AI — A to Z

Everything a CIO Needs to Know to Enable AI Across the Enterprise

GS Jha | CIO & CISO, QuantumScape

ITx Collective • Backbone Forum • March 2026

30 minutes | 3 live polls | 26 things that matter | 1 Monday action

AI: 1950 — 2050

The 100-year arc — from logic machines to autonomous intelligence

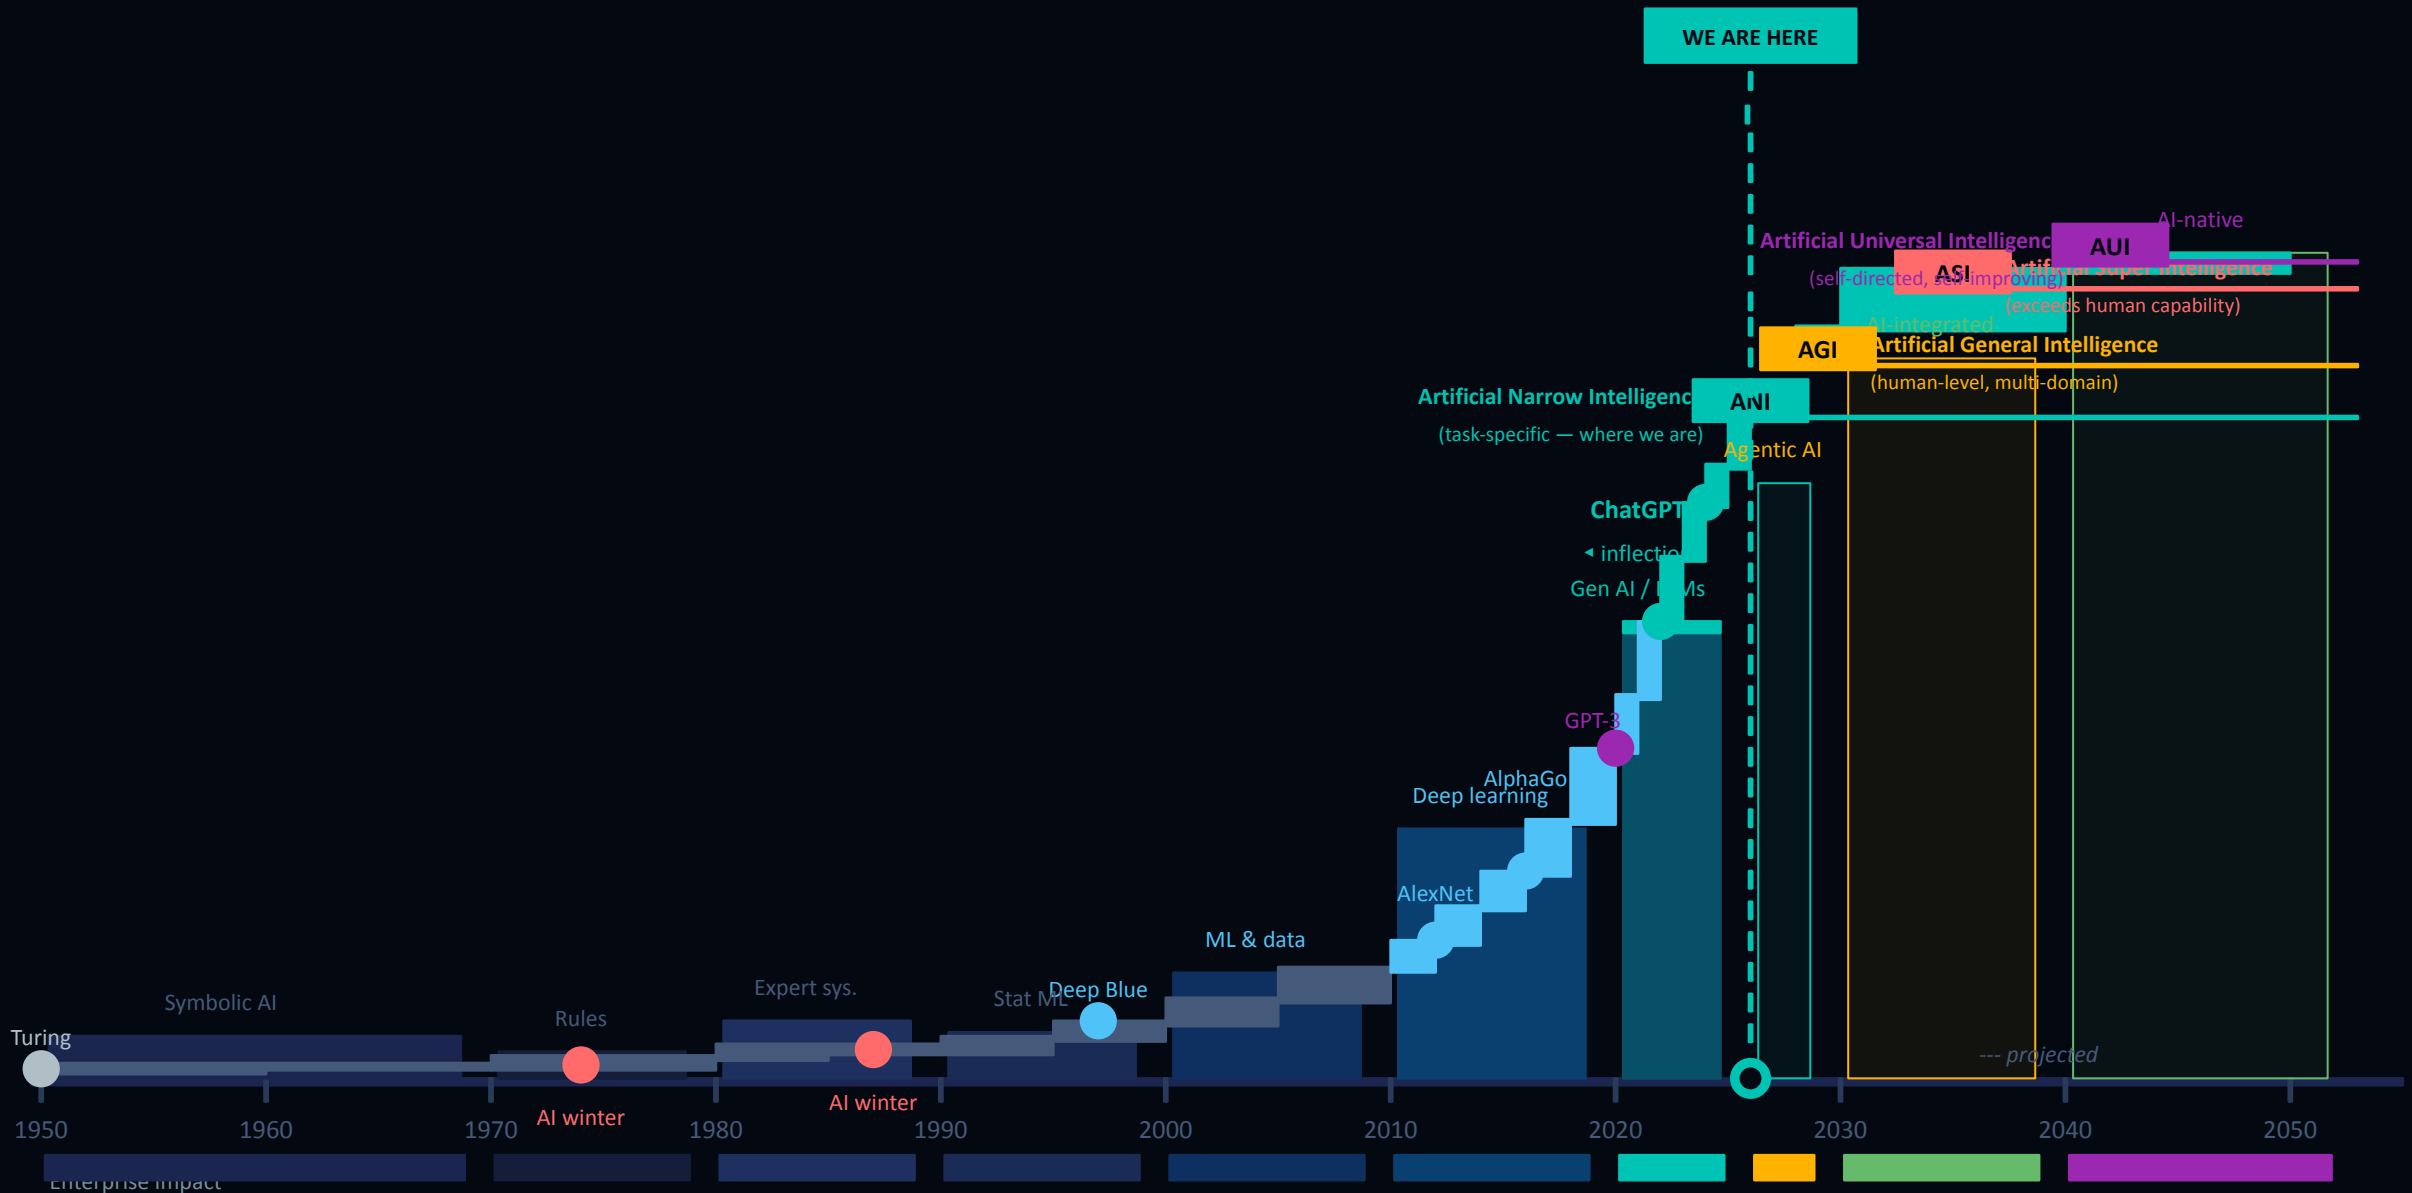

AI is not transforming one part of the enterprise - it is reshaping all of them at the same time, from the center out.

