

Designing Generative AI Deployments for Success

Right Structure at the Design Stage to Maximise Impact

“GenAI programs succeed not because of the models they choose, but because of the structural decisions made before deployment.”

56%

Projects at Scale

83%

Centralised Governance

28%

Security as #1 Barrier

42%

Facing Talent Gap in Architecture and Design

Cambridge Service Alliance

University of Cambridge

× HCLTech

Professor Mohamed Zaki

Deputy Director,
Cambridge Service Alliance

Helen Zhao

Researcher,
Cambridge Service Alliance

Archie Upadhyay

AI Product Strategist,
HCLTech-Cambridge Service Alliance

Vinod Tete

Head of Innovation,
HCLTech

Mohit Mathur

AI & Innovation Lead,
HCLTech

Foreword: The Moment of Shift

Enterprises are investing heavily in GenAI-driven service transformations; yet translating these investments into consistent business and customer outcomes remains a challenge. Through surveys, interviews and data analysis, the CSA–HCLTech research examines this gap, revealing how organizations are reimagining service delivery to drive impact, growth and resilience in the GenAI era.

The purpose of this report is to describe how organizations are using generative AI (GenAI) and how early design decisions shape the results they achieve. The study looks at real GenAI implementations across industries and examines how teams move from experimentation to larger-scale use. It focuses on how organizations balance efficiency and experience goals, manage compliance and data needs, choose technology, handle risks and require the right skills. The aim is to provide a clear, evidence-based view of the patterns that are emerging, as GenAI becomes part of everyday work. It brings together findings from a cross-industry survey and in-depth interviews to understand how teams are planning GenAI work, what structures they are using, and how design-stage decisions shape later deployment. The study covers a range of industries and functions, keeping all data and exhibits as originally reported.



“Emerging evidence from our multi-method inquiry suggests that GenAI’s true contribution lies not in algorithmic novelty but in how rigorously an organization structures its surrounding ecosystem. When data governance, model observability, and human-centered oversight are deliberately codified, AI systems exhibit greater reliability, reproducibility, and alignment with intended outcomes. Under these conditions, GenAI evolves from a technological enhancement into a measurable, theory-grounded socio-technical capability that can be iteratively validated and scaled.”

– Professor Mohamed Zaki, Deputy Director, Cambridge Service Alliance

Across the data reviewed, five shifts stand out:

- **From pilots to operations.** What began as small, contained experiments is now moving into day-to-day workflows across functions and industries, with GenAI becoming part of regular operations.
- **From efficiency to experience.** Early efforts focused on saving time and reducing effort. Today, organisations are using GenAI to elevate how work feels for employees, how interactions land with customers, and how end-to-end journeys come together.
- **From scattered to centralised.** Most initiatives now run through central teams or follow clear, central policies. Fully independent, uncoordinated approaches are largely absent in the evidence reviewed.
- **From cost concerns to compliance concerns.** Security and compliance are not front and centre. As GenAI touches sensitive information and critical workflows, questions of safety, privacy, and responsible use move to the core of decision-making.
- **From hiring to reshaping.** Many organisations find it hard to staff GenAI projects with the right mix of skills, particularly in architecture, design, and specialist technical roles. At the same time, jobs themselves are changing: people are expected to apply judgement, shape narratives, and refine AI-generated outputs.

Technology choices reflect this phase of maturity. Teams typically start with widely used large language models (LLMs) because they are stable and quickly to deploy, while interest in more advanced, agent-style systems is building. As GenAI moves from pilots to scale, organisations are widening goals from efficiency to experience, consolidating work under centralised and policy-guided models, adopting standard LLMs first while exploring agents, encountering security and compliance as the new bottlenecks, and reshaping roles to close talent gaps. Throughout the findings, one idea repeats: outcomes improve when these choices are structured at the design stage, so adoption, experience, operating model, technology path, risk, and talent design all reinforce each other as programmes scale



“The question today in enterprises is not if GenAI and agentic systems work — they’re grappling with how fast it can reshape customer experience and revenue models. The task of reimagining your systems and processes with Agentic and GenAI are non-trivial and is at the core of deploying these systems at scale in enterprises to derive value.”

