

MARKET PERSPECTIVE

Beyond Cost and Compliance: How Red Hat Is Shaping AI, Sovereignty, and Modernization

Jennifer Thomson Stuart Wilson Leonardo Freitas Neil Ward-Dutton
 Filippo Vanara Carla Arend

EXECUTIVE SNAPSHOT

FIGURE 1

Beyond Cost and Compliance: How Red Hat Is Shaping AI, Sovereignty, and Modernization

Red Hat's European Analyst Day, held in September 2025, provided an overview of its strategy, focusing on three major market trends: AI productization, virtualization cost and lock-in, and digital sovereignty. The vendor's core strategy emphasizes giving customers open choice and flexibility through RHEL and OpenShift. Central to the discussion was Red Hat's AI strategy, aimed at accelerating adoption from experimentation to production using open, portable tooling.

Key Takeaways

- Red Hat is focusing on three enterprise priorities: accelerating AI adoption from experimentation to production, enabling virtualization migration amid cost and lock-in concerns, and addressing digital sovereignty through open hybrid cloud architectures.
- The company's AI strategy emphasizes Kubernetes-native, open source solutions such as OpenShift AI and AI Inference Server, aiming to standardize high-performance inference and integrate automation via Ansible for governance and resilience.
- Red Hat leverages its open source ecosystem and partnerships to meet European sovereignty mandates, regulations, and sustainability goals, positioning sovereignty as a procurement differentiator in regulated markets.
- Red Hat has an opportunity to strengthen its position by formalizing AI governance frameworks, sovereignty blueprints, and provision of prescriptive migration and modernization playbooks.

Recommended Actions

- **Publish standardized migration playbooks.** Develop clear workload classification guidance (maintain/migrate/modernize) and best practices to accelerate virtualization migration and reduce risk.
- **Operationalize AI governance frameworks.** Develop governance blueprints with approval gates, rollback procedures, SLA definitions, and attestation workflows for AI artifacts to ensure trust and compliance.
- **Deliver regional sovereignty blueprints.** Create deployment guides tailored to EU, U.K., and APAC compliance regimes, covering data residency, GPU availability, and access controls to accelerate AI adoption in regulated markets.
- **Define measurable AI KPIs.** Introduce quantifiable inferencing performance metrics such as latency, throughput, and cost per query to help enterprises benchmark outcomes and justify AI investments.

Source: IDC, 2025

NEW MARKET DEVELOPMENTS AND DYNAMICS

This IDC Market Perspective provides IDC's assessment of Red Hat's strategic direction, product focus, and market positioning after attending the Red Hat Analyst Day in London in September 2025.

The company's strategy involves capitalizing on three dominant market trends: the explosive demand for enterprise AI, widespread customer migration from legacy virtualization platforms, and the critical need for digital sovereignty. Red Hat's core value proposition is delivering an open, flexible, and consistent hybrid cloud platform that provides customer choice and mitigates vendor lock-in.

Dominant market trends that compel strategic re-evaluation can be synthesized into three primary customer imperatives:

- **Accelerated AI adoption.** Demand for AI is growing exponentially, with a strong push by enterprises to move artificial intelligence from isolated experiments into core business operations. Customers are under intense C-suite pressure to leverage AI in driving measurable productivity gains, which requires platforms that can provide control and choice across the entire AI life cycle at enterprise scale.
- **Rethinking virtualization.** A confluence of rising license costs, complex licensing models, and concerns about long-term vendor relationships and potential lock-in has elevated virtualization strategy to a top priority for IT leaders. Organizations are actively seeking alternatives that reduce costs and provide a bridge to modern, cloud-native architectures.
- **The rise of digital sovereignty.** Sovereignty has evolved from a niche concern to a critical driver of investment decisions. Encompassing data residency, technological autonomy, and supply chain resilience, this trend is compelling organizations to re-evaluate their cloud strategies, repatriate critical workloads, and prioritize solutions that offer greater control and operational choice.