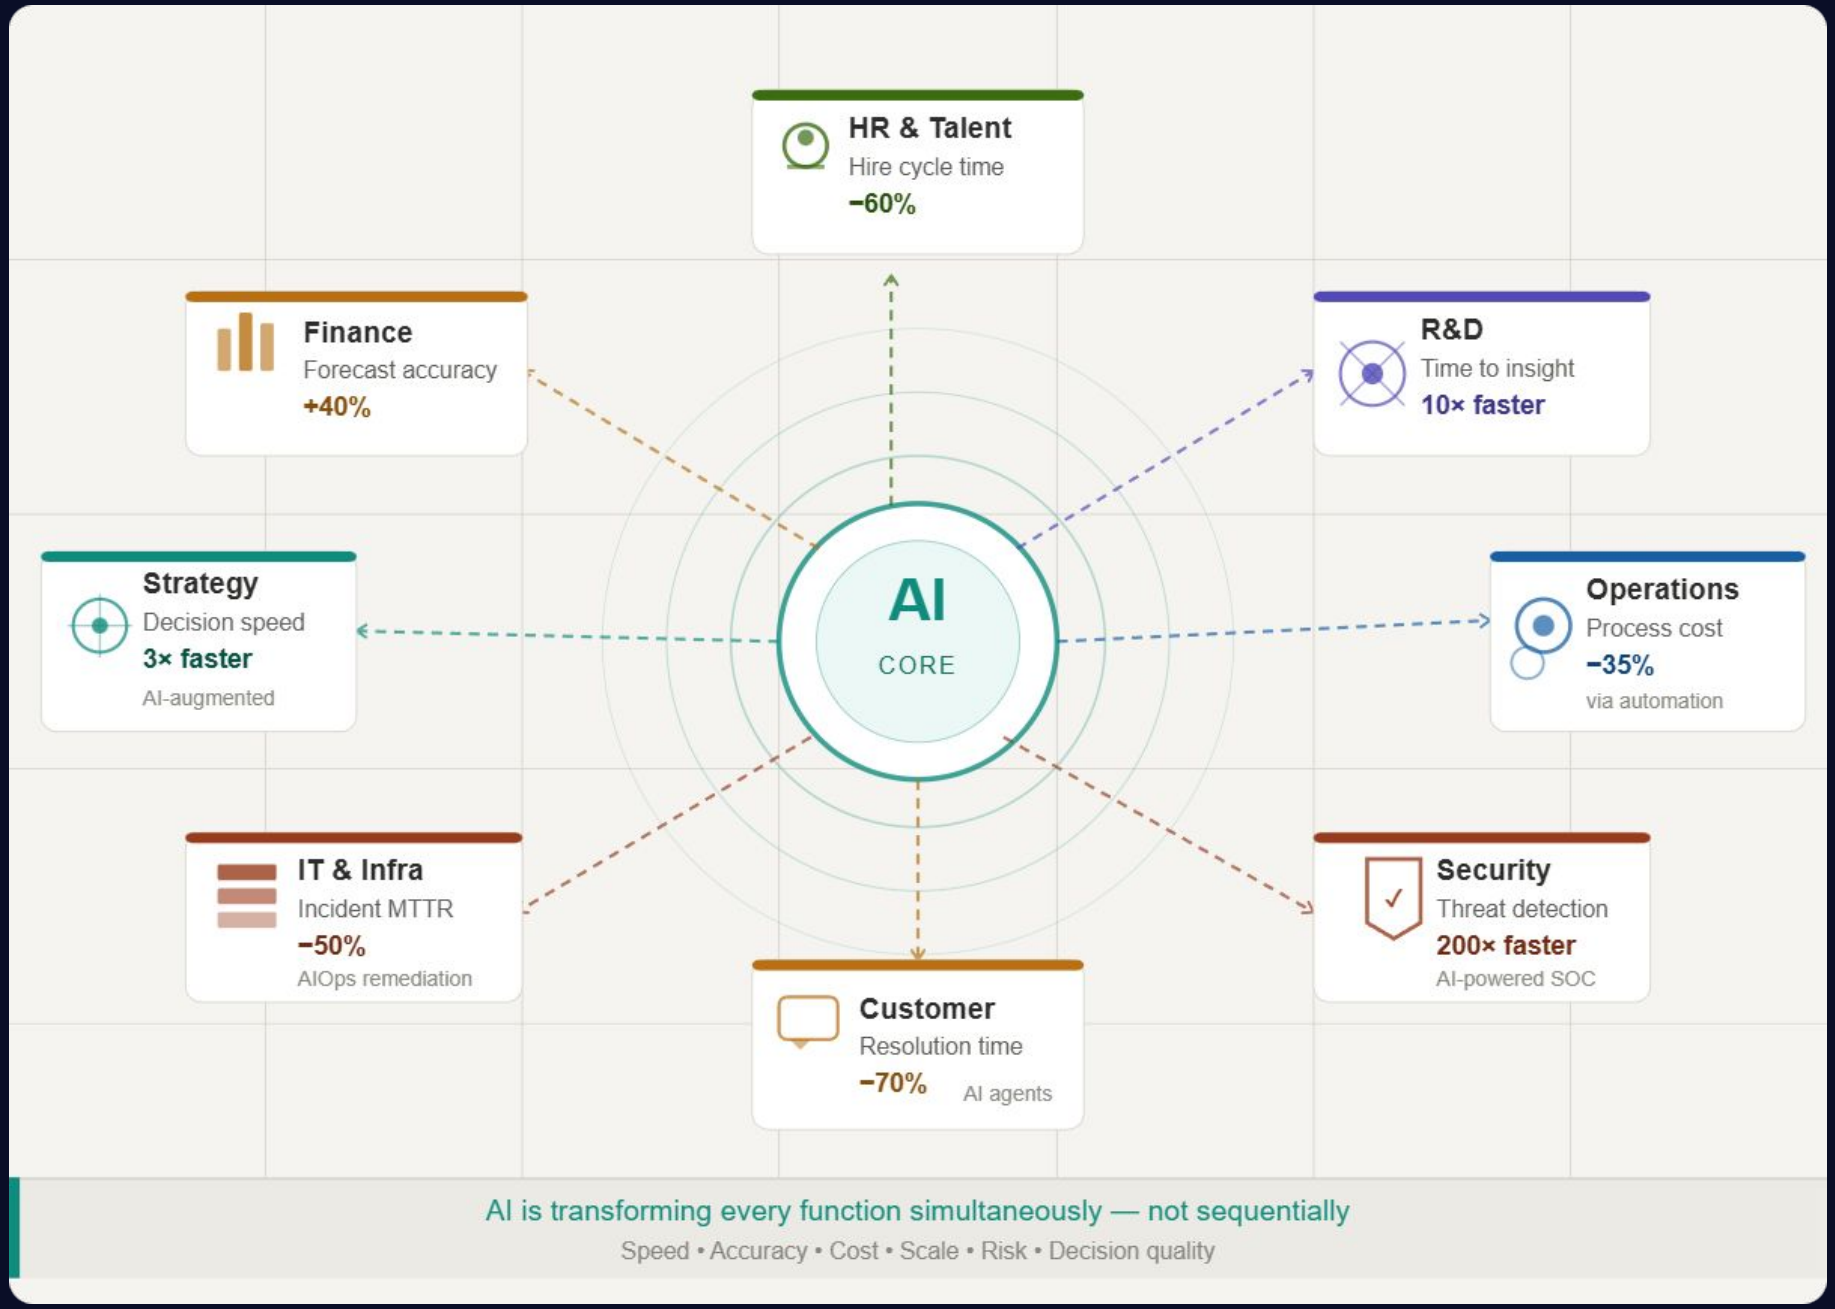

The world changed.

