What Silicon Valley builds today – enterprises face tomorrow

February 2026

PRESENTED BY:

David Cushman, Executive Research Leader

Contents

HFS © 2026 HFS Research. All Rights Reserved.

Tech holds the key to your next leap in growth

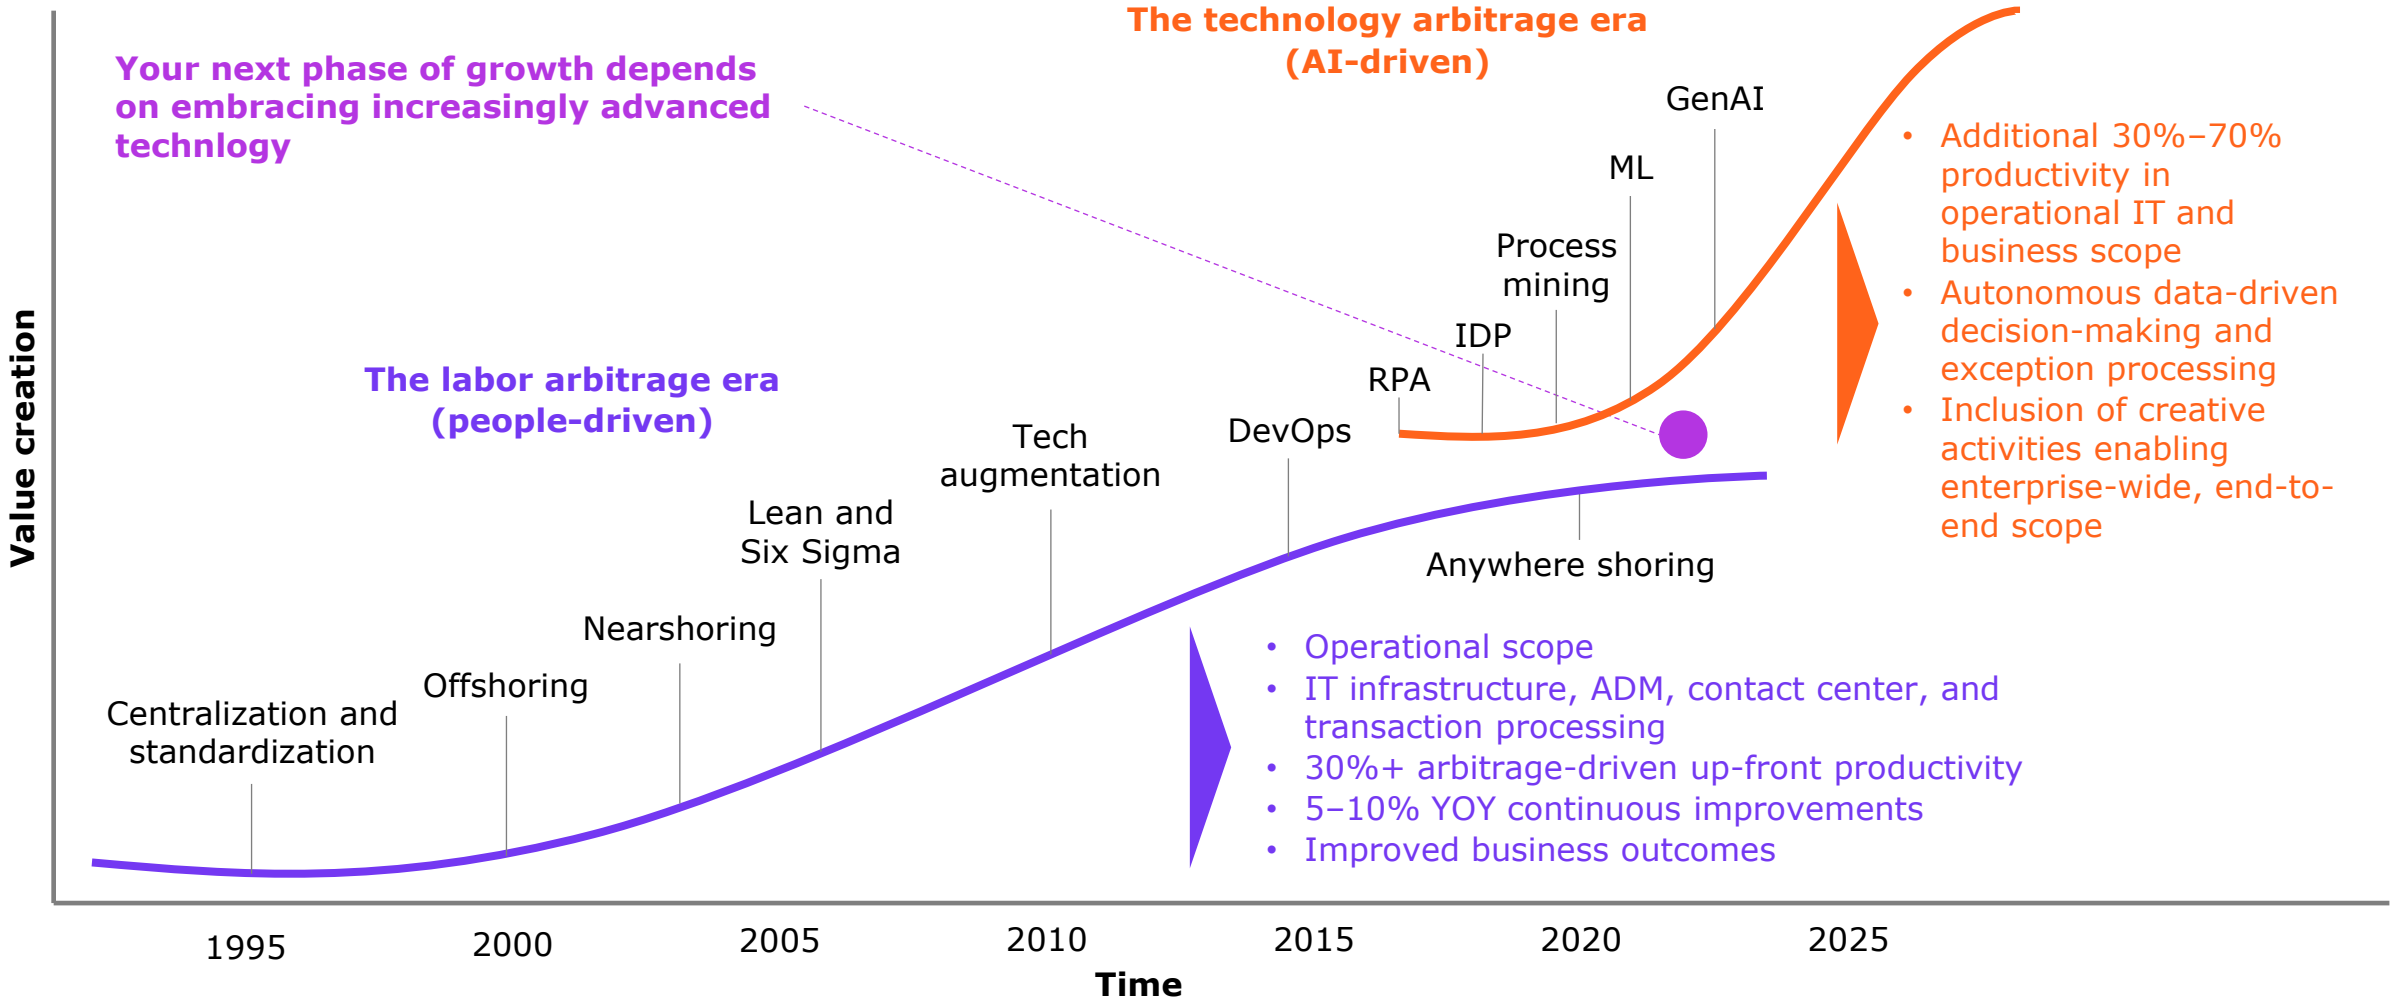

HFS © 2026 HFS Research. All Rights Reserved.

How the leap happens: HFS Services and Ops Tech Vision 2028

Human | Machine

Staff augmentation

- Enables companies to quickly fill skill gaps, scale teams up or down as needed, and maintain control over project execution without the long-term commitment associated with permanent hires.

Key features:

- Flexibility: Easily adjusts team size based on project needs.
- Expertise: Access to specialized skills not available in-house.
- Control: Maintains direct oversight of projects and processes.

Typical commercial model:
Rate card

Technology-enabled services

- Primarily driven by people but supported by proprietary solution accelerators, tools, and software.
- Most service providers use this model to optimize processes and deliver value efficiently; examples include Cognizant Neuro, Infosys Topaz, TCS WisdomNext, and Wipro Lab45.

Key features:

- Human-centric: Primarily driven by skilled professionals.
- Tool-supported: Utilizes a variety of technology tools and accelerators.
- Efficient: Enhances service delivery through technology integration.

Typical commercial model:
FTE-based pricing

Platform-led services

- Leverage built-in delivery platforms to enhance service delivery and efficiency.
- Examples include Accenture SynOps, TCS Cognix, and Cognizant TriZetto—which streamline operations and provide consistent, scalable solutions.

Key features:

- Integrated platforms: Uses cohesive platforms for service delivery
- Scalability: Easily scalable and consistent across various operations
- Efficiency: Enhances productivity and efficiency through platform support

Typical commercial model:
Transaction-based pricing

AI-led agentic services

- Augment human capabilities through smart AI agents to optimize processes and decision-making.
- Examples include Amazon Q, GitHub, Lyzr, Copilot, Replit Ghostwriter, Google Gemini, Einstein Agent, and Mindcorp.
- Organizations such as IBM and the Big 4 consulting firms are increasingly adopting this model.

Key features:

- AI-augmented: Combines human expertise with AI agents
- Cost-effectiveness: Achieves lower TCO through optimization
- Enhanced capabilities: Expands service potential with AI-driven insights

Typical commercial model:
Augmented FTE-based pricing or outcome-driven performance pricing

Services-as-software

- Unlike traditional software-as-a-service (SaaS), this model focuses on delivering services primarily through technology, minimizing human intervention, and maximizing efficiency.
- Examples include startups such as rhino.ai, Now Platform, and Builder.ai.

Key features:

- Technology-driven: Primarily led by advanced software solutions
- Minimal human intervention: Reduces reliance on human resources
- Efficient and scalable: Provides efficient, scalable, and consistent service delivery

Typical commercial model:
License/subscription-based pricing

Current state
2000–2025

Emerging
2025–2028

Source: HFS Research, 2026

© 2026 HFS Research. All Rights Reserved.

Revenue / employee as proxy for impact of Services as Software

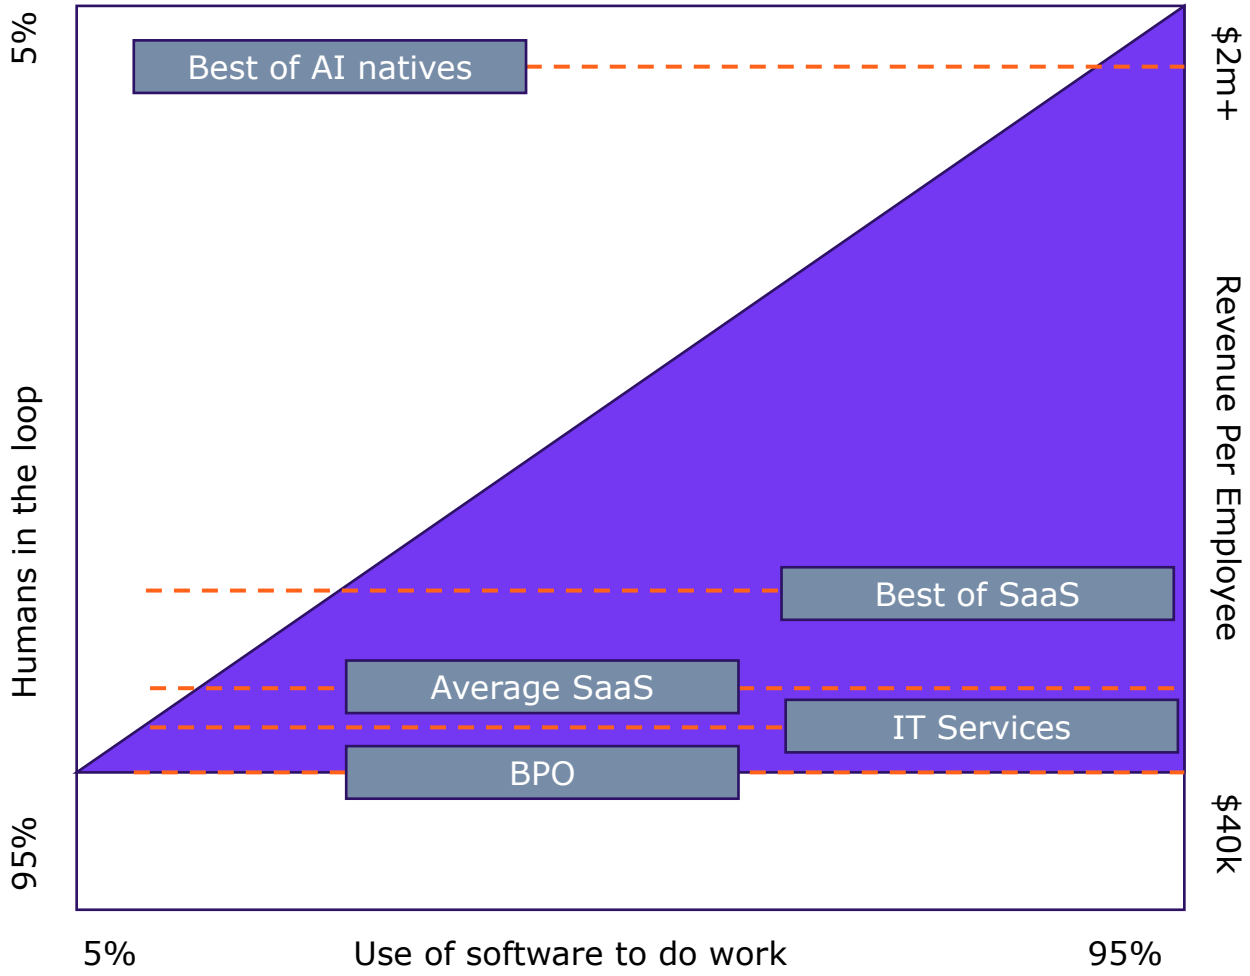

HFS © 2026 HFS Research. All Rights Reserved.