The AI Imperative: From Experimentation to Production

Red Hat's strategic goal for AI is to accelerate customer adoption from isolated experimentation to full-scale, integrated production, providing architecture and deployment choice and control. The company's approach is guided by the mantra "Any model, any accelerator, any cloud," reinforcing its commitment to providing an open, flexible, and portable platform for AI workloads that avoids lock-in to any single-model provider, hardware vendor, or cloud environment. Red Hat is evolving its portfolio to support this principle, with developments such as:

- **Red Hat OpenShift AI.** This solution provides the platform for building, deploying, and managing AI models and agentic workflows at scale.

- **Red Hat AI Inference Server.** This solution is built on vLLM (a performance-optimized inference engine for large language models or LLMs) and llm-d (a high-performance distributed LLM inference framework) and provides integrated access to open models via Hugging Face. The strategic aim of Red Hat AI Inference Server is to establish a de facto standard deployment platform for high-performance inference, akin to the roles Linux and Kubernetes play in their domains.
- **Red Hat Lightspeed suite.** This integrates context-aware generative AI (GenAI) assistance directly into Red Hat Enterprise Linux (RHEL), Ansible, and OpenShift to simplify operations and close talent gaps.
- **Neural Magic acquisition.** Red Hat's acquisition of AI optimization start-up Neural Magic enhances AI acceleration and inference optimization capabilities, supporting more efficient model execution.

Customer evidence highlights the real-world impact of this strategy. Examples shared during the event include the following:

- A global financial institution is building a secure, multitenant AI platform for both experimentation and production, with an ambitious goal of generating significant new revenue through AI-infused applications.
- In the energy sector, a leading enterprise is leveraging generative AI in its upstream operations to improve efficiency and achieve substantial cost reductions in its core business processes.
- In the public sector, a major European city is enhancing citizen services by integrating AI into its digital portfolio, including chatbots and automated workflows. This initiative began on managed cloud services and has since expanded to on-premises deployments for greater control.
- Similarly, a large international airline is using AI to improve passenger experience and fleet inventory management. By adopting a modern AI platform, the airline has doubled its model deployment speed and reduced data processing times, improving critical operations such as boarding and baggage turnaround.

Red Hat's pivot toward inference and agentic AI reflects a pragmatic understanding of where enterprise value lies — it is operationalizing AI at scale, not just training models. Its open source-first approach and Kubernetes-native architecture resonate strongly with CIOs seeking cost control and sovereignty in AI deployments. This emphasis aligns with a growing European priority: AI sovereignty. Organizations increasingly demand choice and control over data, models, and infrastructure, viewing sovereignty as a spectrum of control rather than an all-or-nothing model. Meeting this need requires deliberate architectural decisions, such as deploying hybrid cloud and composable systems, to mitigate risk while enabling flexibility.

The Virtualization Crossroads: From Cost Savings to Strategic Modernization

Customer interest in virtual machine (VM) migration is being catalyzed by urgent cost pressures and contract uncertainty. Red Hat recognizes the entry point is cost takeout, but its strategy intentionally elevates the buyer dialogue to innovation, security, and road map trust. Red Hat OpenShift Virtualization is being positioned not as a like-for-like hypervisor swap but as a strategic bridge to modernization, enabling VMs and containers to coexist on a single modern platform, preserving existing investments while accelerating adoption of container-native practices and AI integration without wholesale rewrites.

In terms of the VMware opportunity, IDC expects a wave of platform churn as enterprises reassess virtualization spend, lock-in, and operating models. Buyers want immediate TCO relief but are equally motivated by control, portability, and developer velocity. Against this backdrop, Red Hat's differentiation is strongest where migration and modernization converge: a single platform for VMs and containers; built-in GitOps and supply-chain security; and deployment choice across on-prem and managed cloud locations (ROSA/ARO/OpenShift Dedicated or OSD/Red Hat OpenShift on IBM Cloud or RHOIC). Competitively, Nutanix is capturing near-term VMware exits on HCI economics and migration tooling; Red Hat's opportunity is to position around "modernize while you migrate" and cloud portability, with emphasis on modernizing at the customer's speed. Red Hat empowers enterprises to modernize on their terms, migrating traditional workloads to a unified hybrid cloud platform, then evolving toward cloud-native and AI-driven innovation at their own pace, all while maintaining operational consistency across on-prem, cloud, and edge.