Did our enterprise?

⚡ Speed

GPT-4 launched March 2023. By end of 2024, it was obsolete. The model that leads today will be a footnote in 18 months. Enterprise cycles run in years. AI runs in weeks.

😰 Fear

41% of workers fear AI will eliminate their job. Our best people are already using AI tools we haven't approved. Both facts are true simultaneously.

🔓 Enablement

The organizations moving fastest aren't the biggest. They're the ones with the clearest strategy — use case discipline, data governance, and P&L-owned pilots.

🔒 Cybersecurity

Shadow AI is the new shadow IT — except the blast radius is larger. Prompt injection, data leakage, unvalidated outputs. The attack surface just expanded.

WHERE THINGS STAND

70%

of AI pilots never reach production

$4.4T

projected AI economic impact by 2030

3x

more spent on data prep than on the model

1 in 4

enterprise employees already using shadow AI

This session is our map.

26 letters. Every decision a CIO needs to own to enable AI across the enterprise.

The AI Architecture Reality — Five Systems, One Enterprise

Challenges & Opportunities from Coexistence

🏛️ Legacy Systems	🤖 Pure-Play Gen AI	🎯 Fit-for-Purpose AI	⚙️ AI in Legacy Apps	🔧 Homegrown AI
ERP, CRM, Mainframe	ChatGPT, Copilot, Claude standalone	Veeva, ServiceNow AI, domain tools	SAP AI, Salesforce Einstein, Oracle AI	Custom LLMs, internal tools, RAG

⟵ Increasing AI nativeness — decreasing control & predictability ⟶

🔀 Mixed Architecture

⚠️ Challenges

→ Data fragmented across 5+ system tiers — no single source of truth

→ AI outputs from different systems conflict — no reconciliation layer

→ Integration debt grows faster than AI ROI can offset it

→ Legacy APIs not designed for real-time AI inference patterns

✓ Opportunities

→ Orchestration layer becomes our strategic differentiator

→ RAG bridges legacy data stores without costly migration

→ Middleware modernization pays compound dividends

🏛️ Governance

⚠️ Challenges

→ Who owns the output when 3 AI systems contributed to one decision?

→ Audit trails broken across system boundaries

→ Data classification policies designed for one system era

→ Vendor ToS conflicts create unmanaged IP exposure

✓ Opportunities

→ Cross-system AI governance forces enterprise-wide data maturity

→ Unified policy layer reduces compliance surface area

→ AI accountability frameworks raise the bar for all decisions

♻️ Sustain

⚠️ Challenges

→ Model drift in homegrown AI invisible until a failure event

→ Legacy vendors on 18-month AI release cycles — we fall behind

→ Retraining costs compound as data volumes grow

→ Shadow AI tools cycle in/out faster than governance can track

✓ Opportunities

→ Standardized evaluation cadence across all AI tiers

→ Vendor roadmap alignment as a procurement lever

→ Architecture consolidation becomes justified ROI initiative

Where AI Stands Today — and Where It Is Heading

2022–2023	2024–2025	2026 NOW	2027–2028	2029+
Experimentation	**Early Production**	**Scaling & Proving**	**Agentic Enterprise**	**AI-Native Business**

▲ YOU ARE HERE

From Assistants to Agents

Now: AI answers questions and drafts content on demand

Next: AI agents take multi-step actions autonomously — booking, coding, approving, escalating

⚡ *Every workflow with defined rules becomes automatable. Governance must evolve ahead of deployment.*

From Monolith to Mosaic

Now: One vendor, one model, one platform per use case

Next: Orchestrated multi-model systems — best model per task, assembled at runtime

⚡ *Vendor lock-in risk peaks then resolves. Architecture flexibility becomes a strategic asset.*

From Prompt to Process

Now: Individual users prompt AI for one-off tasks

Next: AI embedded in end-to-end business processes — from intake to outcome, with humans at defined gates

⚡ *ROI shifts from individual productivity gains to process re-engineering returns.*

