

The Innovation Lab as a Strategic Sales Motion

Analyzing the GTM Strategies of AWS, Microsoft, and Google

Katharine English • November 2025

Adapted from an internal strategy paper. Proprietary data, internal program details and company-specific recommendations have been removed.

Portions of the background research and data aggregation for this document were facilitated by Large Language Models (LLMs). While these tools provided a foundation for efficiency, the final analysis, critical thinking, validation, authorship and editorial decisions reflect the independent judgment of the author.

I. Executive Summary

The “Innovation Lab” model deployed by the cloud hyperscalers is not a research initiative; it is the new pinnacle of the strategic sales motion. These programs are sophisticated, top-of-funnel conversion engines designed to bypass traditional procurement, embed platform technologies, and secure long-term, high-margin consumption revenue. They represent a fundamental shift in enterprise sales, where the initial consulting is subsidized by the platform to accelerate adoption and secure platform lock-in.

For a hyperscaler, the Co-Innovation Lab is a **Loss Leader**. They invest the time of design thinkers, engineers and architects at no cost to the customer because the *output* of the lab (a validated, scalable architecture and accompanying solution(s)) ensures a long-term stream of high-margin consumption that far outweighs the cost of the sprint.

Co-innovation labs transform the customer relationship from **Transaction-based** (buying credits) to **Outcome-based** (burning credits to solve problems). By compressing the innovation cycle, labs reduce the risk of consuming new services and create a programmatic pipeline for cloud adoption.

This market is defined by three divergent strategic bets:

1. **AWS’s Two-Pronged Approach:** AWS focuses on two different models: (1) the **Generative AI Innovation Center** (GenAIIIC), which deploys a high-touch, “no-cost” management consulting approach; and (2) **Digital Innovation Leads**, who deliver rapid solutioning sessions to quickly produce working prototypes. AWS does not typically use the term “Innovation Lab” as a physical destination. Instead, it offers a mobile approach to its **programs** that brings the innovation process to customers (though frequently at an AWS site).
 - **GenAIIIC Enterprise Strategists:** These are typically former CIOs, CTOs, or CISOs hired by AWS. They don’t just “sell”; they are deployed to customer executives to help build the *business case* for cloud and AI. They are “embedded” in the sense that they act as high-level consultants to the VP and C-suite to unblock culture or strategy issues, but they do not build prototypes. The goal is

to de-risk transformation and build deep executive alignment *before* a proof of concept (PoC) is built, driving high-value, long-term cloud commitments.¹

- **Digital Innovation Leads (the “Working Backwards” specialists):** These teams (a mix of AI strategists and hands-on technical developers) run “**Working Backwards**” **workshops** with customers. They don’t start with technology; they start with a customer problem and write a “Press Release” (PR) and “Frequently Asked Questions” (FAQ) document, a well-known Amazon mechanism. The goal is to define exactly what the “Minimum Lovable Product” (MLP) is before building a working prototype.
 - **The Model:** “Co-Build” engagements. They do not build *for* the customer; they build *with* the customer’s developers.
 - **The Outcome:** A working prototype (not production code, but functional) and a “build book” or asset transfer so the customer’s team can take it to production.
 - **Cost:** Fully or partially funded by AWS (using credits or investment funds) to unlock consumption.
- 2. **Microsoft’s “Rapid Sprint” Model:** Microsoft’s traditional **Technology Centers (MTCs)** have been augmented by a newer tier called **AI Co-Innovation Labs**: a high-velocity, engineering-led, “complimentary” engagement. A one-week sprint delivers a customer-owned, platform-locked prototype, creating a powerful flywheel of public case studies and rapid Azure adoption.³
- 3. **Google’s “Platform-as-a-Lab” Model:** Google Cloud does not have a formal, public lab program comparable to its competitors’, instead leveraging its unified Vertex AI platform as a holistic “innovation environment.” This product-led strategy aims to drive adoption through model choice and scales through a partner-centric “Sell, Service, Build” ecosystem, supported by executive briefings.⁵

This whitepaper details how these “free” or “complimentary” engagements are a calculated Customer Acquisition Cost (CAC) designed to generate a threefold revenue stream: (1) securing long-term **cloud commitments**, (2) driving **consumption revenue** (e.g., Amazon Bedrock, Azure OpenAI Service, Vertex AI and BigQuery), and (3) scaling through a franchised partner ecosystem.