2 Why The Valley Matters

© 2026 HFS Research. All Rights Reserved.

Tech is the growth lever, AI its beating heart, The Valley its home

67% of all VC AI investment is in the Bay Area: ~ $135 billion in 2024 vs- ~$65 billion RoTW

40% of the world's unicorns have their home in the bay area (that's around 600 firms)

4000+ The number of AI events in the valley in the last three years – and the volume is rising

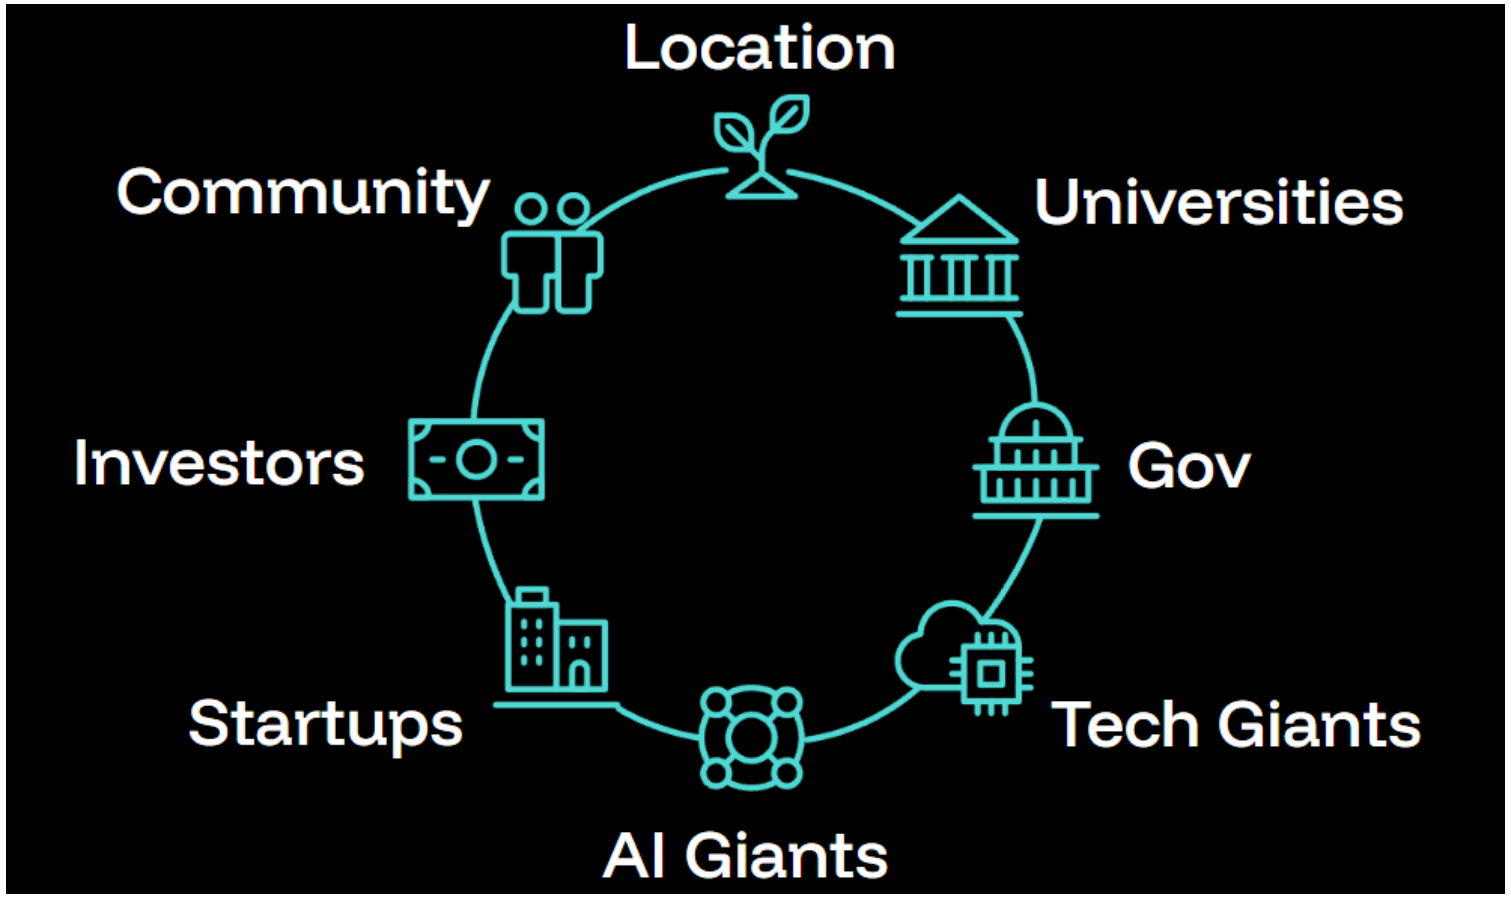

Source: Jeremiah Owyang

HFS © 2026 HFS Research. All Rights Reserved.

Stay close to what's happening – lagging is getting more costly

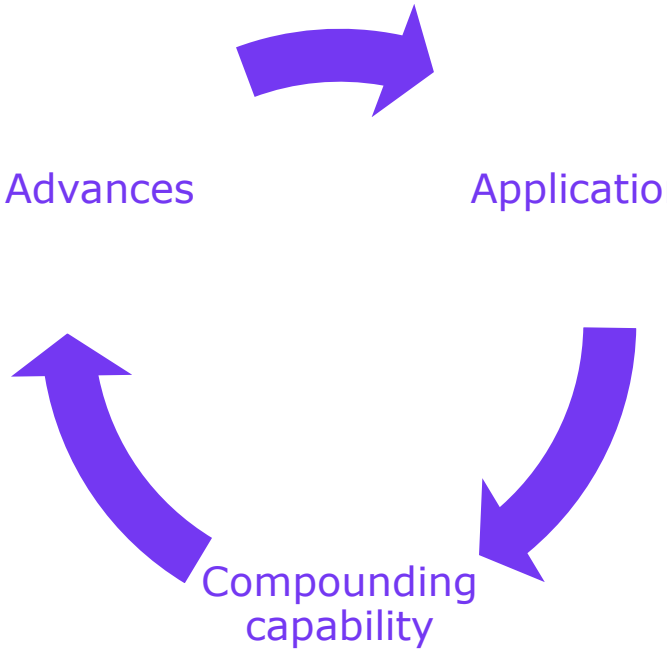

280x The fall in cost of querying a top-tier AI - in 18 months.