From Centralized to Distributed

Now: Central IT owns AI platforms and deployments

Next: Business units operate AI with IT providing guardrails, not gates — federated ownership at scale

⚡ *CIO role evolves from builder to platform provider and governance architect.*

From Cloud to Edge & Private

Now: Most AI inference runs on public cloud APIs

Next: Sensitive workloads move on-premise or to private cloud — driven by regulation and IP protection

⚡ *Infrastructure strategy and AI strategy converge. Separate planning cycles are no longer viable.*

From Output to Accountability

Now: AI outputs evaluated informally by users

Next: Regulatory frameworks (EU AI Act, sector rules) require documented accountability for AI decisions

⚡ *Governance is no longer optional. Auditability, explainability, and human oversight become compliance requirements.*

The organizations that lead in 2028 are making architectural and governance decisions today. The window to shape our trajectory is now.

The Enterprise AI Map — All 26 Letters

A Architecture	**B** Build vs Buy	**C** Culture Change	**D** Data Strategy	**E** Ethics & Bias	**F** Foundation Models	**G** Governance
H Human-in-Loop	**I** Integration	**J** Job Redesign	**K** KPIs & Metrics	**L** Legal & Compliance	**M** Model Selection	**N** Network & Infra
O Operating Model	**P** Pilots to Production	**Q** Quality & Evaluation	**R** ROI & Business Case	**S** Security & Risk	**T** Talent & Upskilling	**U** Use Case Selection
V Vendor Strategy	**W** Workforce Readiness	**X** X-Functional Ownership	**Y** Your Data Moat	**Z** Zero Trust AI		

We'll cover these in 6 themed sections — with 3 live polls built in

Six Themes. 26 Decisions. One Enterprise AI Enablement Map.

A – F 🧱 Strategy & Foundation

→ Architecture: cloud, private, or hybrid — decide before we deploy

→ Build vs Buy: build our moat, buy the commodity

→ Culture: leaders must model AI usage or adoption will stall

→ Data Strategy: clean, governed data is our only durable advantage

→ Ethics & Bias: we are accountable for every output our AI produces

→ Foundation Models: evaluate on our data, not vendor benchmarks

G – L ⚙️ Deployment & Operations

→ Governance: approval process, data access matrix, named owners

→ Human-in-Loop: design the review step before go-live, not after an incident

→ Integration: AI in the workflow gets used; AI beside it gets ignored

→ Job Redesign: map every role change before launch, not after backlash

→ KPIs: baseline + AI metric + business outcome — all three, always

→ Legal: EU AI Act and sector regulations apply now

M – R 📊 ROI & Business Value

→ Model Selection: run a bake-off on our own data, not public benchmarks

→ Network & Infra: AI is hungry — audit capacity before we scale

→ Operating Model: hybrid CoE — central standards, decentralized deployment

→ Pilots to Production: a P&L owner makes it a product; IT ownership keeps it a pilot

→ Quality & Evaluation: models drift — build systematic evaluation into OpEx

→ ROI: measure the baseline before we deploy. 30 minutes. Non-negotiable.

S – V 🔒 Security, Risk & Compliance

→ Security: shadow AI audit first — we have 3-5x more exposure than we think

→ Talent: 3-tier literacy (users, builders, governors) — train each differently

→ Use Case Selection: high-volume, low-judgment, high-consequence-for-error first

→ Vendor Strategy: build model-agnostic orchestration — always maintain the ability to leave

W – X 🤝 People & Organization

→ Workforce Readiness: it's a trust problem, not a training problem — address the fear directly

→ X-Functional Ownership: IT enables, business unit owns, steering committee governs

Y – Z 🚀 The Horizon

→ Your Data Moat: models commoditize — our proprietary data doesn't

→ Zero Trust AI: verify AI outputs; apply least-privilege to every AI agent

Strategy & Foundation

The decisions we make before we write a line of code

1

A–F: Strategy & Foundation

The six decisions that determine whether everything else works

A Architecture

Choose early: public cloud AI, private deployment, or hybrid. This decision affects cost, security, latency, and IP protection for years.

B Build vs Buy

Build when the process is our competitive moat. Buy when it's commodity. 80% of enterprise AI should be bought or configured — not built.

C Culture Change

AI fails most often not technically but culturally. Leaders who don't model AI usage, teams that fear replacement, and workflows that resist change kill more initiatives than bad models.