It closes with the implications for any platform vendor that relies on partners alone for early-stage innovation, and makes the case for a dedicated rapid-prototyping team that handles the lower-risk, high-speed “Day 0 to Day 5” work to secure the technical win, then hands the “Day 6+” production build and scale-out to professional services or a partner. **The vendor owns the relationship and the architecture; the partner owns the scale.**

II. The Rise of the “Innovation Lab” as a Strategic Sales Motion

Deconstructing GTM: From Executive Engagement to Consumption Revenue

The generative AI market is in a state of high stakes and high confusion. While 67% of AI decision-makers plan to increase investment in GenAI, CIOs report they are “drowning in AI consulting pitches” from every vendor.⁸

The “Innovation Lab” model is a direct GTM strategy to cut through this noise. It is a top-down, C-suite-

sponsored engagement that *replaces* the traditional consulting pitch with a tangible, high-value, *subsidized* engagement.

This strategy fundamentally alters the competitive landscape. The hyperscaler is effectively using its high-margin platform revenue to *subsidize* (or give away for free) the initial, high-cost consulting phase: use-case identification and prototyping. **This creates a branded value proposition that traditional System Integrators (SIs) and consulting firms cannot match.** As a result, this motion forces consultants to “invest more in alliance relationships and co-innovation capabilities,”⁹ effectively subordinating their services to the hyperscaler’s platform. **The “lab” is the tool by which the hyperscaler takes control of the enterprise’s transformation roadmap, moving from technology vendor to strategic partner.**

The Threefold Revenue Path: Commitments, Consumption, and Consulting

The “free” lab engagement is a calculated investment designed to generate revenue through three distinct channels.

1. **Commitments:** The initial lab engagement is designed to *manufacture the business case*. Methodologies like the AWS GenAIIIC’s “Five Vs Framework” have a “Value” phase dedicated to identifying ROI *before* building.¹⁰ This jointly developed ROI justification gives the C-suite the business case required to sign multi-year Enterprise Discount Programs (EDPs) or other large-scale cloud commitment contracts.
2. **Consumption Revenue:** This is the primary target. The PoCs built in these labs are *not* platform-agnostic. They are constructed using specific, high-margin, proprietary services (e.g., Amazon Bedrock, Azure OpenAI Service, Google Vertex AI and BigQuery). Forrester notes the boom in generative AI is “obvious” for hyperscalers who “host GPUs.”¹¹ The lab is the GTM motion to fill that GPU capacity. **The “free” PoC is the hook that pulls in recurring consumption revenue for these high-margin AI and data services.**
3. **Consulting Revenue:** The hyperscalers’ in-house expert teams (like AWS’s Working Backwards teams) are expensive and unscalable. The *real* long-term consulting revenue is *partner-led*. Programs like the AWS “Partner Agent Factory”¹² and Google’s “Partner Advantage”⁵ are designed to *train and franchise* the lab’s methodology. This allows partners to execute the full-scale build-out and implementation, which in turn drives consumption revenue back to the hyperscaler’s platform.

III. The AWS Model

From Diagnosis to Deployment: A Deep Dive into a Three-Phase Engagement

The AWS methodology is fundamentally consultative, mirroring a high-end management consulting engagement.¹ The process is gated to ensure resources are spent on high-potential customers.

1. **Phase 1: Diagnostic:** An initial discovery session to learn more about business needs and identify stakeholders.¹ This is the qualification gate where AWS assesses the customer’s *potential* (i.e., long-term consumption value) to decide whether it warrants the no-cost investment, and whether the customer is ready to commit the required IT and business resources to see the engagement through.

2. **Phase 2: Discovery:** A multi-day workshop where AWS experts “identify and prioritize AI use cases” and “develop an execution plan.”¹ This is where strategic, C-suite buy-in is built.
3. **Phase 3: Prototyping:** The rapid PoC phase, where AWS specialists partner with the customer to co-build a prototype or PoC, provide advisory support, and/or customize a model.¹

The key differentiators are the “deeply embedded in-house strategist & data science teams” and the co-build rapid prototyping teams. The “product” AWS is selling is subsidized access to its “specialist AWS Applied Scientists and ML engineers.”¹ This is a higher-cost model, but one that de-risks longer-term pilots and builds immense executive trust by focusing on business strategy first.

The “Five Vs” Framework: A Methodology for Business Value

The core intellectual property of the GenAIIC approach is its “Five Vs Framework” (Value, Visualize, Validate, Verify, Venture), a field-tested methodology that shifts the focus from “What can AI do?” to “What *do we need* AI to do?”¹⁰ This framework is a direct *antidote* to the “backfiring” productivity initiatives Forrester warns about.⁷

The very first step, **Value** (“Is this worth solving?”¹⁰), acts as a critical de-risking mechanism. AWS is willing to move more slowly and (potentially) *not* build a PoC if the business value isn’t clearly identifiable. This is a defensive strategy: a failed, low-value PoC or pilot creates a *negative* customer experience and jeopardizes the long-term, multi-million-dollar consumption contract. This “value- and prototype-first” approach is designed to protect the lifetime value (LTV) of the enterprise account.