– Ashish Kumar Gupta, Chief Growth Officer, HCLTech



“Enterprises can shape the frontier of AI driven growth by pairing platform standardization with targeted market experimentation. With security, data contracts, and observability as firm foundations, they can innovate faster and scale more safely. Real advantage arises when governance provides stability while edge teams push AI’s boundaries.”

— Swapan Johri, President, Growth Markets, HCLTech


EXECUTIVE OVERVIEW

The GenAI Inflection Point: From Experiment to Enterprise

The study examined how leading organisations across financial services, manufacturing, retail, healthcare, logistics, and energy are designing their GenAI programmes for scale. Drawing on a cross-industry survey and in-depth executive interviews, we surface the design-stage decisions that most strongly predict whether GenAI investments generate sustained, measurable business value -- or stall in perpetual pilot mode.

Outcomes improve when adoption, experience, operating model, technology path, risk, and talent design are structured at the design stage, so they reinforce each other as programmes scale.

SIX KEY FINDINGS AT A GLANCE



Design-Stage Imperatives for IT & AI Leaders

- 1 Treat Architecture as Strategy**

The first 90 days of design determine your ceiling. Choices about data governance, model observability, and human oversight made now define what is achievable at scale. Resist the urgency to deploy before the architecture is sound.
- 2 Shift Your Investment Thesis: Efficiency Pilots + Experience Transformation**

Organisations showing the highest EBIT impact redesign full workflows rather than laying in point tools. Experience-led use cases force process redesign, which is the strongest predictor of bottom-line value. Build a dual portfolio from the outset.
- 3 Centralise Governance; Decentralise Execution**

83% of successful programmes use central IT or policy-guided models. Build the enterprise spine with shared platforms, common guardrails, and reusable components, then allow business units to sprint. The goal is guided autonomy that empowers teams while maintaining coherence.

4

Native Security and Compliance

Security and compliance (28%) now outrank cost as the #1 bottleneck. Design governance in from day one with risk tiers, model monitoring, privacy by design, and human-in-the-loop frameworks. Retrofitting this onto live systems costs 3-10x more than building it at design stage.

5

Start with Standard LLMs; Architect for Agentic Complexity

78% start with standard LLMs for good reason, such as speed or manageability. But build your architecture anticipating multi-step agentic systems. Retrofitting observability and guardrails onto agents is significantly harder than building it in from the start.

6

Invest Now in the Roles That Are Already Needed

61% face skill gaps in the roles they need most: solution architects, AI designers, domain-aware curators. These are present bottlenecks, and upskilling programs and ecosystem partnerships must start immediately.


FINDING I

Adoption Is Shifting from Pilots to Scale

GenAI is entering the operational layer — the pilot era is ending

56% of organisations report GenAI initiatives at steady-state or scaled implementation. Teams are no longer treating GenAI projects as isolated trials but are embedding them into operational workflows. Much of this momentum comes from the edges, where teams closest to delivery experiment first, validate outcomes, and expand as confidence grows.

GenAI Project Maturity Stage (%)



Deployment by Business Function (%)




Exhibit 1: Most GenAI projects are now at POC or steady-state | Exhibit 2: Customer Service (23%) and HR (17%) lead deployment by function

Customer service and HR lead because they benefit most from conversational interfaces and content automation. This reflects a broader adoption pattern: enterprises deploy GenAI first where impact is direct and visible, before scaling to more complex, regulated, or mission-critical domains.