1 million step agents for real work

Open source models are catching up.

'Best-in-breed' is rarely at the top for long.

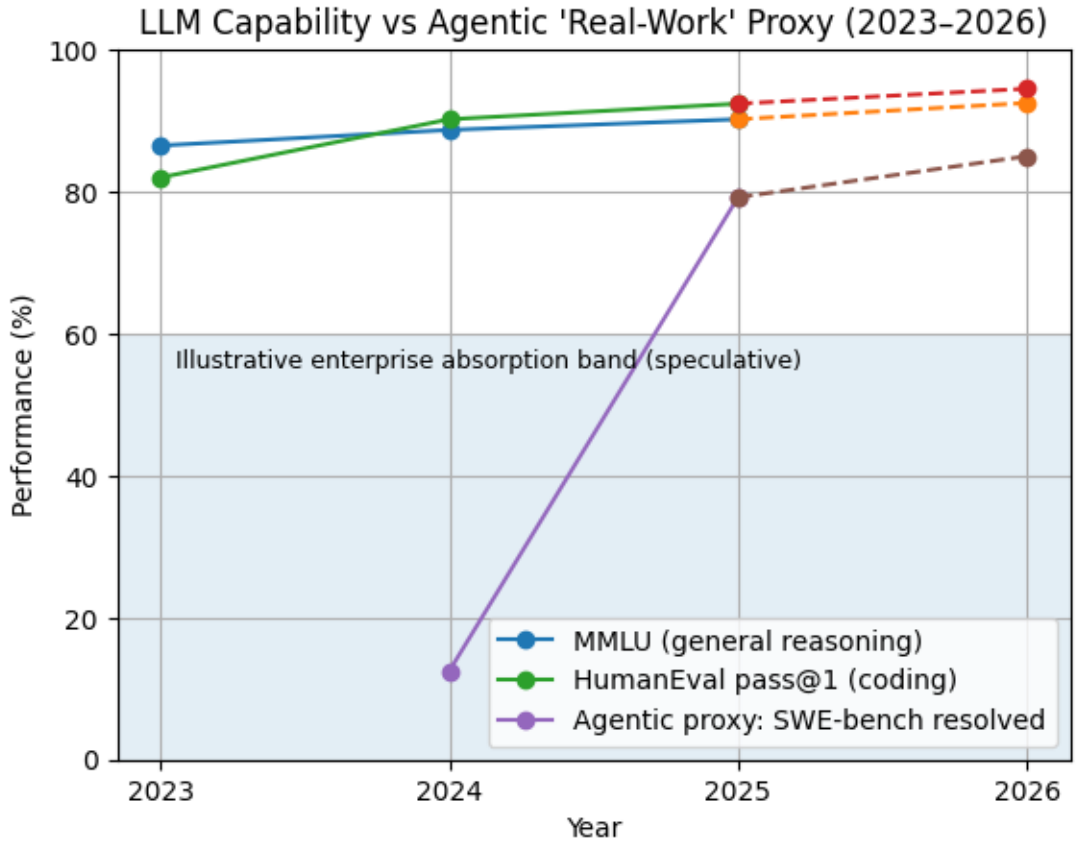

Most enterprises' *operational change capacity – the absorption band* - isn't keeping up with rising capability

HFS © 2026 HFS Research. All Rights Reserved.

What this means for the enterprise

The pace of improvement emphasises the need to avoid long-term vendor lock-in while selecting safe, productive LLMs.

Contracting is changing – HFS's ***AI deal lab*** is for firms seeking to balance flexibility with predictability.

Both tech and business architectures must become more modular and design for low-friction when swopping elements in and out.

This is the age of impermanence.

HFS © 2026 HFS Research. All Rights Reserved.

3

Multi Agent Systems

© 2026 HFS Research. All Rights Reserved.

The valley put agents on the team - you need to catch up

The Valley adopted agents as co-working colleagues instantly. It's allowed start-ups to scale faster.

To make the journey from agentic 'primitive', to enterprise-grade, required new solutions such as MCP (Model-Context-Protocol).

1 Week The time it took for MCP to spread across The Valley.

1 Year The time agents have gone from primitive to Multi-Agent-Systems

Few enterprise technologies have ramped this quickly.

HIGHLIGHT

Ops leaders must embrace DoubleO's agentic-by-default approach

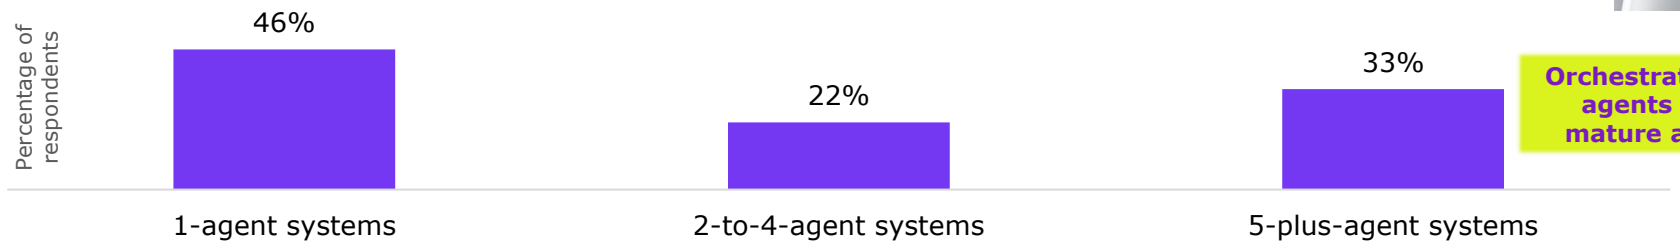

Orchestrators are already averaging 12 agents in production in their most mature and successful deployments.

© 2026 HFS Research. All Rights Reserved.

What this means for the enterprise

Primitives are being developed to enterprise-grade at increasing pace – The Valley is paying attention to enterprise (not just consumer) need.

Workflows and org charts must be reviewed.

Don’t fall into the productivity trap.

Get set to build an agent OS – not just more agents.

HFS © 2026 HFS Research. All Rights Reserved.

4

Vibe coding

© 2026 HFS Research. All Rights Reserved.

Vibe coding is rampant in the valley – dormant in most enterprises

Vibe coding is the production engine to make Services-as-Software real. It will deliver great productivity uplifts if you develop AI-native talent and redesign leadership behaviours to enable decentralised innovation.

Firms such as Cursor (San Francisco) and Replit (Foster City) lead.