The Economics of the \$100M Investment: How “No-Cost” Engagements Convert

In 2025, AWS announced an *additional* \$100 million investment in the GenAIIC, two years after its 2023 launch.² This \$100M is effectively a *Customer Acquisition Cost (CAC) fund*. It pays the salaries of the global team of AI scientists, strategists, and engineers² who “embed directly with customer teams.”¹

The “no-cost” engagement is the *mechanism* for deploying this fund. AWS bets an illustrative \$50,000 no-cost engagement (in the form of expert time) on a qualified, high-potential customer like Ryanair, Formula 1, or S&P Global.² The *return* on that bet could be (e.g.) \$5 million in *new* annual consumption of Amazon Bedrock,¹³ Amazon SageMaker, and foundational data services.

This model is predicated on the fact that “**the most successful AI implementations start with a strong data and cloud foundation.**”² The GenAIIC engagement often *starts* with an AI discussion but *results* in “centralizing their data lakes...and modernizing their analytics capabilities,”² which is the *real*, high-volume consumption prize.

Scaling via Partners: AWS’s “Partner Agent Factory” (PAF) and Its Revenue Model

The embedded-scientist model is highly effective but unscalable. The “Partner Agent Factory” (PAF) is the solution.¹² The PAF is a co-innovation program with “Partner Innovation Alliance” (PIA) members.¹² This is a *franchise model* in practice:

1. AWS trains a PIA partner on the Working Backwards and GenAIIC methodologies.¹²

2. AWS and the partner co-innovate.¹²
3. “Production-ready solutions” are then *published to the AWS Marketplace*.¹²

This creates two new revenue streams for AWS:

1. **Consulting Revenue:** The partner, now an expert in this *repeatable* solution, can sell it to *dozens* of other customers *without* direct AWS involvement.
2. **Marketplace Revenue:** AWS takes a cut of every sale on the Marketplace and, more importantly, *every* new customer the partner lands drives consumption revenue back to AWS. This model scales the high-touch Working Backwards and GenAIIIC methodologies far beyond AWS’s internal capacity.

Evidence and Wins: Analyzing Engagements

The cited case studies reveal the true nature of the wins. The *entry point* is AI, but the *business result* is large-scale cloud migration, data lake centralization, and analytics modernization: the core, high-volume AWS services.

- **Experian:** “Accelerates .NET modernization using Agentic AI.”¹⁴ The win is *modernization*, with AI as the catalyst.
- **BMW Group:** “Manages 23 million vehicles...hosted on AWS”² and “migrates data lake to AWS cloud.”¹⁴ The win is the *massive* data lake and cloud back-end contract.
- **Jabil:** “Achieved a 74% reduction in data processing times.”² The win is a foundational *data* optimization.

IV. Microsoft’s “Rapid Engineering Sprint”: The Co-Innovation Labs

The Four-Step Journey: From Application to One-Week Complimentary Sprint

Microsoft’s model is built for sales velocity, designed to move from “envisioning” to a tangible prototype in weeks.³ It is the *inverse* of AWS’s “strategist-led” approach, opting instead for an “engineering-led” sprint.³

1. **Step 1: Application:** A simple online form.³
2. **Step 2: Onboarding:** A “Streamlined Onboarding” journey that takes two to three weeks for qualification and scoping.³
3. **Step 3: Design:** The customer and Microsoft “schedule and design the sprint.”³
4. **Step 4: Sprint (one week):** A “complimentary, personalized development sprint” where Microsoft “AI experts collaborate directly with your team...to build a working prototype.”³

This model skips the “Value” and “Visualize” phases of the AWS model and jumps directly to “Validate.” The goal is speed: “reduce time to market”³ and deliver a *tangible asset* in five days.

The “Complimentary” Revenue Engine: IP Ownership and Azure Lock-In

The business model is the program’s most disruptive component. As detailed at the UWM lab opening, “Participation is free of charge, and companies retain full ownership of their intellectual property.”¹⁵

This “free IP” model is a carefully designed psychological and strategic mechanism that converts a “complimentary” engagement into bankable revenue:

1. **“Free”** eliminates procurement friction and budget approvals.
2. **“Co-innovation”** (getting the customer’s own *engineers* into the lab³) creates internal champions and *transfers* product knowledge.
3. **“Full IP ownership”** is the linchpin. The customer *cannot* simply throw the prototype away; it is now a *company asset* that a manager has to justify in their next review.
4. **The Lock-In:** This new company asset (e.g., the Renaissance AI assistant⁴) was built *natively and exclusively* on Microsoft’s platform, using services like Azure AI Speech and Azure OpenAI.⁴
5. **The Result:** To use the IP they *legally own*, customers *must* deploy it on Azure, driving immediate and long-term **consumption revenue**. The “complimentary” sprint is a relatively low CAC for a mostly *guaranteed* new production workload on Azure.