Red Hat is already packaging the migration path through its services-led Virtualization Migration Assessment (VMA) and partner-enabled programs. As the Red Hat builds on this foundation, IDC sees strong potential for the vendor to capture meaningful share from VMware exits and drive sustained adoption of OpenShift Virtualization as the landing zone for legacy VMs, particularly among organizations prioritizing sovereignty and multicloud strategies.

Key customer examples shared during the analyst day underscored this momentum:

- A major financial institution is undertaking a large-scale migration of its infrastructure, including core banking systems and thousands of virtual machines, to a modern virtualization platform. This move supports its multicloud strategy, leveraging multiple providers while maintaining a significant on-premises presence to ensure resilience and address geopolitical considerations.

- In the ecommerce sector, a leading regional platform adopted an open source–based virtualization solution to avoid the high costs of renewing a proprietary contract, improve platform availability, and reduce the risk of future vendor lock-in. This strategic shift reflects a growing trend among organizations seeking cost efficiency and architectural flexibility in their digital transformation journeys.

Automation as the Next Horizon

Modernization is only one part of the transformation story. To fully realize the benefits of hybrid cloud and container-native architectures, organizations must also automate at scale. Red Hat positions Ansible Automation Platform as the operational backbone for this part of the journey, extending automation from infrastructure provisioning to application deployment and compliance. By embedding event-driven workflows, policy as code, and integrations with tools such as HashiCorp Terraform, Red Hat aims to enable enterprises to move beyond ad hoc scripting toward mission-critical automation that enforces governance, accelerates remediation, and supports platform engineering practices. The company underscored the following announcements:

- **Event-driven Ansible.** This enables closed-loop remediation by automatically triggering playbooks in response to events from third-party monitoring tools.
- **Policy as code.** This provides guardrails and human-in-the-loop controls for GenAI-initiated automation, ensuring governance and compliance.
- **HashiCorp integration.** Ongoing work is focused on improving interoperability between Ansible and Terraform, as well as broader integrations between OpenShift, Ansible, and Vault.

Red Hat’s ongoing work to improve interoperability between Ansible Automation Platform and HashiCorp Terraform, alongside integration with Vault, positions it well to address this need. Features such as event-driven Ansible and policy as code reinforce Red Hat’s focus on enterprise-scale automation with governance and resilience built in.

Following IBM’s acquisition of HashiCorp, IDC expects the synergy between Ansible and Terraform to become a strategic differentiator. Customers are increasingly seeking unified workflows that span infrastructure as code, compliance enforcement, and closed-loop remediation. Red Hat confirms active development on interoperability and Vault integration. To fully capitalize on this opportunity, Red Hat and IBM must articulate a clear integration trajectory and demonstrate how joint innovation will simplify automation complexity for global enterprises.

The Sovereignty Mandate: Enabling Digital Autonomy

Digital sovereignty is no longer a niche concern — it is becoming a defining principle for technology strategy in Europe. The EU AI Factories initiative, with €12 billion in combined EU (€4 billion) and local government (€8 billion) funding, aims to build

sovereign AI infrastructure using high-performance computing (HPC) centers. Red Hat is positioned as a key partner in this transition, bridging the gap between traditional open source and enterprise-ready commercial models. At the same time, commercial organizations such as leading U.K. banks are recalibrating their risk profiles and shifting toward a more balanced on-premises footprint to strengthen resilience against hyperscaler dependency.

Discussions at IDC's European CIO Xchange in October 2025 confirmed this trend, with many organizations expressing concern over regulations that could enable government access to data. IDC survey data reinforces this sentiment: nearly nine in 10 organizations globally want choice and control when deploying AI at scale, and more than half prefer open models over closed, proprietary ones. Sovereignty is not just about compliance; it is about strategic flexibility, control, and trust.

