre-tested Lenovo Validated Design with Centific AI Data Foundry.



NVIDIA acceleration

Enjoy world-class GPU performance for low-latency inferencing and real-time personalization, as well as support for complex vision-language workloads such as multilingual concierge services.



Cost-effective scalability

Begin with a small deployment and expand effortlessly to additional properties without a full infrastructure redesign.



Enhanced security and compliance

Benefit from advanced controls that protect guest data while meeting stringent hospitality privacy and safety standards.


By unifying AI-driven personalization with operational automation—and harnessing the strength of Lenovo infrastructure and NVIDIA acceleration—Centific AI Data Foundry empowers hotels and resorts to **boost guest loyalty, optimize efficiency, and drive profitability** in an ever-competitive market.

Lenovo Validated Design Architecture

Built on Lenovo ThinkSystem hardware with NVIDIA GPU acceleration, this Validated Design is fully assembled, performance-tuned and exercised against real hotel data flows—video analytics, multilingual check-ins, RAG queries—before it ever reaches your site. Unlike generic reference architectures, every node, driver and firmware level is locked, documented and backed by Lenovo Support, so you can drop it in, power on and start serving guests within hours instead of weeks.



Lenovo ThinkSystem SR655 V3/SR635 V3
Single-socket servers used for control, orchestration, and general workloads. Perfect for running Kubernetes-based platforms and managing real-time data pipelines without compromising speed or reliability.



Lenovo ThinkSystem SR675 V3
A GPU-rich rack server built for high-intensity AI. Supports up to 8 NVIDIA L40S or H200 NVL GPUs for advanced multimodal workloads—ideal for real-time guest personalization and surveillance.

Lenovo’s AI infrastructure scales seamlessly from Starter deployments to Scalable Unit (SU) deployments, leveraging L40S or H200 NVL GPUs for high-performance workloads. This LVD scalable model enables cost-effective AI model refinement before full-scale deployment. The following outlines scalability across different tiers.

	ThinkSystem platform	GPU	AI compute nodes	Concurrent video requests
Starter	SR675 V3	L40S (8 count)	Single	10
Starter/Intermediate	SR675 V3/SR635 V3	H200 NVL (8-16 count)	Dual	20-40
Enterprise (one scalable unit deployed)	SR675 V3/SR655 V3	H200 NVL (16-32 count)	4 nodes	40-80
Multiple scalable unit deployment	SR675 V3/SR655 V3	H200 NVL (32+ count)	Multiple SU racks	80+

Why Lenovo

Lenovo’s portfolio is engineered for compute-intense AI in challenging environments—whether in a resort, at the urban edge, or across large-scale property clusters. With Lenovo Validated Designs, customers benefit from:

- Faster implementation: Pre-tested NVIDIA blueprints accelerate deployment.
- Flexible, modular scaling: Expand seamlessly with consistent system management.
- Built-in sustainability: Energy-efficient components and advanced cooling solutions support greener operations in hospitality.

Why NVIDIA

NVIDIA is a pioneer in accelerated computing, delivering cutting-edge AI capabilities such as real-time video analytics, multilingual speech processing, and large language model deployment. With NVIDIA technologies, you gain:

- Advanced Vision & Speech AI: Leveraging NVIDIA Metropolis, NVIDIA Riva, and NVIDIA NeMo for robust applications.
- GPU-Accelerated Performance: Real time AI inferencing is essential for responsive and personalized guest interactions, ensuring a seamless experience.

[Download technical guide](#)

Lenovo Validated Design: Multimodal Agentic AI Solutions with Centific AI Data Foundry