Case Study Deep Dive: The Renaissance Logistics Prototype

The Renaissance case study is the textbook example of Microsoft’s model in action.⁴

- **Business Problem:** A small, tactical, *solvable* problem: “drivers were sometimes getting stuck during check-in.”⁴
- **Sprint Type:** “Rapid prototyping engagement.”⁴
- **Solution:** A “voice-enabled AI assistant” built in one week using Azure AI Speech Services and Azure OpenAI Service.⁴

Renaissant’s Chief Product Officer said the lab “accelerated our roadmap and enhanced the quality of our outcomes.”⁴ Microsoft traded one “free” week of engineering time for a new, permanent, consumption-driving feature in Renaissance’s core platform.

Market Evidence: Why Microsoft Dominates Public Case Studies

The quantitative *proof* of this strategy’s success comes from an October 2024 IoT Analytics report.¹⁶

Table 1: Publicly Cited Cloud GenAI Case Study Market Share (Q4 2024)

Hyperscaler	Number of Public Case Studies	Percentage Share
Microsoft	127	62%
Google	37	18%
AWS	33	16%

Data sourced from IoT Analytics Global Cloud Projects Report 2024.¹⁶

The “Rapid Sprint” model is a *case study generation machine*. A one-week sprint with a defined prototype output is far more *repeatable* and *publishable* than AWS’s longer, strategic, and often confidential Value workshops. This creates a *powerful marketing flywheel*:

1. The Co-Innovation Lab generates a quick, tangible win.⁴

2. The win is published as a case study.
3. The market fills with evidence positioning Microsoft as the leader (62% share¹⁶).
4. An executive sees this overwhelming social proof and is *more likely* to choose Microsoft for their AI project, feeding the flywheel.

The analysis supports this, noting that Microsoft's case study share significantly outperforms its actual cloud market share, *suggesting the lab is an effective, asymmetric GTM weapon*.¹⁶

V. Google's "Platform-as-a-Lab": Vertex AI and the Partner-Led CoE

A Strategy of Choice: The Four Pillars of Google's GenAI GTM

The research revealed no formal, public "Google Co-Innovation Lab" program analogous to its competitors'. Google's core GTM strategy, as articulated by Google Cloud CEO Thomas Kurian, is a *platform-centric* and *product-led* strategy.⁶ Google is betting that, in the long run, enterprises want *choice*, *openness*, and a *superior, unified platform* (Vertex AI). Its "lab" is the Vertex AI platform itself.

The "Four Reasons" customers choose Google are:⁶

1. **Vertex AI:** The "only unified platform" for the entire AI lifecycle.
2. **Agents:** Purpose-built agents for productivity.
3. **Model Customizability & Choice:** "First-party, third-party, and open models."
4. **Grounding:** Connecting models to "enterprise truth" (robust data and apps).

Incentivizing the Ecosystem: The "Sell, Service, Build" Partner Advantage Model

If the "lab" is the platform, the *consulting* is primarily outsourced to partners. The Google Cloud Partner Advantage program is the scaling engine.⁵ The model has three tracks: "Sell," "Service," and "Build."⁵

1. Google uses partner discounts and other financial incentives⁵ to encourage partners to build solutions on Vertex AI.
2. Google then co-markets these partner solutions.
3. Customers engage the *partner* for the "Service" (consulting), and the partner's solution drives **consumption revenue** for Google (Vertex AI, BigQuery). Google is *enabling* a partner-led consulting ecosystem rather than *competing* with it.

Engaging Executives: "Gemini at Work" and Platform-Centric Wins

Instead of vendor-led labs, Google's executive engagement model is built primarily on executive briefings and strategy workshops, exemplified by events like "Gemini at Work."¹⁸ This is a *product-led growth (PLG)* motion scaled to the enterprise. Google's pitch is "Here's what Mercedes-Benz, Virgin Voyages, and Mercari *built on our platform*."¹⁸ The implication: "You can too."

Case Studies: Mercedes-Benz and Virgin Voyages as Platform Wins

The case studies confirm this PLG strategy, and the narrative is notably different from AWS's or Microsoft's. Customer wins are attributed to the *power of the platform and its models*, not to a *consultative engagement*.