Red Hat's philosophy is clear: sovereignty is achieved through community and innovation, not isolation. By leveraging transparency and control inherent in open source, Red Hat aims to enable customers to own their technology stacks, secure their software supply chains, and build a resilient digital future without lock-in.

Red Hat's sovereignty strategy is powered by its core platform technologies:

- **RHEL as the core.** Red Hat Enterprise Linux 10 introduces Image Mode, enabling the OS to be managed as an immutable, bootable container image. This cloud-native approach addresses regulatory pressure by reducing configuration drift, improving patch consistency, and accelerating upgrades to meet strict compliance timelines.
- **Automation as mission critical.** The Ansible Automation Platform provides event-driven remediation and policy as code, embedding governance into automation workflows initiated by GenAI tools.
- **Developer experience and supply chain.** Red Hat Developer Hub modernizes developer workflows, while initiatives such as the AI Bill of Materials (AI BOM) strengthen supply chain transparency for AI artifacts.
- **Sovereign AI as a strategic necessity.** Red Hat AI utilizes sovereign, self-managed deployments with open source, container-native tools that support GPU orchestration and multitenancy across clouds and on-premises stacks. This provides flexibility and efficiency. The AI platform utilizes vLLM and llm-d, offering high-performance, cost-effective, and scalable inference.
- **Ecosystem and partnerships.** Collaborations with OEMs such as HPE, Dell, and Lenovo ensure certified hardware for hybrid deployments, while interoperability with partners such as HashiCorp supports seamless integration into existing tool chains.

Sovereignty is emerging as a procurement differentiator in Europe. Organizations that combine open source transparency, hybrid flexibility, and strong governance

will be best positioned to meet compliance requirements while maintaining innovation velocity.

Other Initiatives Strengthening Red Hat's Market Position in Europe

Regulatory Landscape and Compliance

Regulation remains a critical focus area, with the EU Cyber Resilience Act emerging as one of the most significant concerns in EMEA. Red Hat is actively engaging with policymakers to provide feedback and ensure that the legislation is practical and does not inadvertently harm the open source ecosystem. IDC believes this proactive stance is essential to safeguard innovation while meeting compliance requirements. Recent discussions indicate that the situation is improving, thanks to these efforts — a positive sign for vendors and customers navigating complex regulatory landscapes.

Sustainability in AI and Computing

Sustainability continues to be a core priority, particularly as it becomes a key requirement in EU tenders as well as a growing concern due to the high energy consumption of GPUs. IDC views sustainability as a differentiator in enterprise AI strategies, not just a compliance checkbox. To address this, Red Hat employs techniques such as model quantization and compression to create smaller, more efficient LLMs, reducing energy use and costs. A customer achieved substantial savings through this approach, validating IDC's position that efficiency gains can directly support ESG goals. Infrastructure innovation also plays a role: for example, a customer operates a solar-powered datacenter running container platforms and issuing sustainability certificates, demonstrating how green computing can be integrated into enterprise environments.

IDC's 2025 *EMEA Multicloud Survey* reinforces this trend: cost optimization (39.6%) is the top driver for sustainability in cloud services, followed closely by IT sustainability strategies (37.5%) and regulatory compliance (25.8%). This indicates that organizations increasingly view sustainability as a lever for cost control and compliance, not just as an ESG obligation.

Red Hat is also investing in cloud-native sustainability tooling. Its open source Kepler project (launched in 2023), developed with IBM Research, estimates Kubernetes workload power consumption using eBPF and machine learning. Kepler provides actionable metrics for energy-aware scheduling, carbon footprint estimation, and sustainable CI/CD pipelines, enabling enterprises to optimize resource allocation and align with sustainability goals. These capabilities position Red Hat to influence emerging FinOps and GreenOps practices, which IDC expects to gain traction in Europe as organizations seek cost control, compliance, and ESG alignment.

