

ECOSYSTEM VENDORSPPHERE

AWS Analyst Summit: Highlights from Singapore

At the AWS Analyst Summit in Singapore, AWS leadership outlined their vision for AI and how they are helping customers keep pace with rapid technological change as AI — particularly agentic AI — moves out of the hype cycle and into practical deployment.

Ecosystem Analysts [Sash Mukherjee](#) and [Louise Francis](#) share their take on the announcements, demos, and messaging, highlighting what resonated, what stood out, and what it means for the future.

Here are 5 key takeaways from our analysts:

01. AWS view of AI as an enterprise operating layer

AWS is treating AI as the operating layer for enterprise systems, embedding intelligence directly into workflows, system interactions, and decision flows rather than isolating it at the interface. The emphasis is on connecting and governing AI across the stack so it can operate reliably in production, shifting the focus from standalone AI tools to environments where intelligence is integrated into core operations.

This shows up in AWS's framing of the shift from models to integrated systems; from copilots to agents that can coordinate actions and retain context over time; and from isolated pilots to AI embedded into continuous production workflows. The assumption is broad adoption; the differentiator is execution at scale within existing enterprise constraints.

AWS's stack is positioned around that reality rather than individual products.

- [Amazon Bedrock](#). Model access layer providing multiple foundation models (including Anthropic, Meta, Mistral, Cohere, Stability AI, Amazon Nova, and OpenAI in limited preview), allowing organisations to switch or combine models without changing core integrations and accelerating application deployment. Bedrock Guardrails adds configurable safeguards for responsible AI, while Bedrock Knowledge Bases enables RAG grounded in enterprise data.
- [AgentCore](#). Runtime layer for AI agents, adding memory, observability, and policy controls so agents can operate across systems with governance and monitoring in production.
- [Kiro](#). AI-native development environment where agents generate, structure, and modify code from specifications, reducing design and documentation effort. Kiro succeeds Amazon Q Developer as AWS's primary AI-assisted development tool, using spec-driven development to convert natural language into structured requirements, design documents, and implementation tasks.
- [Amazon Quick](#) (including Desktop App). Enterprise AI assistant for daily work, supporting search, summarisation, workflow execution, and content creation, with a desktop layer that brings enterprise data into the user's workspace for outputs such as summaries and presentations.

- [AWS Trainium & Inferentia](#). Custom AI silicon designed for AI training and inference workloads, intended to improve price-performance efficiency (up to 50%) compared with conventional GPU-based infrastructure.
- [Amazon SageMaker HyperPod](#). Managed infrastructure for large-scale distributed model training with built-in resilience, automatic fault recovery, and optimised cluster management.

ECOSYSTEM OPINION

AWS's approach broadens the competitive landscape beyond models and infrastructure into workflow platforms, SaaS ecosystems, and systems integrators that shape how AI is embedded into enterprise operations. Rather than competing only at the model layer, AWS is positioning itself around an integrated agentic AI stack spanning infrastructure, orchestration, tools, and workflows.

The key question is how effectively this evolves into a widely adopted enterprise AI experience layer, especially as competitors increasingly embed AI into productivity, collaboration, and customer experience tools. For AWS, the opportunity lies in making their agentic foundation more accessible and pervasive across enterprise environments.

A parallel constraint is organisational readiness. Many enterprises are still adapting their data, integration, governance, and operating models for agent-led workflows. This makes AWS's positioning relevant not just as infrastructure, but as a connective layer for operationalising AI across fragmented systems.

02. Agents as an accelerator for modernisation & migration

AWS positions agentic AI as infrastructure for change, going beyond the productivity layer. Agents are treated as persistent systems that operate across enterprise environments, supporting legacy modernisation, workflow coordination, self-monitoring, and adaptation within defined policy guardrails.

The focus is on operational continuity across fragmented systems rather than conversational interfaces.

This is reflected in the stack:

- [AgentCore](#)
Runtime layer for agents, providing memory, observability, and control mechanisms.
- [AWS Transform](#)
Agentic automation for large-scale migration and modernisation, including VMware estates, mainframes, Windows environments, and code transformation.
- [NLX \(acquired by AWS in April 2026\)](#)
No-code conversational AI platform integrated into Amazon Connect, enabling enterprises to build and deploy AI-powered self-service experiences with less implementation effort and shorter deployment timelines, while simplifying the scaling of conversational and agentic systems.