- **Mercedes-Benz:** Is “building cars that can converse with their drivers” using “Gemini via Vertex AI.”¹⁸ This is a *product feature* win for Vertex AI.
- **Virgin Voyages:** Is “using Veo’s text-to-video features to create thousands of hyper-personalized ads.”¹⁸ This is an *application* win for Google’s Veo model.

VI. Competing Methodologies: The Broader Innovation Landscape

The hyperscalers’ “platform-subsidized” GTM motion is best understood in contrast with their more traditional competitors.

The Paid Engagement Model: IBM’s “Co-Create, Co-Execute, Co-Operate” Garage

The IBM Garage¹⁹ and its “Co-Create, Co-Execute, Co-Operate” methodology¹⁹ represent the traditional, paid consulting model. IBM is selling consulting as its primary product.²² The hyperscalers are disrupting this model by giving away the initial consulting (the PoC or workshop) because they make their money on back-end platform consumption. IBM cannot easily do the same, because consulting is its revenue stream.

The Vertical Focus: Oracle’s Industry-Specific Co-Creation Labs

Oracle’s “Industry Innovation Lab”²³ is a defensive GTM motion. The labs are highly vertical-specific, with physical locations for construction, utilities, and healthcare.²³ This is an upsell and cross-sell motion designed to prevent customers from defecting. The goal is to show a healthcare customer²⁴ or a Fusion ERP customer²⁵ how Oracle’s AI can improve the Oracle application they already pay for.

The Upskilling Hub: Salesforce’s UK AI Centre and CRM-Centric Model

Salesforce’s UK AI Centre²⁶ is an even more focused defensive play. It is not a PoC lab; it is a training center. Its primary focus is “critical upskilling opportunities”²⁷ within its own ecosystem. Salesforce’s goal is to ensure its 500,000+ ecosystem jobs²⁷ are trained on its AI platform, Einstein.²⁸ This is a GTM motion to drive upsell revenue for the “#1 AI CRM”²⁹ by embedding AI into existing licenses.

VII. Strategic Analysis

What Works: The Power of Rapid Prototyping and Embedded Expertise

- **Microsoft’s Rapid Sprint:** The “Rapid Sprint” model³ works exceptionally well for generating GTM velocity, customer buy-in (via the IP hook¹⁵), and overwhelming market proof (62% case study share¹⁶). It is the most effective *sales conversion* and *marketing* model.
- **AWS’s Embedded Strategist:** The “Value-first”¹⁰ consultative model works for large, complex, risk-averse enterprises. By embedding scientists,¹ AWS de-risks massive, multi-year transformations, building C-suite trust and securing high-LTV, foundational modernization contracts.² **Paired with a**

Working Backwards engagement, AWS combines high-level strategy with technical proof, but this approach requires more time and a greater commitment from customers.

- **The Hybrid Model:** The most effective format blends the two. An “operationalized,” abridged design-thinking session identifies business value in a few focused areas (based on pre-workshop discovery), followed by a three-to-five-day design and co-development sprint that produces one or more working prototypes in the customer’s own cloud environment. The customer owns the IP, and the prototypes serve as a pull-through mechanism not only for accelerated production deployments, but also for follow-on workload migrations, exactly as the Microsoft model is designed to do.

The Gap in a Partner-Only Model

A platform vendor that leaves early-stage prototyping entirely to partners cedes its “**Innovation Wedge**” to competitors willing to invest engineering hours (not just credits) to win the long-term platform war. That creates a structural gap in the GTM motion.

- **The Gap:** Partners (GSIs and RSIs) are incentivized by **billable hours** and **margin**. They naturally gravitate toward known solutions (migrations, data warehouse or ERP implementations) where risk is low and profit is predictable.
- **The Consequence:** It is hard to motivate partners to deliver rapid prototyping because, unlike an internal team, where the vendor and the customer are learning, building and investing together, the technical risk is higher and the immediate margin is low. And when a field team secures funding for a partner-led prototyping workshop, it shrinks the pot of money available for production deployment, with no guarantee the partner will get it right.
- **The Result:** Without an internal team to de-risk these “art of the possible” prototypes, the most innovative use cases, the ones that actually drive future consumption, often die on the whiteboard.

The “Free” Fallacy: Comparative Cost and Revenue Models

“Free” is a misnomer; it means “platform-subsidized customer acquisition.” The table below synthesizes the divergent strategies and their true revenue targets.