D Data Strategy

Our data is the only thing competitors cannot replicate. Clean it, classify it, govern it, and build access controls before we touch a model.

E Ethics & Bias

Every model reflects the data it was trained on. We are accountable for its outputs — including discriminatory results in hiring, lending, and operations.

F Foundation Models

GPT-4o, Claude, Gemini, Llama — these are our starting points. Evaluate on: accuracy, cost, privacy, context window, and whether the vendor's ToS lets us keep our IP.

LIVE POLL #1 — Strategy Reality Check

Which of these is your biggest strategic gap right now?

A

Data Strategy

We don't have clean, governed, AI-ready data

B

Culture Change

Our people aren't ready or willing to adopt AI

C

Build vs Buy

We're building things we should be buying

D

Governance

We have no formal AI governance or ethics framework

Raise your hand — pick the one that keeps you up at night

Deployment & Operations

How we move from pilot to production to self-funding

2

G–L: Deployment & Operations

The operational decisions that separate pilots from production systems

G Governance

Our AI governance framework must answer: who can approve new AI deployments, what data can AI access, who reviews outputs, and who is accountable when it goes wrong.

H Human-in-Loop

Define before deployment: which AI outputs require human review before action. High-stakes domains — legal, medical, financial, HR — need mandatory review. Design the workflow around it.

I Integration

AI embedded in the workflow gets used. AI bolted on beside the workflow gets ignored. The integration decision is more important than the model selection decision.

J Job Redesign

AI changes jobs — it doesn't just eliminate them. Map every role in our AI deployment: what tasks get augmented, what gets eliminated, what new tasks get created. Do this before go-live.

K KPIs & Metrics

We need three metrics for every AI initiative: a pre-AI baseline, an AI performance metric, and a business outcome metric. All three. Missing one means we cannot prove value.

L Legal & Compliance

GDPR, CCPA, EU AI Act, HIPAA, SOX — AI creates new exposure in every regulated domain. Legal must review every AI deployment that touches customer data, employment decisions, or financial reporting.

ROI & Business Value

How to measure, prove, and compound AI returns

M–R: ROI & Business Value

The framework that turns pilots into board-level proof points

M — Model Selection

Evaluate models on: task accuracy, cost/token at volume, latency, data residency, and fine-tuning options. Run head-to-head evaluations on our own data — not vendor benchmarks.

N — Network & Infra

AI workloads are compute-intensive and latency-sensitive. Audit our network capacity, GPU availability (cloud or on-prem), and data pipeline architecture before scaling any AI system.

O — Operating Model

Decide: centralized AI CoE, federated model with business unit AI leads, or hybrid. Each has trade-offs on speed, governance, and talent. Most large enterprises land on hybrid.

P — Pilots to Production

70% of pilots never reach production. The gap is not technical — it's governance, change management, and business ownership. A pilot that IT owns is a science project. A pilot with a P&L owner is a product.

Q — Quality & Evaluation

AI quality degrades over time as data and usage patterns drift. Build systematic evaluation: regular accuracy testing, human review sampling, and an incident process for when outputs go wrong.

R — ROI & Business Case

Build our ROI case on three pillars: cost avoidance (what manual work was automated), revenue impact (what new capability was enabled), and risk reduction (what exposure was mitigated).

LIVE POLL #2 — ROI Honest Check

Do you have a documented pre-AI baseline for our current AI initiatives?

A

Yes, rigorously

Documented, dated, agreed with Finance

B

Informally

We have a rough sense of it

C

Not really

We know it's better but can't prove it

D

What baseline?

Nobody asked for one when we started

No judgment — this is the most commonly skipped step in enterprise AI

Security, Risk & Compliance

The work that makes AI defensible — and sustainable

4

S–V: Security, Risk & Compliance

The four pillars of enterprise AI trust

S Security & Risk

Shadow AI: Employees are using unapproved AI tools today. Audit what tools are in use before building policy.

Data Classification: Every AI system needs a data tier: what can go to public APIs, what stays on-prem, what never touches AI at all.

Prompt Injection: Attackers can embed instructions in documents or inputs that manipulate our AI agents. Design defenses into agentic systems from day one.

T Talent & Upskilling

AI Literacy Tiers: Not everyone needs to prompt engineer. Define three tiers: AI users, AI builders, AI governance. Train each differently.