■ What This Means for IT Leaders

- Edge-led innovation requires a central spine. Build the shared platform infrastructure that allows local experiments to connect, replicate, and scale.
- If your programme is still in perpetual POC, the bottleneck is likely governance or talent design rather than technology readiness. Diagnose before adding more pilots.
- Prioritise customer service and HR for fastest time-to-scale, then use those learnings to accelerate entry into higher-complexity domains.

“Most of our early GenAI work started within functional teams. The people closest to delivery experimented first, tested what worked, and refined their own use cases. Once those smaller experiments showed promise, we expanded them into broader workflows. This bottom-up start helped us move from pilots into real operational use much faster.”

— Technical Lead, Global Media & Communications Company

FINDING II

From Efficiency to Experience (and Growth)

"GenAI is no longer just about speed — it is about how interactions feel"

While efficiency-oriented gains remain the most common use (56%), experience-led applications (44%) are rapidly gaining ground. **The strategic conversation is shifting from "How much faster?" to "How much better?"** Organisations are using GenAI to elevate how work feels for employees, how interactions land with customers, and how end-to-end journeys come together.

USE CASE CLUSTER	SHARE	PRIMARY APPLICATIONS
Internal Operations & Infrastructure	39%	HR process automation, employee comms, workflow streamlining
Customer Service & Interaction	22%	AI chatbots, customer-facing support, service personalisation
Logistics, Planning & Operations	22%	Transport planning, risk management, operational automation
Knowledge & Analytics	17%	Document summarisation, knowledge retrieval, decision support

Exhibit 3: Four GenAI use case clusters, with internal operations & infrastructure leading at 39%

Why experience-led use cases deserve priority. Experience improvements (NPS/CSAT uplift) correlate with higher retention, cross-sell, and lower cost-to-serve. High performers pair efficiency with growth objectives and achieve materially higher EBIT impact because they redesign full workflows. Experience-led use cases also create an adoption flywheel. When journeys improve perceptibly, frontline adoption accelerates, raising data quality and compounding efficiency gains.

■ What This Means for IT Leaders

- Build a dual portfolio: efficiency pilots for near-term ROI justification, and experience transformation bets for competitive differentiation.
- For experience-led use cases, engage the business in workflow redesign, not just tool procurement. GenAI embedded in redesigned processes delivers 3-5x the impact of GenAI added onto unchanged workflows.
- Track NPS/CSAT and adoption rates alongside cost metrics. Experience is now a measurable, board-level outcome.



"For years, digital transformation was measured through efficiency — reduced cycle times, automated tasks, cost optimisation. GenAI is shifting the paradigm. The most sustained enterprise value comes from experience transformation: how interactions feel, how workflows operate, how intelligently systems anticipate needs. Efficiency may ignite the journey, but experience fuels long-term competitive differentiation."

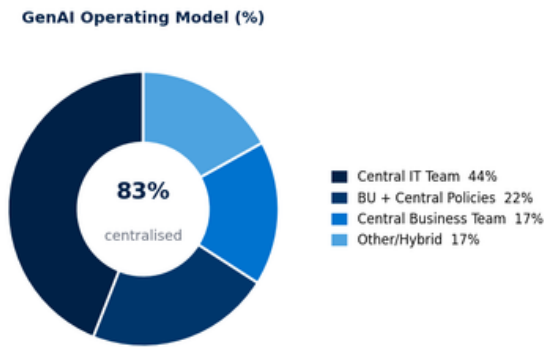
— Vinod Tete, Head of Innovation, HCLTech

FINDING III

Centralised IT and Policy-Guided Models Dominate

Zero organisations run fully federated GenAI by deliberate design

83% of GenAI programmes run through a central IT team (44%), central business team (17%), or individual BUs guided by central policies (22%). No respondents report a fully federated model. This is not coincidence. It reflects the enterprise reality that ungoverned AI proliferation creates unacceptable risk to data integrity, compliance, and strategic coherence.