Table 2: Comparative Analysis of Hyperscaler and Enterprise Vendor GenAI Innovation Models

Provider	Program	Primary Methodology	Investment Model (Customer Cost)	IP Ownership	Key Differentiator	Primary Revenue Target
AWS	Generative AI Innovation Center	Embedded Strategist: high-touch, value-first, consultative engagement. ¹	“No-cost” for qualified, high-potential customers.	Co-developed.	Deeply embedded in-house strategist & data science teams. ¹	Platform consumption (Bedrock, SageMaker) & partner consulting (PAF) ¹²
Microsoft	AI Co-Innovation Labs	Rapid Engineering Sprint: high-velocity, one-week prototyping sprint. ³	“Complimentary” for qualified customers. ³	Customer retains 100% of IP. ¹⁵	Physical labs for intense, short engineering sprints; IP ownership drives lock-in. ³	Platform consumption (Azure OpenAI, Azure AI) ⁴

Provider	Program	Primary Methodology	Investment Model (Customer Cost)	IP Ownership	Key Differentiator	Primary Revenue Target
Google Cloud	GenAI CoE & executive engagements	Platform-as-a-Lab: product-led, focused on Vertex AI and model choice. ⁶	Often funded via joint partner programs or platform credits. ³¹	Varies by engagement (often partner-owned).	Strong leverage of partner ecosystem; unified Vertex AI platform. ⁵	Platform adoption (Vertex AI, BigQuery) & partner-led consulting ¹⁷
IBM	IBM Garage for Generative AI	Paid Consulting: “Co-Create, Co-Execute, Co-Operate” methodology. ¹⁹	Paid engagement.	Co-developed as part of a paid SOW.	Proven, mature methodology; deep integration with IBM Consulting. ²¹	Paid consulting revenue & IBM platform/software ²²
Oracle	Oracle Industry Innovation Lab	Vertical Co-Creation: industry-specific, physical lab workshops. ²³	Subsidized for strategic customers.	Co-developed.	Deep, “real-world” simulation for specific industries (e.g., construction, healthcare). ²³	Application upsell (Fusion, Oracle Health) & platform defense ²⁵

VIII. Why a Lab Approach Creates a “Consumption Culture”

While hyperscalers sell potential (capacity), they only recognize revenue through utilization. A “consumption culture” is one where the customer habitually solves business problems by spinning up cloud services and embracing low-risk experimentation, rather than waiting for software release cycles. **The following rationale details why rapid prototyping sprints are the most effective tool for fostering this culture.**

1. The “Shelfware” Prevention Mechanism

The Problem: Large enterprises often sign multi-million-dollar commit contracts to get discounts, but then fail to migrate workloads quickly enough to burn down that commitment. This creates “cloud shelfware”: credits that sit unused due to organizational inertia, fear of failure, or technical debt.

The Sprint Rationale (Unblocking Inertia): A three-to-five-day rapid prototyping sprint forces a “bias for action.” It removes the paralysis of months-long planning cycles.

- **Immediate Burn:** By building a functional PoC in a week, the customer immediately moves from “planning” to “provisioning.” Even a small prototype starts the meter running on data storage, compute, and AI services.




2. De-risking High-Value Services (AI & ML)

The Problem: Traditional “lift and shift” migration yields low-margin consumption when done incrementally. The real revenue growth for hyperscalers lies in high-value, “sticky” services like AI, data storage and processing, application modernization and serverless computing. However, customers are often tied to legacy technologies and reluctant to invest in a high-risk, full migration.

The Sprint Rationale (Low Cost of Failure): In a co-innovation lab, a customer can test a complex AI or data hypothesis in a few days with hyperscaler engineers. If it fails, they lose one week, not one year, and they learn how to pivot and try again.

- **Confidence to Scale:** When a prototype succeeds, the customer has a verified proof of concept. This gives them the confidence to turn the “prototype” into “production,” triggering the large-scale growth (and consumption) that hyperscalers want.



3. Transforming “Gatekeepers” into “Builders”

The Problem: Traditional IT cultures are “gatekeepers.” Their job is to protect the existing stack and on-premises stability, which naturally resists the constant change that meaningful cloud consumption requires.

The Sprint Rationale (Cultural Change): Rapid prototyping introduces a “builder” mindset. When internal teams see they can build a working app in days rather than months, the culture shifts from gatekeeping resources to consuming resources to solve problems.

- **Upskilling by Osmosis:** Working side by side with hyperscaler engineers (e.g., in Microsoft’s AI Co-Innovation Labs or AWS’s Working Backwards co-builds) upskills the customer’s team. A more skilled team can consume more complex (and expensive) cloud services on its own in the future.

4. Creating “Pull-Through” Revenue

The Problem: Selling a single service (e.g., storage) is hard. It is easier to sell a solution (e.g., predictive maintenance).