10x – Replit's growth since the launch of its vibe coding agent (~$2m rev/employee)

$6m – Cursor's revenue/employee – with ARR doubling every two months

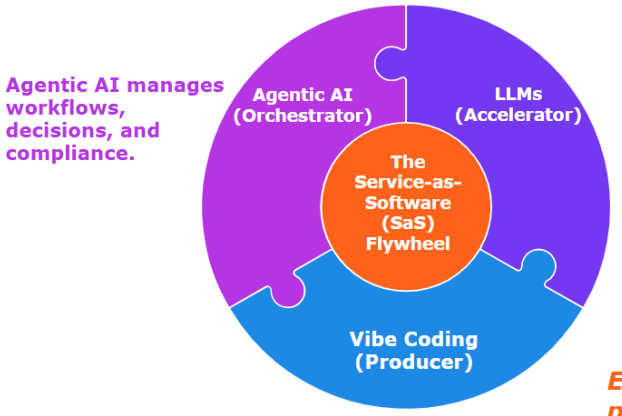

Embrace "vibe coding" as the production engine that makes Services-as-Software real…

Q1: To what extent do you agree that your organization is culturally and organizationally ready to realize productivity gains from vibe coding?

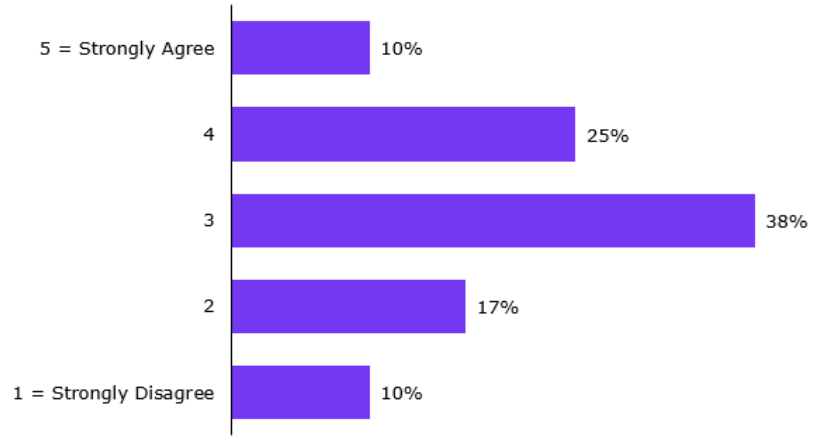

Sample: 100 UK&I business leaders
Source: HFS Research, 2026

HFS © 2026 HFS Research. All Rights Reserved.

What this means for the enterprise

When global players such as Cognizant partner with the likes of Lovable, Cursor, and CrewAI to support internal (recording breaking) coding initiatives, it's time to take vibe coding seriously.

1. **Most firms are not culturally and organisationally ready for Vibe Coding.** 27% actively don't feel ready, and the greatest number (38%) are still unsure.
2. **Digital-savvy grads and early career technologists hold the key.** 88% of respondents recognize AI-native talent is critical to moving the needle on cultural readiness.
3. **Cultural drag hampers adoption.** No more than 22% of workforces will be using vibe coding within six months – slowed by risk aversion and lack of talent.
4. **Leaders need to set clear rules.** The biggest shackles on progress with vibe coding are blockers from legal, security and compliance functions.
5. **Productivity first.** Acceleration is good – but transformation is better. Firms must embrace the latter to create real value with vibe coding.
6. **Cursor is king:** Despite the rise of Lovable and Replit, it's Cursor that is the current defacto choice among enterprise vibe coding tools.

- We are at the edge of a new culture in which impermanence takes precedence over software as structure: Software for the moment and for the need, rather than built for the long run.
- Does Vibe Coding need to be perfect? future proof? Endlessly maintainable? Rising stars such as **Prompt Driven** reject such notions - arguing code is so cheap and fast to build there is no point trying to maintain or fix creaking code bases.

HFS © 2026 HFS Research. All Rights Reserved.

Agents + vibe coding are transforming how the valley works

Thousands of first-time entrepreneurs are leaving great jobs and dropping out of fantastic university courses because it's now so easy, cheap, and fast to make their business dreams come true.

YC teams have dropped from (on average) six people in 2022 to just two or three in 2024, with **a surge in solo founders**.

Agents and vibe coding usher in an era in which we are all potential founders.

You must harness this new reality within your own organisation.

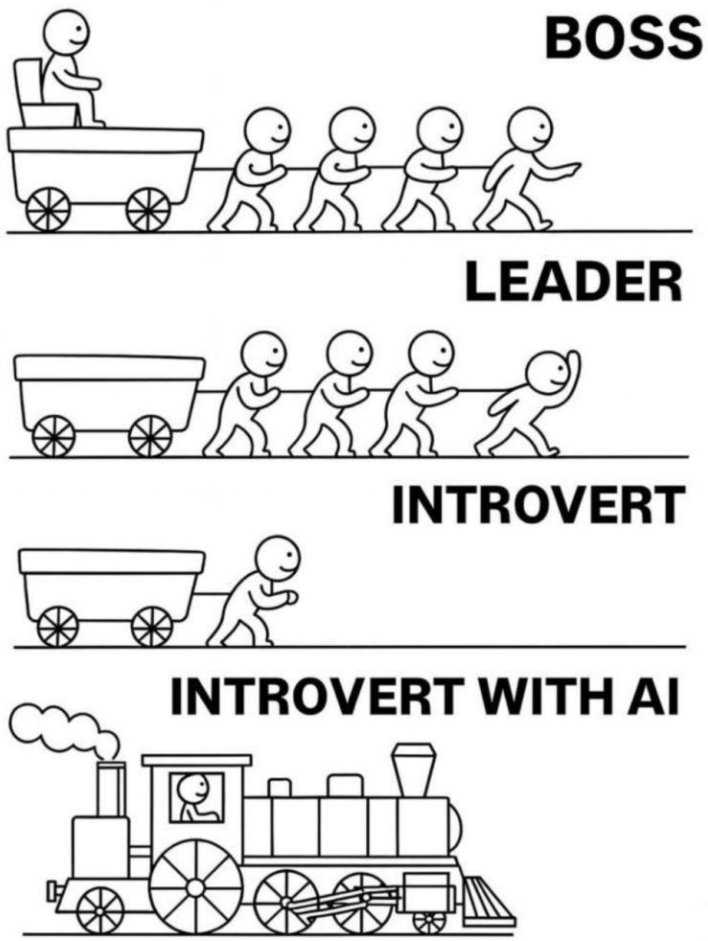

Credit: Alina Krasnobrizha PhD

What this means for the enterprise:

It's now easier to make business dreams real outside your org than within it.

Reverse that by providing the tools to enable everyone to create, direct, and deploy, their own agentic workforce.

Enable your people to be the human at the helm of the AI operation of end-to-end processes.

HFS © 2026 HFS Research. All Rights Reserved.

5 Vertical AI

© 2026 HFS Research. All Rights Reserved.

Vertical AI places domain expertise front and center

Deep expertise, rich relationships, and real context are becoming a recipe for out-performing even the most heavily-funded frontier LLMs as AI becomes more grounded in the real work of industries.