Why Centralisation Works

Consistent security: Standards applied uniformly across all deployments

Reusable components: Reduce duplication; accelerate future deployment

Coherent compliance: Single governance posture satisfies regulators

Shared observability: Unified monitoring and performance measurement

Guided autonomy: BUs retain freedom; operate within common guardrails

Exhibit 4: 83% of organisations use centralised or policy-guided models; none reported fully federated

■ What This Means for IT Leaders

- Establish a central AI operating model before deployment proliferates. Retrofitting governance onto live systems is 3-10x more costly than building it at design stage.
- The "guided autonomy" model is the proven sweet spot: define non-negotiables (security, data handling, evaluation standards) and give teams freedom on everything else.
- Bring business, technology, and governance stakeholders together at design.

“Balancing competing priorities across the organisation is often the hardest part, so we rely on a structured evaluation process that brings business, technology, and governance stakeholders together early.”

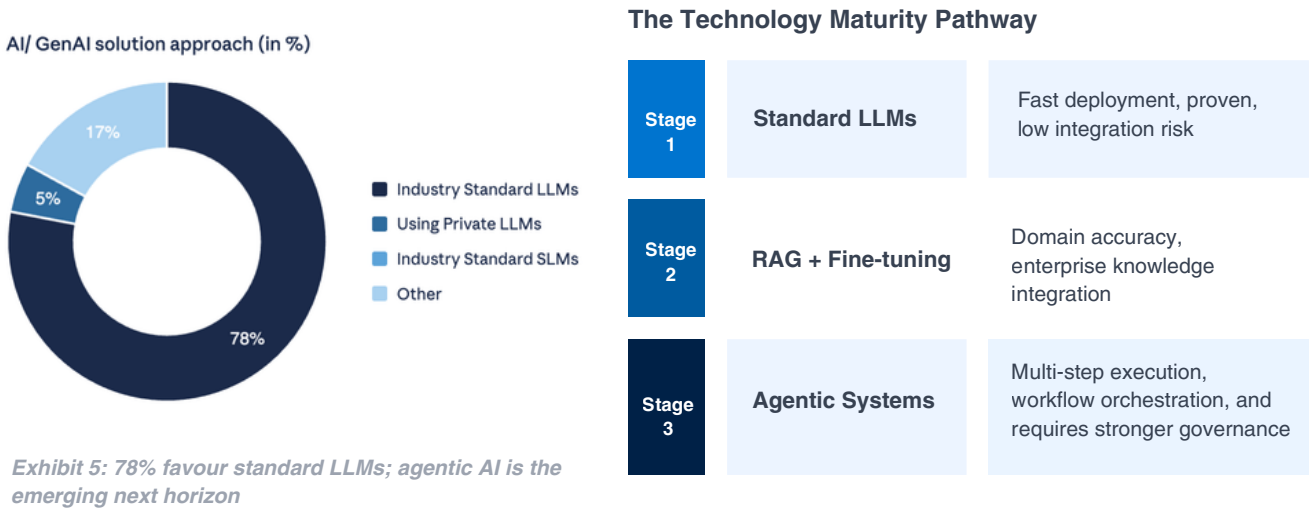
— Senior Product Manager, Telecommunications Provider

FINDING IV

Technology Choices: Standard LLMs First, Agents Emerging

Start proven; architect for the complexity ahead

78% of organisations start with industry-standard large language models. These offer accessible entry points — stable, quick to deploy, widely supported. At the same time, agentic AI interest is building rapidly for document analysis, enterprise search, service support, and multi-step workflow orchestration.



The critical architectural insight: AI agents deliver significantly greater value when embedded in redesigned processes with strong guardrails and observability. The shift to agentic systems is not a technology upgrade — it requires fundamentally stronger governance infrastructure. Multi-agent systems that handle multi-step tasks are harder to monitor, audit, and control. This must be anticipated at design, not retrofitted at deployment.