The Sprint Rationale (The Solution Wedge): A sprint focuses on a business outcome (e.g., “reduce factory downtime”). To solve it during the prototype, the solution naturally pulls in multiple consumption levers: IoT services to ingest data, plus a data lake to store it, plus SageMaker or Vertex AI to analyze it.

- **Ecosystem Entrenchment:** Once a prototype is built using native cloud tools, refactoring it to run elsewhere is difficult. This locks in consumption for the lifecycle of that application.

5. Shortening the “Time-to-Consumption”

The Problem: In traditional migrations or waterfall development, a cloud environment might not be fully provisioned and used until month nine of a project.

The Sprint Rationale (Day 1 Usage): Sprints require environments to spin up immediately, typically on Day 1 of the lab.

- **Viral Adoption:** A successful prototype serves as an internal marketing tool for the customer. Business units show off the new app to other divisions, sparking demand for similar sprints and further consumption across the enterprise. In private equity portfolios, a single visible win can generate interest from sister companies across the portfolio.

Summary: Traditional vs. Consumption Culture

Feature	Traditional IT Culture	Consumption Culture (Fostered by Sprints)
Metric	Uptime / stability	Velocity / feature release
Risk Model	Avoid failure at all costs	Fail fast, learn cheap
Resource Model	Fixed capacity (CapEx)	Elastic scaling (OpEx)
Development	Waterfall (months/years)	Agile sprints (days/weeks)
Impact on Platform Revenue	Stagnant revenue	Compounding consumption growth

IX. Implications for Platform Vendors

The analysis points to a clear model for any platform vendor competing for AI workloads: **a core specialist co-innovation and rapid prototyping team, supported by a lean regional hub-and-spoke model and a small set of selected, embedded, certified partners.**

The objective of such a team is to use hands-on co-innovation to convert C-suite interest into long-term, high-margin platform consumption. **Success is not defined by the one-week proof of concept; it is defined by the multi-year production workloads that follow.** The team should be measured accordingly: on deal velocity, net-new production workloads, and win rates for opportunities that include a prototyping sprint, compared with those that don't.

Why a Dedicated Team: The “Unicorn Fallacy”

It is tempting to assume that existing pre-sales engineers can absorb rapid prototyping alongside their account work. Three constraints make that unlikely.

1. The Bandwidth Constraint (the “Context Switch” Tax)

- **The Myth:** “Sales engineers have downtime; they can code between meetings.”
- **The Reality:** Deep coding requires a flow state. An engineer managing five to ten accounts is constantly interrupted by email, business reviews, support tickets, and urgent customer needs.
- **The Result:** A generalist might take weeks to write code that a dedicated prototyper could write in a few days. Time-to-value drags out, and the customer loses momentum.

2. The Skillset Gap (Demos vs. Working Prototypes)

- **The Myth:** “Our sales engineers are technical; they can build apps.”
- **The Reality:** Most pre-sales engineers excel at **architecture** and **configuration**. Fewer are full-stack developers able to write the Python, React or Go code that stitches those components into a custom app.
- **The Result:** Generalists often build demos with standardized functionality, which is fine for opening doors but may not reflect the customer's specific business needs. Dedicated prototypers build PoCs that solve the customer's actual use cases.

3. The Incentive Mismatch (the Deal vs. the Build)

- **The Myth:** “Sales engineers are motivated to close the deal, so they will build the best prototype.”
- **The Reality:** A sales engineer is motivated to secure the *technical win*. Once the customer says yes, they hand off to the account team for commercials and move to the next opportunity to hit their number.
- **The Result:** Without a dedicated overlay team and a programmatic path from PoC to production (via partners or professional services), the prototype can be abandoned, leaving the customer to restart from scratch. A dedicated team would be measured on **completion** and **deployment**, not just commitments.

