

HOT TECH

HFS **Services-as-Software™** **Hot Tech: Hazel AI**

Author and program lead:

David Cushman, Hot Tech Editor-in-Chief, Executive
Research Leader

Accelerate legacy estate transformation through Hazel AI's agentic services



Hazel AI is an early-stage agentic modernization platform taking an HFS Services-as-Software approach to data platform and ERP modernizations, where enterprises burn money because “understanding the work” or “understanding the institutional knowledge” takes longer than building the work.

Hazel AI encodes the organization’s workflow and best practices into specialized agents with its proprietary context engine to produce deterministic, reviewable outputs that cut the time and cost to modernize.

The initial focus is on modernizing data platforms (e.g., legacy systems-of-record to cloud data platform) that are information-dense, repetitive, and slow because humans must extract, normalize, and validate inputs across many steps. Where some solutions focus on automating slices of the modernization process with AI coding tools, Hazel AI is built to orchestrate the end-to-end outcome.

Cut your “discovery” phase to increase the speed to outcome

Hazel AI starts with a problem every CIO and COO recognizes: most legacy modernization programs slow down not because code is hard to generate, but because enterprise context is hard to reconstruct. The real constraint is the discovery and translation phase, which involves surfacing tribal knowledge, undocumented business rules, legacy process logic, and stakeholder-specific exceptions that sit across old systems, files, workflows, and people.

The typical pattern is this:

- Enterprises and system integrators extract and load structured and unstructured data from complex legacy estates and systems-of-record (files, legacy tables, ERP, third-party integrations).
- They migrate it to modern platforms (e.g., SAP S/4HANA, MongoDB, GCP AlloyDB, Databricks) that power business functions.
- The expensive part is the discovery and iteration loop: figuring out what data means, mapping it correctly, and proving migration are accurate. This is where timelines stretch 9–18 months and budgets land between US\$250,000 to over US\$1 million depending on scope.

Hazel AI claims to compress this by taking the end-state intent (KPIs, outcomes, user roles) and generating the artifacts that typically require senior people doing weeks of business analysis and architecture work.

Start with the intent to deliver full-lifecycle modernization

Hazel AI’s approach is to start from intent (e.g., “modernize my banking system of legacy Oracle Forms and PL/SQL into Java microservices and a cloud database”). Without workshopping explicit and specific prompts, the platform automatically generates a detailed product/business specification (complete with legacy technical and business logic, roles, journeys, user stories, KPIs and drill-down needs).

From this, it generates high-level and low-level designs (integration architecture, mappings, cost/performance considerations) for review in minutes, a fraction of the time (often months) required with manual processes. Customers then review and amend as necessary, and those designs effectively become the specifying instructions to generate end-to-end code, tests, and deployment scripts. These can then be handed to the customer as a deployment pipeline. The whole process, from capture to deployment, is cut to from a 9–18 months of engagement to nearer 10 weeks end-to-end.

Today, specification and compliance compression; tomorrow, autonomous deployment

Today, customers manually run the deployment pipeline through staging and promote to production after co-verifying outcomes. Hazel AI's stated direction is to move toward customers granting deeper cloud access so it can deploy and maintain workflows continuously and autonomously. Enterprise leaders should therefore treat Hazel AI today as a delivery acceleration and specification/compliance compression play.

Pay per workflow, pay per use

Hazel AI's business model is an upfront per-workflow charge to build the workflow (outcome-based rather than tied to headcount or time consumed) with a recurring usage-based fee for running workflows.

This may prove cheaper than most CIOs have experienced, but Hazel AI's real value lies in outcomes such as speed to business value, higher-grade outputs, and making scarce experts scale across more work.

If Hazel AI can prove value consistently, this becomes a practical Services-as-Software engine: codifying delivery into repeatable agentic workflows that scale without scaling headcount.

Compressed delivery timelines for Exceego

Yunus Mohammed Khan, Senior Director of Operations at Exceego Infolabs, said Hazel AI has materially accelerated their EHSWatch roadmap by turning detailed domain-specific requirements into documentation, prototypes, and code far faster than traditional delivery. Work completed in roughly three months, and potentially two without delays, would normally have taken seven to eight months, sharply compressing the path from requirement capture to usable AI features.

Yunus iterated that Hazel AI's speed and ability to understand EHSWatch requirements helped embed AI across modules, improving go-to-market readiness and customer outcomes. He would like to see the company expand its time zone coverage or regional support to reduce response lag for his Gulf-based teams.

HFS' take

Hazel AI is tackling one of the key constraints in enterprise transformation: institutional knowledge trapped in people and process. Its claim that agents can generate production-grade specs, designs, and delivery assets is compelling because it attacks the budget sink: figuring out what should already be known and documented, but rarely is.

But enterprise leaders should be clear-eyed. If you let an agentic platform encode your workflow, you can scale faster. But you can also

industrialize the wrong assumptions at machine speed. Hazel AI must prove governance, auditability, and security to earn the right to operate inside regulated, high-stakes environments.

If Hazel AI delivers on its promises, it is a direct threat to time-and-materials delivery models and a potential unlock for enterprises stuck in modernization backlogs. The real test if it is ready to show you repeatable outcomes beyond impressive first runs.


Vendor fact sheet

- **Headquarters:** San Jose, California
- **Founded:** 2025
- **Key executives:** CEO and Co-founder Sneha Shah; CTO and Co-Founder Ram Balagurunathan
- **Funding:** Undisclosed
- **Partners and ecosystem:** All Frontier AI labs (Open AI, Anthropic), cloud platforms (GCP, AWS, Microsoft, SAP, MongoDB, etc.)
- **Clients:** Has already partnered with Global 2000 enterprises, leading SI partners, and cloud service providers across industry verticals for modernization and migration
- **Solutions set:**
 - SAP S/4HANA modernization – greenfield and brownfield
 - Legacy applications, databases into modern app, cloud databases
- **Industry coverage:**
 - Data platforms and database modernization ecosystems
 - Cross-Industry for SAP

The HFS Hot Tech designation and Services-as-Software

HFS Hot Tech is an exclusive group of emerging players, each with a differentiated value proposition aligned with creating momentum toward Services-as-Software™. HFS analysts regularly speak with numerous exciting startups and emerging players. We designate a select few as HFS Hot Tech based on their offerings' distinctiveness, ecosystem robustness, client impact, financial position, and, in this case, their impact on Services-as-Software. HFS Hot Tech companies may not have the scale and size of more established rivals, but they have the vision and strategy to impact and disrupt the market. In the rapidly changing AI-led operations space, enterprises realize they can't be everything to everyone.

HFS Hot Techs offer a range of approaches toward Services-as-Software



Source: HFS Research, 2026

HFS Hot Tech report author



David Cushman
Hot Tech Editor-in-Chief,
Executive Research Leader

David is an Executive Research Leader at HFS and Editor-in-Chief of the HFS Hot Tech program. He also leads our OneOffice Emerging Technology Practice, is our strategic lead on Generative AI, Web3, and metaverse, and covers automation and employee experience. He is a published author (*The 10 Principles of Open Business*, Palgrave-Macmillan), a former tier-1 consulting director, and a digital strategy and innovation expert with leadership experience in start-up, scale-up, and enterprise digital transformation.

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](https://www.linkedin.com/company/hfsresearch)



www.horsesforsources.com



www.horsesmouthpodcast.com