What This Means for IT Leaders

- Start with standard LLMs for justified reasons, but architect today as if you will need agentic capabilities within 18 months. Retrofitting observability onto agents is significantly harder than building it in.
- Follow a ‘vital few’ approach to use case selection: cluster possibilities into priority themes and secure executive sponsorship for must-win battles before scaling.
- Use LLMs to analyse operational data such as tickets and user feedback to surface highest-value use cases. This is evidence that builds internal alignment faster than intuition.

“Identifying the right GenAI use cases is challenging because the volume of potential scenarios can easily lead to fragmentation. We follow a vital-few approach: generating possibilities, clustering into priority themes, and securing executive sponsorship for must-win battles. We use LLMs to analyse tickets and user feedback — evidence that makes it easier to build internal alignment and justify investment.”

— Vice President, Major Healthcare Services Group

FINDING V

Security & Compliance Are the New Scalability Constraint

Governance is not a cost of GenAI but the foundation on which enterprise-scale AI is built

Security and compliance (28%) now outrank cost overruns (17%) as the most frequently cited barrier. As GenAI touches sensitive information and critical workflows, questions of safety, privacy, and responsible use have moved from the sidelines to the core of decision-making. This marks a transition to mature deployment where GenAI must operate reliably within existing controls and regulatory frameworks.

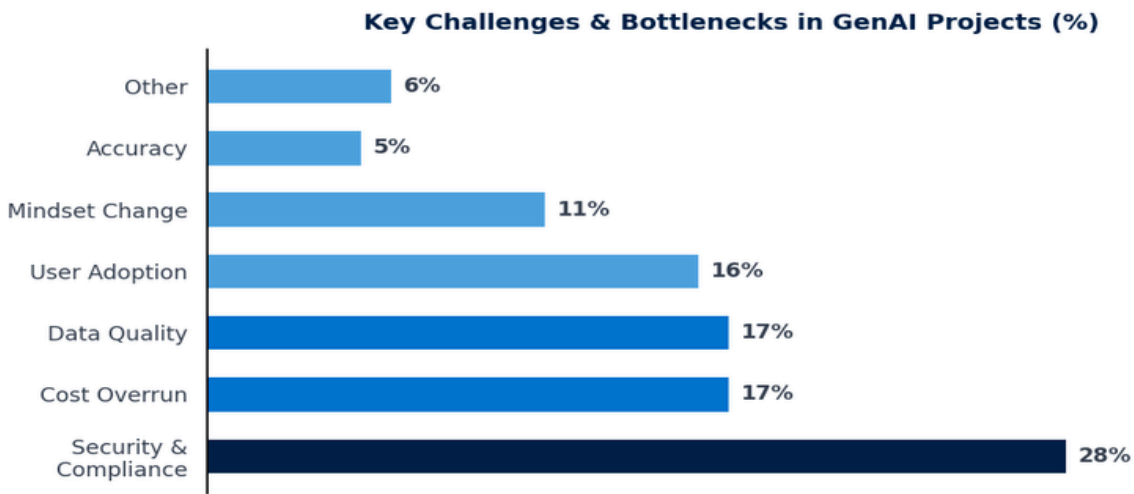


Exhibit 6: Security and compliance (28%) is the leading barrier — ahead of cost, data quality, and adoption

Five pillars of GenAI governance that distinguish high-performing organisations

RISK TIERING	MODEL MONITORING	PRIVACY BY DESIGN	HUMAN IN THE LOOP	AUDIT TRAILS
Classify use cases by sensitivity & consequence of failure	Continuous performance tracking & output drift detection	Data governance from day one, not post-deployment	Mandatory review for high-stakes outputs; human accountability	Every output traceable for explainability & regulatory compliance

■ What This Means for IT Leaders

- Governance is a velocity enabler. Organisations with clear risk tiers and transparent model monitoring move faster because confidence replaces hesitation at every stage-gate.
- Treat security and compliance as architecture requirements. The cost of retrofitting is 3-10x higher than building it in at design stage.
- For high-stakes content, maintain mandatory human review. Technology executes the task; humans must retain accountability for the output.