Agents are the link between AI capability and enterprise modernisation, shifting migration, refactoring, and workflow integration from discrete projects to continuous execution.

ECOSYSTEM OPINION

Agentic AI adoption is likely to accelerate first within technology functions, where familiarity with automation, APIs, and system integration is already more established. Ecosystem research shows that 68% of organisations in Asia Pacific are exploring agentic AI for technology functions, followed by customer experience (62%) and HR (48%), highlighting that adoption remains stronger among technical teams compared to other business functions.

Scaling beyond early adoption will depend on cost (including token and compute economics), infrastructure readiness, and clear accountability for autonomous actions in production environments. These constraints will determine how quickly agent-led models move from pilots to enterprise-wide deployment. While AWS presents a strong agentic AI narrative, some customers may still view their positioning as infrastructure-led, with narrower industry applicability and uneven availability across parts of Asia Pacific.

03. Amazon Connect as an operational engagement layer

Amazon Connect has moved beyond its traditional contact centre role into a broader layer for coordinating enterprise operations. It now spans environments where work is distributed across people and systems and responsiveness, timing, and consistency are critical. The target is human-intensive workflows that require real-time interaction and cross-system coordination.

Areas of Expansion:

Amazon Connect Decisions

Agentic AI supply chain planning and intelligence solution, succeeding AWS Supply Chain. AI teammates are designed to align with operating procedures, integrate with existing systems, and learn from practitioner decisions. Capabilities include demand forecasting, constraint-aware supply planning, exception detection, and operational recommendations, with applications across retail, CPG, automotive, and manufacturing.

Amazon Connect Talent

Designed to support recruitment and workforce coordination for high-volume, time-bound hiring such as seasonal demand, where manual processes do not scale. AI agents conduct voice interviews around the clock and score candidates on competencies; recruiters review AI-generated scores, full transcripts, and supporting reasoning before making final decisions. The system is based on Amazon's hiring science derived from processing 250,000+ seasonal hires.

Amazon Connect Health

Developed to reduce administrative workload in healthcare workflows, including patient verification, appointment scheduling, ambient clinical documentation, medical coding, and post-visit summarisation, integrating with existing EHR systems. It also enables pre-consultation patient insights to support clinicians, with deployments in the US and selective expansion into regulated markets.

Connect is being applied in operationally dense environments where the constraint is coordination across people, processes, and systems, rather than within a single functional domain.

ECOSYSTEM OPINION

AWS is following a familiar pattern: internal Amazon systems originally built to operate at massive scale are being productised for external customers, particularly in areas where organisations struggle to move from experimentation to enterprise-scale outcomes. Amazon Connect reflects this approach, extending capabilities developed for Amazon's own distributed operational environment into enterprise use cases.

The advantage is clear: these are proven, high-scale systems now applied to broader enterprise domains. However, as AWS expands further into workflow-intensive areas such as talent, healthcare, and enterprise operations, it increasingly overlaps with established SaaS and industry software ecosystems. Success will depend on how effectively AWS balances leveraging its internal IP while coexisting alongside, rather than displacing, existing partners in these markets.

04. Sovereignty & resilience as foundations for production AI

Across sessions, sovereignty and resilience were consistently treated as prerequisites for AI in production rather than compliance considerations. Digital sovereignty was discussed in operational terms: data control, encryption and key ownership, regional isolation, access transparency, and incident detection and response in increasingly distributed and autonomous systems.

This was reinforced through broader discussions on data foundations and operational readiness. Sovereignty combines data sovereignty and operational sovereignty — who controls data and where it resides, and whether systems remain resilient, observable, and independently operable under stress. The underlying message was that scaling innovation depends on this balance, not a trade-off between control and capability.