Vertical AI is designed to understand the messy, domain-specific workflows inside (for example) hospitals, networks, or even PE portfolio operations.

Proponents (for example Veltris) address vertical specific challenges as orchestrators of vertical AI solutions – embedding industry-specific workflows with a focus on maximising ROI.

What this means for the enterprise:

As the application of AI becomes more pragmatic an effective route to monetization is to develop and operationalize verticalized models with industry-experienced partners.

Build: Construct the engagement to build or innovate a new product or platform, or adapt an existing one to a new commercial model.

Modernize: Evolve to AI-first architectures, integrate with existing systems, harden security, unify data platform for AI readiness, and automate pipelines to modernize the surrounding systems.

Monetize: Accrue ROI from KPI-bearing features that will be shipped to end users and deliver insights, for new revenue streams. Enable non-linear scale by making the platform available with APIs, etc.

HFS © 2026 HFS Research. All Rights Reserved.

6 Efficiency breakthroughs

© 2026 HFS Research. All Rights Reserved.

The race is on for better intelligence/watt/dollar

2025

2026

DeepSeek

DeepSeek's R1 model pushed the valley to make efficiency the rising star of 2025.

Us too

California could focus beyond brute-force IQ AI, too - with examples such as **WRITER** AI, and responses from the frontier models).

RL returns

Reinforcement Learning is back in vogue, being applied to support reliable agentic systems

mCH

Now DeepSeek's **Manifold-Constrained Hyper-Connections** paper ups the ante. mHCs combine architecture and systems to increase the effective capacity of a model without ramping up computer processing power. And that means more intelligence/watt/dollar - and less need for high-end GPUs.

What this means for the enterprise:

- **More choice** among an increasing diversity of vertical specialists training mid-size foundational models
- **Downward pricing pressure** on access to 'good enough' capability, and sovereign options becoming more plausible.
- **Emphasised** need to shift from 'which single model' to a portfolio approach.

Every enterprise CTO should track efficiency like they do major model releases.

Ideas like mCH - and those that will follow - will make significant impacts on cost curves.

HFS © 2026 HFS Research. All Rights Reserved.

7 Physical AI

© 2026 HFS Research. All Rights Reserved.

Physical AI is on the rise

Physical AI is where intelligence enters the real world – requiring a step-change from the ability of LLMs to predict the next word, to AI capable of predicting the next atom.

To make the leap to hardware we need real-world robot data - video, actions, sensors, failures. This is where the Valley leads. China makes the hardware, the Valley makes the intelligence.

Researchers, startups, and large players are now releasing real-robot datasets including NVIDIA's LeRobot initiative, and a rapidly growing maker community.

This surge is also enabled by cheaper video storage, better tooling, and an open-source AI culture now spilling into the physical world.

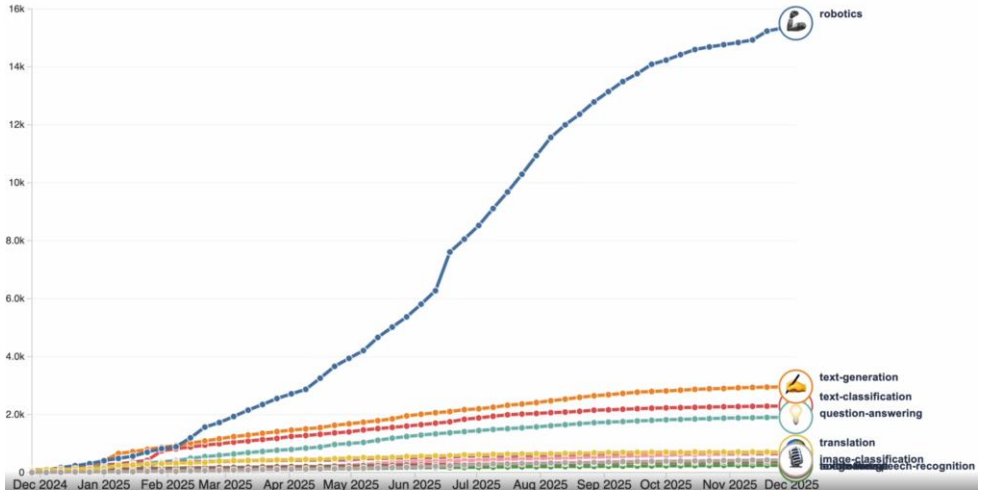

Robotics is the fastest-growing segment among Hugging Face datasets – rising from 1k to 27K in 2025.

HFS © 2026 HFS Research. All Rights Reserved.

What this means for the enterprise

Assisted physical AI

- Programmed to perform specific tasks
- Humans are on standby for exceptions
- AI is used for offline programming and simulation
- Good for pilots, training, and safety baselining
- Implemented by several automobile companies in factories

Conditional physical AI

- Sensor-based autonomy with AI at the edge
- Humans for exceptions
- Edge AI for local decision-making
- Used for real-world PoCs and early production
- BMW: Uses it for faulty part detection using edge computer vision

Autonomous physical AI

- Fusion of multimodal sensors
- Zero-shot learning with self-learning, near real-time decisions
- Used for proven industrial scaling
- Hyundai: Deployed a quadruped robot for quality and safety inspection in its factory

HFS defines physical AI as software-led services embodied in physical systems (robots, humanoids, vehicles, industrial equipment) that perceive, decide, and act in the real world, operating autonomously, and delivered and monetized as productized services.

Bots in the home: eg NEO, Eggie etc are here but don't even qualify as 'assisted physical AI'. They are essentially human-operated. As are all the humanoid robots I've seen to date. These will accelerate the growth of datasets we need to achieve physical AI

HFS © 2026 HFS Research. All Rights Reserved.

8

AGI – definitions and progress

© 2026 HFS Research. All Rights Reserved.

What we mean by AGI

RPA

"I follow instructions exactly"

RPA is the **task automation** that eliminates manual effort wasted on repetitive tasks

Key Characteristics:

- Executes structured, rule-based processes
- Performs repetitive digital tasks
- Operates within defined system boundaries
- Follows exact step-by-step procedures

GenAI

"I can create based on prompts"

GenAI is a **productivity amplifier** that accelerates creative and analytical work that bottlenecks humans

Key Characteristics:

- Assists with specific tasks (writing, analysis, coding)
- Requires human direction and oversight
- Improves individual productivity
- Works within existing job roles

Agentic AI

"I can understand goals and figure out how to achieve them"

Agentic AI is a **collaborative actor** that removes the need for constant human oversight of complex processes

Key Characteristics:

- Acts as virtual coworker completing end-to-end processes
- Self-directs and coordinates multiple tasks
- Transforms entire workflows
- Creates new organizational paradigms

Artificial General intelligence (AGI)

"I can think, reason, and learn like a human in any domain"

AGI is a **self-directed intelligence** that overcomes human cognitive limitations across all domains

Key Characteristics:

- General problem-solving ability across domains
- Autonomous learning and adaptation
- Human-level reasoning and understanding
- Transfer learning between different types of tasks
- Self-improvement capabilities

Artificial Super Intelligence (ASI)

"I can outperform human intelligence."