Upskilling the Open Source Community

Open source is central for Europe's determination to achieve digital sovereignty, as its efforts ultimately rest on reducing its dependency on external technology providers and prioritizing transparent, interoperable, and secure solutions.

High-profile initiatives, such as The Important Projects of Common European Interest: Next Generation Cloud Infrastructure and Services (IPCEI-CIS) projects coordinated by the European Commission, as well as investments in domestic digital infrastructure by member states, have been advancing these ambitions for over a decade. These efforts offer European organizations a way to retain ownership and control over their technology stacks while ensuring regulatory compliance and fostering trust. In this scenario, upskilling within Europe's open source community is essential as digital technologies redefine enterprises' needs (flexibility on data ownership, lock-in avoidance) as well as societal needs (such as geopolitical).

However, the continuing gap between skill supply and new technology demands — particularly in cloud, AI, and cybersecurity — risk creating digital divides and curbing innovation. During the analyst summit, Red Hat reinforced its position to support these developments through strategic partnerships and investment in the European learning and development ecosystem for open source, especially because it sees the technology as a potential key enabler for AI adoption going forward.

Ecosystem and Partnerships

A strong ecosystem underpins the company's European strategy. Partnerships with leading hardware providers ensure certified solutions for on-premises and hybrid deployments, while collaborations with major cloud providers deliver optimized images for enterprise operating systems. IDC believes these alliances are vital for meeting diverse customer requirements across hybrid cloud and multicloud environments. The network extends to independent software vendors and global system integrators, enabling comprehensive industry solutions and accelerating adoption. Ecosystem strength is central to success in the AI era since customers increasingly demand integrated, interoperable solutions rather than isolated products.

Red Hat's ecosystem success increasingly hinges on its ability to drive customer engagement and alliance execution in tandem. While its open hybrid cloud model offers flexibility and sovereignty, it also challenges traditional partner incentives — particularly for global system integrators (GSIs) and hyperscalers whose business models often rely on deeper integration and workload retention. Red Hat's European Analyst Day discussions underscored that the vendor must lead with customer momentum to pull partners in, not rely on partners to generate demand.

This is especially true in regulated markets, where sovereignty, data residency, and infrastructure control are driving architectural decisions. Red Hat's partnerships with

OEMs, ISVs, GSIs, and hyperscalers remain vital, but the company must navigate a delicate balance: enabling partner profitability while preserving customer choice.

IDC sees an opportunity for Red Hat to strengthen its position by investing in vertical go-to-market strategies, demand orchestration, and partner monetization models that align with open source principles. The resurgence of local SIs and MSPs also signals a shift — ecosystem value is no longer concentrated solely in global alliances. To scale influence, Red Hat must ensure its platform is technically interoperable as well as commercially compelling across partner types, especially as AI, virtualization, and sovereignty reshape enterprise priorities.

ADVICE FOR THE TECHNOLOGY SUPPLIER

Clarify the Virtualization Strategy

While OpenShift Virtualization is positioned as a modernization bridge, IDC sees an opportunity for Red Hat to strengthen customer confidence through a standardized workload classification playbook (maintain/migrate/modernize). Such guidance would accelerate migration timelines and help customers manage costs predictably. IDC recommends that Red Hat focus on migration tooling and best practices to reinforce its differentiation against VMware alternatives and ensure smoother execution at scale.

Build Governance Blueprints for Scalable AI

Red Hat should accelerate the development of governance blueprints to maintain trust and predictability in AI scaling. The vendor's AI strategy demonstrates strong technical ambition, but there is a need for more evidence and visibility of a formal governance framework for cross-environment deployments. Without clear approval gates, rollback procedures, and SLA definitions, enterprises risk misalignment between experimentation and executive expectations.

