AI summary

Sovereignty, Tokenomics and the New Agentic Operating Model

ATX ENTERPRISE VOXO

Sovereignty, Tokenomics and the New Agentic Operating Model
Wednesday, 20 May 2026

Participants

Steve Shirkey
Director, APAC AI Platform, Red Hat

Sovereignty, Tokenomics and the New Agentic Operating Model
Wednesday, 20 May 2026

Summary

The session explored emerging paradigms in AI and tokenomics, emphasising the shift from traditional labour-intensive models to AI-driven operating frameworks. The speaker discussed the importance of focusing on outcomes rather than workforce size, highlighting the potential of AI to achieve greater efficiency with fewer resources. Citing an 81% success rate of agent hijacking according to an NIST study, the speaker underscored the need for robust guardrails, evaluations, and foundation structures to ensure AI systems align with organisational goals. Red Hat's enterprise story was used as a backdrop to illustrate transitions in the tech landscape, from UNIX to Linux, monolithic to microservices, and now to agentic, AI-native applications.

Red Hat's AI platform was presented as a comprehensive solution offering flexibility and scalability. By abstracting over various accelerators like NVIDIA and supporting multiple operational environments from edge to cloud, the platform provided customers with choices tailored to their needs. The speaker highlighted successful case studies, such as Argentina's RSAT achieving 30% operational cost reductions and the UK Ministry of Defence unifying its AI capabilities. The platform also integrated advanced model options, including Meta's Lama models and open-source solutions, enabling users to operate effectively across diverse infrastructures. Observability, governance, and security were emphasised as critical components to managing agentic AI systems efficiently.

The challenges of cost inefficiencies in AI implementation were also addressed, with the speaker cautioning against uncontrolled spending during scaling and production phases. They recommended an ownership model where organisations manage their AI infrastructure to maintain cost predictability and quality standards. The Red Hat AI Factory, in partnership with NVIDIA, was showcased as a solution to maximise hardware ROI and streamline day-two operations. The session concluded with an invitation to explore demos and resources, reinforcing the importance of adaptability in navigating regulatory and technological shifts. The audience was encouraged to prioritise owning their AI systems to future-proof against evolving constraints and ensure strategic alignment.

Takeaways

Guardrails and Evaluation Are Critical for AI Success

The session emphasised the importance of instituting guardrails and thorough evaluations to prevent AI systems from deviating from their intended purpose. These measures ensure that AI operates within defined boundaries, aligns with organisational goals, and mitigates risks of misaligned or insecure deployments.

AI Ownership and Infrastructure Flexibility Are Essential

The speaker underscored the necessity of owning and controlling AI to navigate regulatory changes and optimise costs. A flexible platform capable of supporting diverse models, accelerators, and deployment environments empowers organisations to adapt quickly and maintain operational continuity.

Cost Efficiency and Observability Are Key to Sustainable AI Operations

The presentation highlighted the risk of unsustainable costs when transitioning AI pilots to production without proper planning. Observability, cost control mechanisms, and thoughtful infrastructure choices were stressed as critical enablers for scalable and efficient AI systems.

AI summary powered by

voxoevent.ai