ASI is a fully autonomous intelligence that **surpasses human cognitive capabilities** and can solve problems and innovate on an exponential scale.

Key Characteristics:

- Surpasses human problem-solving and reasoning
- Forms independent goals and innovates autonomously
- Continuously self-improves and evolves
- Applies superior intelligence across all domains

HFS © 2026 HFS Research. All Rights Reserved.

What the journey to AGI translates as for the enterprise

In enterprise application terms, "closer to AGI" stops meaning "sounds smart" and starts meaning:

- **Higher straight-through completion** on real workflows (less human patching)
- **Lower oversight cost** (review time falls, not just task time)
- **Better recovery** (handles tool failures, ambiguity, missing data)
- **Safer autonomy** (does less dumb stuff *while acting*, not just while answering)

✅ = solidly there (prod-grade in at least some enterprise settings)
🟡 = emerging (works in constrained/assisted setups; not yet reliable at scale)
⚪ = not credibly there yet

'Levels of AGI'	What it means in enterprise terms	OpenAI (GPT-5.2)	Anthropic (Claude)	Google (Gemini/ DeepMind)	Microsoft (Copilot/Agents)	xAI (Grok)	Mistral	Meta (Llama ecosystem)
1. Chatbots	Natural-language interface for knowledge work; drafting, Q&A, summarization	✅	✅	✅	✅	✅	✅	✅
2. Reasoners	Multi-step problem-solving + long-context synthesis with fewer silent failures	✅ (explicitly marketed as long-context reasoning SOTA)	✅/🟡 (strong reasoning; less "official" eval framing)	✅/🟡 (strong, but mixed signals by domain)	🟡 (depends on which underlying model you select)	🟡 (tool-calling claims; uneven trust signals)	🟡 (good models; smaller ecosystem)	🟡 (varies by wrapper; open models)
3. Agents	Can plan + use tools/APIs + execute multi-step tasks with guardrails	✅/🟡 (tool-rich; agentic "projects" feasible; still needs controls)	✅/🟡 (Cowork + advanced tool use; strong real-world task execution)	✅/🟡 (Gemini Agent + "computer use" toolset; Astra still limited testing	✅ (enterprise-grade orchestration + governance story; model-agnostic options)	✅/🟡 (Agent Tools API + tool-calling focus; governance/safety volatility)	✅/🟡 (Agents API exists; ecosystem smaller but coherent)	🟡 (agentic systems exist; maturity depends on integrators)
4. Innovators	Generates novel solutions/designs/research directions that hold up under constraints	🟡 (good co-inventor; not dependable "R&D replacement")	🟡 (strong creative + code; still spiky)	🟡 (DeepMind research edge; productized innovation is uneven)	🟡 (innovation via workflow + data + multi-model choice more than a single model)	🟡 (fast iteration; credibility issues affect adoption)	🟡 (good for constrained innovation workflows)	🟡 (open ecosystem enables innovation, but "system reliability" varies)
5. Organizations (Business Unit as Software)	End-to-end execution of a "business unit" (strategy → ops → governance) with minimal human management	⚪	⚪ Open Claw?	⚪	⚪	⚪	⚪	⚪

HFS © 2026 HFS Research. All Rights Reserved.

9

Turn up the pace of growth by changing how you govern

© 2026 HFS Research. All Rights Reserved.

Update your governance to unleash value

Most enterprises fail to generate top-line growth with AI.

The default is to repeat known processes, a little faster and with a few less humans involved.

As identified at a recent HFS Roundtable – enterprises are deploying AI widely for automation and personalization, yet it contributes less than 1% to top-line growth.

Locked-down and slow-moving governance – focused on governing automation - acts as a brake on growth. Release that brake by applying our RUNWAY framework to govern for innovation.

POINT OF VIEW

HFS

CIOs must adopt the AI–ready governance RUNWAY to make growth fly

LEARN MORE →

https://www.hfsresearch.com/research/cios-adopt-ai-runway-growth-fly/

© 2026 HFS Research. All Rights Reserved.

RUNWAY governance offers unbounded ideation, with bounded execution

Govern actions, not "thought"; preserve exploration; constrain irreversible impact