Codify Governance for AI Scaling

Red Hat's concept of an "AI BOM" is innovative and aligns with the growing need for transparency in AI supply chains. IDC sees an opportunity for the vendor to strengthen this approach by introducing a formal schema, attestation workflows (structured processes for verifying compliance and certifying AI artifacts), and integration with SBOMs to support emerging regulatory requirements. IDC expects that by mid-decade, enterprises will increasingly demand verifiable AI artifact provenance as part of governance and audit readiness. Operationalizing these controls will enable Red Hat to reinforce its leadership in trusted AI and provide customers with confidence in accountability across environments.

Develop Regional Sovereignty Blueprints

IDC views sovereignty as a decisive factor in cloud strategy and sees an opportunity for Red Hat to strengthen its position by publishing region-specific deployment guides covering data residency, GPU availability, and access controls. Delivering sovereignty blueprints aligned with EU, UK, and APAC compliance regimes would accelerate AI scaling in regulated markets and reinforce Red Hat's hybrid cloud value proposition.

Tell an End-to-End AI Sovereignty Story

Red Hat must build on the work already done with Red Hat Inference Server and leverage the digital sovereignty story. The vendor must show how its products, services, and ecosystem partners empower customers to exert sovereign control over AI system design, development, deployment, operation, access, and governance.

Enable Developers and Modernize Playbooks

IDC sees developer experience as a cornerstone of successful migration and modernization. Red Hat can amplify its impact by providing prescriptive playbooks for "migrating while modernizing" and clarifying ownership of Internal Developer Portal rollout. These frameworks, combined with policy guardrails, would reduce misconfigurations and compliance risks while enabling developers to adopt modern practices with confidence.

Define KPIs for AI Inferencing Performance

When it comes to inferencing KPIs, Red Hat's focus on cost and energy efficiency is a strong starting point, but IDC sees an opportunity to deepen enterprise confidence by introducing quantifiable performance metrics, such as latency, throughput, and cost-per-query. These KPIs will help organizations justify AI investments and benchmark operational outcomes, reinforcing Red Hat's position as a trusted partner for scalable AI.

LEARN MORE

Related Research

- *Digital Sovereignty* (IDC #IDC_P45150, August 2025)
- *Red Hat Summit Conference: The Future of AI Is Open* (IDC #US53550425, June 2025)
- *Red Hat Summit 2025: Ushering in a New Era of Open AI-Enabled Platforms and Ecosystem Innovation* (IDC #US53551525, June 2025)
- *Red Hat Finds its Sweet Spot in AI With Boosting Inferencing Efficiency* (IDC #lcUS53544025, May 2025)

- *Red Hat's AI Strategy: Enabling the AI Technology Operating Model for Enterprises' AI Pivots* (IDC #AP51583224, December 2024)

Synopsis

This IDC Market Perspective examines Red Hat's evolving strategy in response to three dominant market trends: accelerating enterprise AI adoption, widespread migration from legacy virtualization platforms, and the growing imperative for digital sovereignty. Following Red Hat Analyst Day 2025, Red Hat is positioning itself as a leader in delivering open, flexible, and consistent hybrid cloud platforms that prioritize customer choice, operational control, and mitigation of vendor lock-in.

"In a market chasing speed, Red Hat's open approach reframes the conversation from lock-in avoidance to long-term resilience, where sovereignty and AI governance become strategic imperatives. Red Hat's strength lies in bridging modernization and sovereignty while operationalizing AI at scale," said Jennifer Thomson, AVP European Research.

ABOUT IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology, IT benchmarking and sourcing, and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives. Founded in 1964, IDC is a wholly owned subsidiary of International Data Group (IDG, Inc.).

IDC Italy

Viale Monza, 14
20127 Milan, Italy
+39.02.28457.1
Twitter: @IDCItaly
idc-insights-community.com
www.idcitalia.com

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, and web conference and conference event proceedings. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/about/worldwideoffices. Please contact IDC at customerservice@idc.com for information on additional copies, web rights, or applying the price of this document toward the purchase of an IDC service.

Copyright 2025 IDC. Reproduction is forbidden unless authorized. All rights reserved.