“In every successful GenAI deployment I have overseen, governance has played a decisive role — not as a constraint, but as a catalyst. With clear risk tiers, transparent model monitoring, and human-in-the-loop protocols, confidence rises, adoption accelerates, and value realisation deepens. Governance is the foundation on which enterprise-scale AI is built.”

— Mohit Mathur, AI & Innovation Lead, HCLTech

FINDING VI

Talent Scarcity and Role Redesign

GenAI skill gaps hit hardest in architecture, engineering, and domain expertise

42% of organisations report difficulty finding GenAI solution architects and designers, with a further 26% flagging gaps in both technical engineering and domain-specific skills. This reflects the reality that GenAI requires combined strengths rather than a narrow set of capabilities. At the same time, roles are shifting fundamentally. As routine tasks automate, human expertise must move toward interpretation, curation, oversight, and judgment.

Note: Multi-response question. Respondents could select more than one option.

SKILL AREA	TALENT SCARCITY	IMPLICATION FOR PROGRAMME DESIGN
GenAI Solution Architects & Designers	42% report gap	Critical bottleneck for design-stage quality; ecosystem partnerships essential
GenAI Technical Skills (Engineering)	26% report gap	Creates deployment delays; internal upskilling delivers fastest ROI
Domain Knowledge for AI Application	26% report gap	Required for experience-led use cases; build cross-functional teams
All Required Skills Generally Available	47% report available	Strong capability; focus on retention and accelerating to advanced use cases

Exhibit 7: Talent gaps are broad. Solution architects and designers are the most constrained and critical role

Roles are evolving from execution to interpretation. The emerging roles such as AI curators, orchestrators, narrative designers, and decision architects are not abstract future ideas. They are already appearing in teams deploying GenAI at scale. As routine tasks automate, what remains is the human capacity to guide, interrogate, and refine AI systems toward meaningful outcomes.

■ What This Means for IT Leaders

- Map talent gaps at the design stage. Solution architects and designers are the most constrained role, so secure or develop this capability before it becomes a critical path bottleneck.
- Invest in hybrid intelligence: people with critical reasoning, domain depth, and ethical judgment who can guide, interrogate, and refine AI systems.
- Design roles around the new work. The organisations most likely to succeed will build the capability to steer AI safely and creatively at every stage of deployment.



“As GenAI scales, execution becomes automated but accountability does not. The workforce of the future will be defined by those who can question, refine, and responsibly steer AI systems.”

— Archie Upadhyay, AI Product Strategist, HCLTech-Cambridge Service Alliance

FINDING VII

Value Realisation: Redefining What Success Looks Like

ROI alone is no longer sufficient. Success is financial, operational, experiential, and technical

Organisations define GenAI success through a multi-dimensional lens. Financial measures (ROI, cost savings) remain important at **38%**, but operational efficiency (**23%**) and user satisfaction and adoption (**23%**) appear just as prominently. This signals a move toward a more balanced understanding of impact, where business value is seen in both efficiency gains and experience improvements.

How organisations define GenAI success (%)

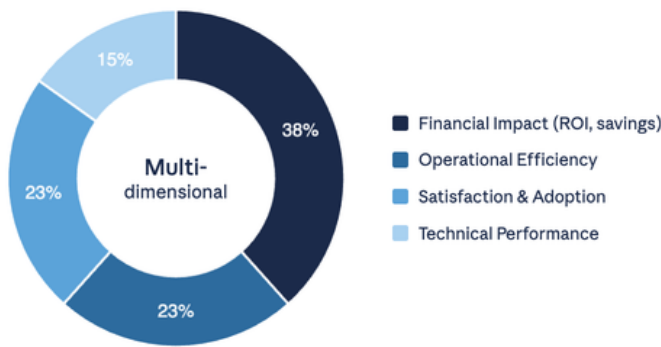


Exhibit 8: Success is multi-dimensional — financial, operational, experiential, and technical

What High-Performing Organisations Measure

- **Financial:** ROI, cost savings, headcount reduction
- **Operational:** Time/effort savings, efficiency gains
- **Experience:** User adoption, employee satisfaction, customer satisfaction
- **Technical:** Prediction accuracy, system reliability

GenAI adoption is strongest in functions where benefits are immediate and visible. Over time, organisations evaluate success not only through cost metrics but also by how smoothly GenAI integrates into daily work and improves user experience.