Works Cited

1. Generative AI Innovation Center - AWS - Amazon AWS, accessed November 6, 2025, <https://aws.amazon.com/ai/generative-ai/innovation-center/>
2. AWS doubles investment in AWS Generative AI Innovation Center, marking two years of customer success, accessed November 6, 2025, <https://aws.amazon.com/blogs/machine-learning/aws-doubles-investment-in-aws-generative-ai-innovation-center-marking-two-years-of-customer-success/>
3. Microsoft AI Co-Innovation Labs, accessed November 6, 2025, <https://aiotlabs.microsoft.com/en>
4. Renaissance , [Case Study](#)
5. Google Cloud: Partner Home, accessed November 6, 2025, <https://partners.cloud.google.com/>
6. Google Cloud CEO: Why Customers Are Betting Their GenAI Future on Google's AI Technology, CRN, accessed November 6, 2025, <https://www.crn.com/news/ai/2024/google-cloud-ceo-why-customers-are-betting-their-genai-future-on-google-s-ai-technology>
7. Predictions 2025: GenAI As A Growth Driver Will Put B2B Executives..., Forrester, accessed November 6, 2025, <https://www.forrester.com/blogs/predictions-2025-b2b-marketing-sales-product/>
8. Generative AI Trends For All Facets of Business - Forrester, accessed November 6, 2025, <https://www.forrester.com/technology/generative-ai/>
9. The Future Of AI Consulting Services Is Disruptively Bright - Forrester, accessed November 6, 2025, <https://www.forrester.com/blogs/the-future-of-ai-consulting-services-is-disruptively-bright/>
10. Generative AI , [Artificial Intelligence - Amazon AWS, accessed November 6, 2025](#)
11. Predictions 2024: AI Will Boost Big Cloud Players — But Threaten Them, Too - Forrester, accessed November 6, 2025, <https://www.forrester.com/blogs/predictions-2024-cloud/>
12. AWS Generative AI Innovation Center Launches Partner Agent Factory, accessed November 6, 2025, <https://aws.amazon.com/blogs/apn/aws-generative-ai-innovation-center-launches-partner-agent-factory/>
13. Introducing the AWS Generative AI Innovation Center's Custom Model Program for Anthropic Claude , [Artificial Intelligence, accessed November 6, 2025](#)
14. Customer Success Stories: Case Studies, Videos, Podcasts, Innovator Stories - Amazon AWS, accessed November 6, 2025, <https://aws.amazon.com/solutions/case-studies/>
15. UWM and Microsoft unveil AI Co-Innovation Lab to drive manufacturing innovation, accessed November 6, 2025, <https://uwm.edu/news/uwm-and-microsoft-unveil-ai-co-innovation-lab-to-drive-manufacturing-innovation/>
16. Who is winning the cloud AI race? Microsoft vs. AWS vs. Google, IoT Analytics, accessed November 6, 2025, <https://iot-analytics.com/who-is-winning-the-cloud-ai-race/>
17. Google Cloud partner digs into generative AI for marketers - Marketing Tech News, accessed November 6, 2025, <https://www.marketingtechnews.net/news/google-cloud-partner-digs-into-generative-ai-for-marketers/>
18. Real-world gen AI use cases from the world's leading organizations, accessed November 6, 2025, <https://cloud.google.com/transform/101-real-world-generative-ai-use-cases-from-industry-leaders>
19. IBM Garage, accessed November 6, 2025, <https://www.ibm.com/garage>
20. IBM Consulting unveils Center of Excellence for generative AI, accessed November 6, 2025, <https://www.ibm.com/new/announcements/ibm-consulting-unveils-center-of-excellence-for-generative-ai>
21. IBM Garage — Phase 1: A Startup is Announced , [by Rachel Reinitz](#)
22. IBM to drive enterprise-ready AI adoption at scale in collaboration with NVIDIA, accessed November 6, 2025, <https://newsroom.ibm.com/IBM-to-drive-enterprise-ready-AI-adoption-at-scale-in-collaboration-with-NVIDIA>
23. Industry Innovation Lab , [Oracle, accessed November 6, 2025](#)
24. Oracle Launches an AI Center of Excellence for Healthcare to Help Customers Maximize the Value of AI Across Clinical, Operational, and Financial Workflows, accessed November 6, 2025, <https://www.oracle.com/news/announcement/oracle-launches-an-ai-center-of-excellence-for-healthcare-2025-09-10/>

25. Oracle AI for Fusion Applications, accessed November 6, 2025, <https://www.oracle.com/applications/fusion-ai/>
26. Salesforce Opens First AI Center with AI Training Event Backed by London City Hall, accessed November 6, 2025, <https://www.salesforce.com/uk/news/stories/ai-center-london-opening/>
27. Salesforce Picks London for First AI Center, Building on \$4B Investment, accessed November 6, 2025, <https://www.salesforce.com/news/press-releases/2024/06/06/ai-center-london/>
28. Artificial Intelligence - Powerful AI Solutions - Salesforce, accessed November 6, 2025, <https://www.salesforce.com/artificial-intelligence/>
29. Salesforce: The #1 AI CRM , [Salesforce](#), accessed November 6, 2025
30. 9 Google Cloud customers on AI implementations, best practices - Constellation Research, accessed November 6, 2025, <https://www.constellationr.com/blog-news/insights/9-google-cloud-customers-ai-implementations-best-practices>
31. Google Developer Program Plans & Pricing, accessed November 6, 2025, <https://developers.google.com/program/plans-and-pricing>