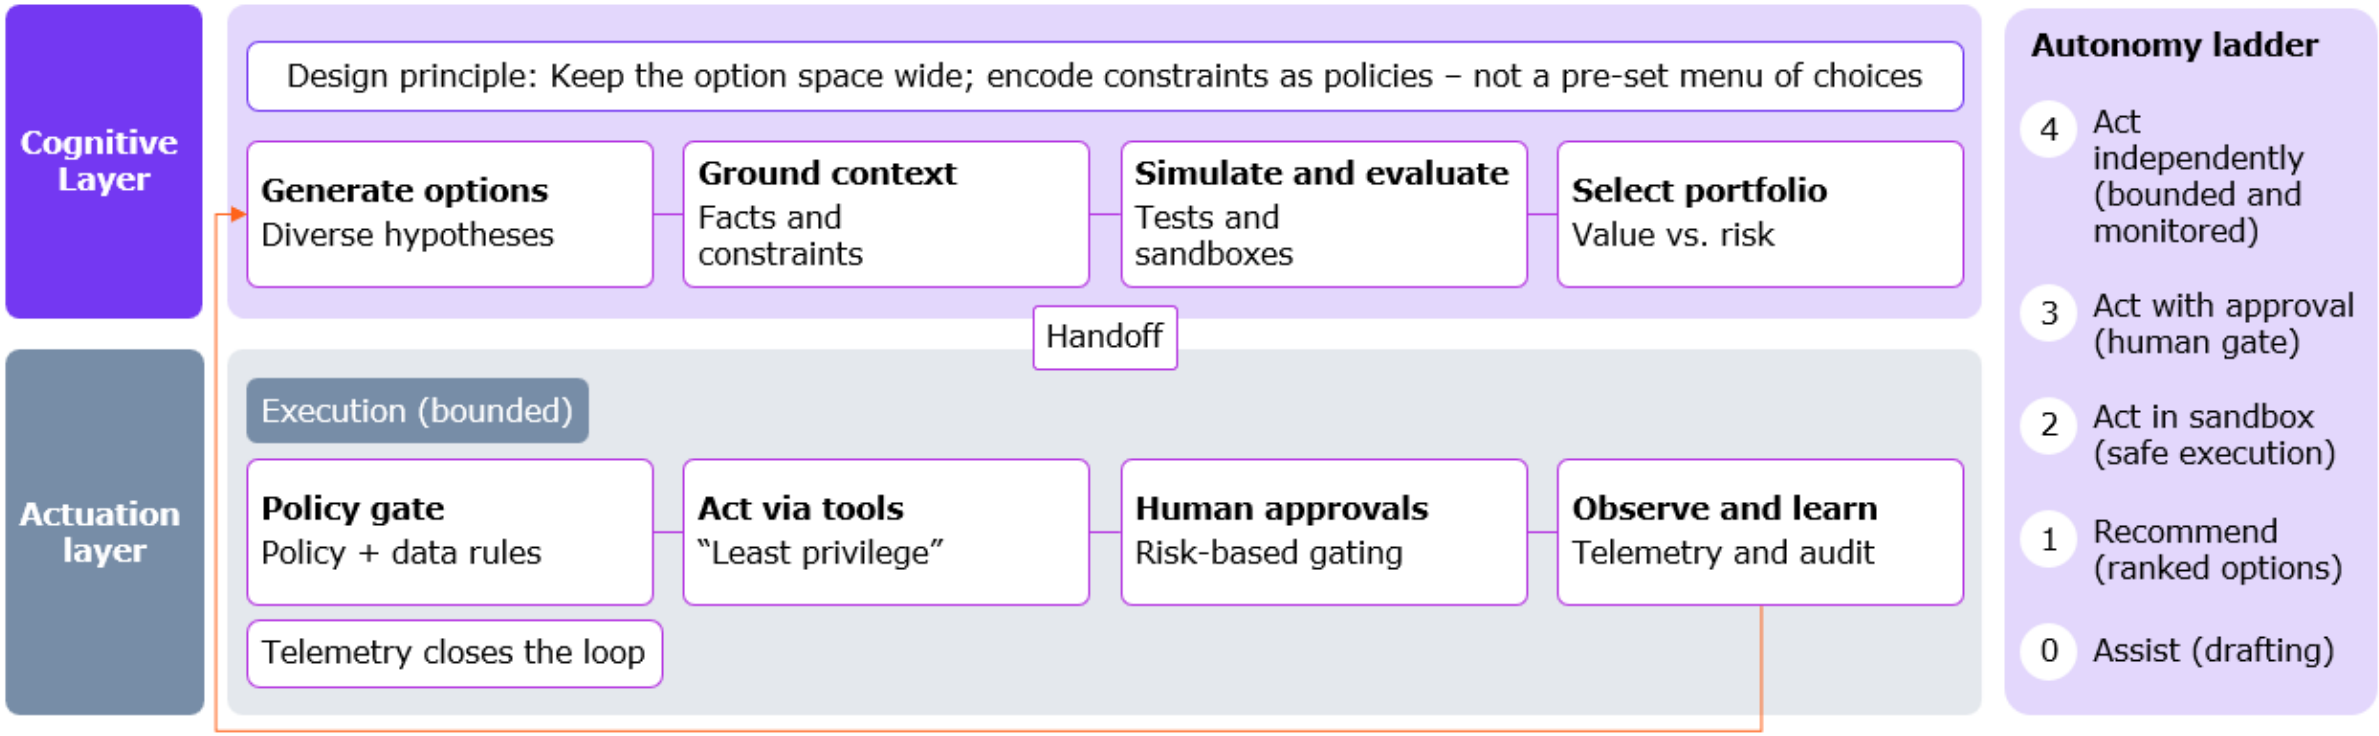

RUNWAY: Risk- segmented autonomy; **Unbounded** ideation; **Novelty** as KPIs **Work** as graphs: **Automated** evaluation; **"Yes**-if governance"

HFS © 2026 HFS Research. All Rights Reserved.

Firms must also accelerate how fast they ingest innovation

Enterprises need the access to innovation that startups offer – without the risk. Startups need enterprises to trust them – without layering on bureaucracy and extended timelines

The two sides of the speed-trust gap		
Enterprise side	**Impact**	**Startup side**
Fragmented AI policy: Legal, procurement, and IT often have divergent AI risk frameworks.	**Enterprises move too slowly to test new tech**	Enterprise readiness unknowns: Founders often do not know what is required for "enterprise-grade" maturity.
Security and compliance: Enterprises demand: SOC 2, ISO 27001, GDPR readiness, and formal risk frameworks.		Certification debt: Achieving compliance certifications is costly and time-consuming.
Procurement overload: Average vendor onboarding can take 6–12 months in Fortune 500 firms.	**Startups burn out before achieving adoption**	Sales friction: Procurement delays kill runway; average startup sales cycle to an enterprise = 9–18 months.
Integration barriers: Legacy data architectures, SSO requirements, and regulatory silos impede proof-of-concept deployment.	**Innovation supply chain breaks**	Integration complexity: Building to APIs, SSO, and data standards varies wildly across enterprises.
Cultural fit: "Fail fast" clashes with enterprise governance and risk tolerance.		Validation Bottleneck: Lack of third-party proof of credibility, resilience, or ROI.

How to solve:

Match your needs to startup solutions.

Streamline your startup onboarding and risk policies.

Develop:

- Readiness playbook
- Gap analysis
- Maturity scoring
- 100-day roadmap

HFS © 2026 HFS Research. All Rights Reserved.

10

Takeaways and Q&A

© 2026 HFS Research. All Rights Reserved.

Takeaways & Q&A

1. Tech drives the leap to the next S-curve of value – leading to Services-as-Software.
2. Stay close to what's happening in the Valley. It's business critical as the change accelerates.
3. Get set to build an agent OS – not just more agents.
4. It's time to take vibe coding seriously – software is now situational, not structural.
5. AI is becoming more pragmatic and industry-focused. Develop and operationalize verticalized AI with industry-experienced partners.
6. Track efficiency breakthroughs like you should major model releases. Ideas like mHC - and those that will follow - will make significant impacts on cost curves.
7. Training data for Physical AI is scaling rapidly – expect breakthroughs for robotics soon.
8. AGI is edging closer to delivering business units as software. OpenClaw is a warning.
9. Update your governance to ingest innovation faster and unleash AI for top-line growth.

HFS © 2026 HFS Research. All Rights Reserved.

Your HFS Research presenter

David Cushman
Executive Research Leader
david.cushman@hfsresearch.com

David is an executive research leader for HFS Research. He has a long-term focus on emerging technology, tracking OneOffice™ and OneEcosystem™ enablers from automation, artificial intelligence (AI), generative AI (GenAI), data and design thinking, Web3 and metaverse, through to process orchestration, workflow, and intelligence. He leads the HFS Hot Tech program, and is a published author (*The 10 Principles of Open Business* – Palgrave-Macmillan)

Follow me on Linkedin:
https://www.linkedin.com/in/davidcushman/

© 2026 HFS Research. All Rights Reserved.

About HFS

- **INNOVATIVE**
- **INTREPID**
- **BOLD**

HFS Research is a leading global research and advisory firm helping Fortune 500 companies through IT and business transformation with bold insights and actionable strategies.

With an unmatched platform to reach, advise, and influence Global 2000 executives, we empower organizations to make decisive technology and service choices. Backed by fearless research and an impartial outside perspective, our insights give you the edge to stay ahead.


www.hfsresearch.com


hfsresearch


www.horsesforsources.com


www.horsesmouthpodcast.com