AWS anchored this through a set of core security and infrastructure services:

- ▶ [AWS Nitro System](#). Secure compute foundation providing hardware-level isolation and workload protection, with no mechanism for AWS operators to access customer instances. Independently validated by NCC Group.
- ▶ [AWS KMS \(Key Management Service\)](#). Centralised encryption key creation and control for data protection
- ▶ [AWS CloudHSM](#). Dedicated hardware security modules for stronger key ownership and compliance needs
- ▶ [AWS Certificate Manager](#). Management of TLS/SSL certificates for secure communications
- ▶ [Dedicated Local Zones](#). In-country/near-user infrastructure for latency-sensitive and residency-controlled workloads
- ▶ [AWS AI Factories](#). Fully managed AI infrastructure deployed in customer-owned data centres. Customers provide data centre space and power; AWS deploys and manages the integrated infrastructure, supporting sovereign AI model training and inference while meeting data residency and regulatory requirements.
- ▶ [AWS Clean Rooms](#). Enables multiple parties to analyse combined datasets without exposing underlying data, supporting privacy-preserving collaboration for AI/ML workloads
- ▶ [Amazon Bedrock Guardrails](#). Configurable safeguards to filter harmful content, enforce topic boundaries, and redact sensitive information in AI model inputs and outputs

This was also reflected in how AWS positioned the cloud more broadly — less as infrastructure, and more as the environment where AI systems are secured, governed, and operated. The idea that “the cloud becomes an agent platform, not just a compute platform” captures this shift, where sovereignty, resilience, and observability become embedded requirements as systems move toward more autonomous, agent-driven operations.

The emphasis shifts from individual capabilities in the stack to whether the overall environment can support trusted, production-scale AI systems operating continuously across data, applications, and infrastructure.

ECOSYSTEM OPINION

This is particularly relevant for Asia Pacific markets, where uneven infrastructure maturity, regulatory diversity, and talent constraints make sovereignty and resilience practical enablers of AI adoption. AWS’s Digital Sovereignty Pledge reflects this direction, focusing on data residency controls, verifiable access, encryption across states, and resilience by design.

This is also evident in enterprise modernisation efforts, such as a banking example shared during the session where approximately 60% of compute has already moved to the cloud, with remaining core systems planned for migration within the next two years. In parallel, data has been separated into dedicated decision layers processing billions of data points daily, supported by AWS-led capability building to help address skills gaps.

AWS’s sovereignty positioning is a competitive strength at a time when geopolitical resilience is a growing priority for organisations. However, many enterprises in Asia Pacific still face foundational challenges around data quality, integration maturity, and security capability. AWS’s success will depend on how effectively it helps customers strengthen these foundations through partner ecosystems, training, and sovereignty-aligned architectures that can adapt to local constraints.

05. AWS vision & ambition in Asia Pacific (APJC)

AWS sees APJC as a core driver of their global strategy, not a downstream growth market. With nearly half of the world’s developers expected to be in the region by 2029, it is increasingly viewed as where cloud and AI patterns will be built, tested, and operationalised at scale.

This is reflected in their commitment of over [USD 40 billion in cloud and AI infrastructure across 14 APEC economies from 2025-2028](#) and a move toward a more unified regional operating model, despite significant variation in regulation, maturity, and infrastructure readiness across countries. Additionally, planned [investments in Southeast Asia alone are expected to reach over USD 33 billion](#) by 2039.

Ecosystem scale is central to execution. [AWS has trained 9.3 million people in the region](#) and continues to anchor adoption through innovation hubs and applied programmes such as hackathons for AI-native development. Partners remain critical in bridging capability gaps, especially in markets where internal skills and delivery capacity are uneven.

The approach is deliberately demand-led, shaped by enterprise and government requirements rather than a single market template. This includes country-specific sovereign-aligned deployments and continuous engagement with regulators and public sector bodies to align on operating requirements.

ECOSYSTEM OPINION

AWS's approach — combining large-scale infrastructure investment, partner-led delivery, and sovereignty-aligned deployments — reflects a key reality in the region: AI adoption in Asia Pacific is shaped as much by ecosystem maturity and delivery capacity as by technology availability.

For AWS, the near-term opportunity lies in enabling workflow transformation through the infrastructure and orchestration layers that underpin agentic systems, particularly across operations, customer engagement, and coordination-heavy enterprise processes. This includes connecting data systems, supporting integration across fragmented environments, and providing the tools and services that allow agents to be embedded into existing enterprise workflows rather than replacing them outright.



About Ecosystem

Ecosystem is a leading technology market analyst and advisory firm that helps stakeholders navigate innovation in the digital economy through data, insights, and expertise. We connect enterprises, technology companies, digital-native founders, investors, and policymakers to enable informed decision-making. With ongoing research and access to top analysts and strategic advisors, we empower business planning, go-to-market activities, thought leadership, and innovation strategy consulting. Visit ecosystem.io