■ What This Means for IT Leaders

- Define your success scorecard at the design stage. Include financial, operational, experience, and technical metrics and align them to business stakeholder priorities before work begins.
- Track adoption metrics alongside ROI. If users are not adopting, ROI calculations are theoretical. User satisfaction is a leading indicator of sustained value.
- Report progress against all four dimensions to the business. Experience and adoption metrics are increasingly what boards and growth leaders want to see.

“For us, success showed up when people actually started using the tool because it made their work easier. When teams saved time, produced cleaner outputs and felt more supported in daily tasks, adoption grew on its own. That’s when we knew the project was working — not just when costs went down, but when the experience improved.”

— **Technical Leader, Renowned Manufacturing Firm**

CONCLUSION

The Design-Stage Dividend

Across all seven findings, one pattern repeats with remarkable consistency: **the decisions made at the design stage strongly influence how well GenAI scales later.** Early choices about where to start, how to govern, which models to use, how to manage risk, and how to build talent shape the long-term trajectory of every programme.

GenAI is maturing into an enterprise-scale capability where success depends less on technology novelty and more on organisational design. The organisations winning are not necessarily those with the most advanced models, but those that structured their programmes most deliberately at the outset.

DESIGN DIMENSION	THE DESIGN-STAGE QUESTION TO ANSWER	IF YOU DEFER THIS...
Adoption Strategy	Where will we start and what does the scale pathway look like?	<i>Fragmented pilots; no clear path to operational integration</i>
Experience Design	Are we redesigning workflows, or adding tools onto unchanged processes?	<i>Efficiency gains only; experience and growth potential unrealised</i>
Operating Model	Who governs, who executes, and what are the non-negotiables?	<i>Duplication, inconsistency, compliance failures at scale</i>
Technology Architecture	Are we architecting for agentic complexity, even if starting simple?	<i>Costly architectural rework when capabilities need to advance</i>
Governance & Security	Are security, compliance, and oversight baked into the design?	<i>Programme halts, data breaches, regulatory exposure</i>
Talent & Roles	Do we have the architects, curators, and oversight roles we need?	<i>Design quality suffers; deployment stalls; adoption fails</i>

Exhibit 9: The GenAI Design-Stage Decision Framework with six dimensions that determine programme success

RESEARCH METHODOLOGY

This study was conducted in two phases. A **quantitative survey** (14 weeks, Microsoft Forms) engaged GenAI Architects, Service Delivery Leaders, Business Function Leaders, IT Leaders, and Data Architects across multiple organisations and cohorts. **Qualitative in-depth interviews** (12 weeks) with senior participants covered strategy alignment, feasibility assessment, deployment decisions, and future skills. The study spans Financial Services, Manufacturing, Automotive, Retail, Life Sciences, Healthcare, Travel, Transportation, Logistics, and Energy & Utilities. This study offers directional insights supported by multiple sources of information. **Any decisions made at board level should consider the full set of evidence in this report.**

ABOUT THE AUTHORS

Professor Mohamed Zaki
Deputy Director,
Cambridge Service Alliance

Helen Zhao
Researcher,
Cambridge Service Alliance

Archie Upadhyay
AI Product Strategist,
HCLTech-Cambridge
Service Alliance

Vinod Tete
Head of Innovation,
HCLTech

Mohit Mathur
AI & Innovation
Lead, HCLTech