Prompt Engineering: This is now a core enterprise skill. Our best-performing employees should be our best prompt engineers. Build the training.

Hiring for AI: New roles needed: AI Product Manager, ML Ops Engineer, AI Ethics Officer, Prompt Engineer. Plan the org chart now.

U Use Case Selection

High-Value Targets: Start with: high-volume, low-judgment, high-consequence-for-error tasks. Document processing, exception handling, first-pass triage.

Avoid These First: Don't start with: customer-facing interactions, HR decisions, anything requiring explainability we can't provide yet.

Use Case Scoring: Score each use case on: data readiness, business impact, implementation complexity, and regulatory risk. Prioritize top-right quadrant.

V Vendor Strategy

Evaluation Criteria: Score every AI vendor on: capability, data privacy, contractual IP protections, exit strategy, and roadmap alignment.

Avoid Lock-in: Build on abstraction layers. Our orchestration code should be model-agnostic. We should be able to swap GPT for Claude without rewriting our application.

Contract Essentials: Required in every AI vendor contract: data ownership clause, no training on our data without consent, SLA for model changes, data deletion rights.

People & Organization

The human system that makes the technical system work

5

W–X: People & Organization

AI is a technology problem until it isn't — then it's a people problem

W Workforce Readiness

Fear is the default:

Most employees assume AI means job cuts. Address this directly and early. Leaders who avoid the conversation amplify the fear.

AI literacy programs:

Every employee needs baseline AI literacy. Build a 3-tier program: Awareness (all staff), Practitioner (power users), Builder (IT and product teams).

Champions network:

Identify 1-2 AI champions in every department. Give them early access, training, and a platform to share wins. Peer credibility drives adoption faster than top-down mandates.

Measure adoption, not just deployment:

A deployed AI tool with 10% adoption is a failure. Build adoption metrics into our success criteria from day one. Track weekly active users, time savings, and satisfaction.

Change management timeline:

Assume 3x longer than our IT delivery timeline for full workforce adoption. Build this into our program plan and executive communication.

X Cross-Functional Ownership

IT is not the owner:

IT is the enabler. The business unit is the owner. Every production AI system needs a named business owner who is accountable for outcomes — not just a technical owner.

The AI Steering Committee:

Establish a cross-functional AI Steering Committee: CIO, CFO, CHRO, General Counsel, and business unit heads. Meets monthly. Approves initiatives. Reviews incidents.

Funding model:

Centralize infrastructure funding. Decentralize use case funding. Business units fund their own AI initiatives with IT providing shared platform, security, and governance.

The 90-day rule:

Any AI initiative that cannot show measurable progress in 90 days should be reviewed for continuation. Not killed — reviewed. With a decision made by the business owner, not IT.

Success sharing:

When AI delivers ROI, make it visible. Share the wins — the specific numbers, the team that delivered them, the process that changed. Celebration drives the next wave of adoption.

LIVE POLL #3 — Organizational Readiness

Where is the biggest people/org gap in your AI program?

A

Workforce Fear

Teams are worried about job displacement

B

No Business Owner

IT owns the AI initiatives, not the business

C

Skills Gap

We don't have AI builders or prompt engineers

D

Leadership

Senior leaders don't model AI usage themselves

Pick the one that most limits our progress — our honest assessment

The Horizon

Our data moat and zero-trust AI — what comes next

Y–Z: The Horizon

The two concepts that define enterprise AI leadership in the next 3 years

Y Your Data Moat

The foundation models are commoditizing. GPT, Claude, Gemini — they are all getting better, faster, and cheaper. In 3 years, the model layer will be nearly free.

What will not commoditize is our proprietary data. Our customer history. Our operational knowledge. Our domain-specific training sets. Our institutional memory.

The organizations that win the AI era will be the ones that invested early in clean, governed, well-labeled proprietary data — and built AI systems that compound on it.

Start now: data labeling programs, knowledge capture, institutional memory systems

Z Zero Trust AI

Zero Trust in cybersecurity means: never trust, always verify. The same principle applies to AI — never assume an AI output is correct, always design verification into the system.

AI Agents & Automation:

As AI moves from generating text to taking actions — executing code, calling APIs, modifying data — the blast radius of an error grows. Apply least-privilege access to every AI agent.

Adversarial Inputs:

Prompt injection, jailbreaking, data poisoning — the attack surface for AI systems is new and evolving. Our security posture must evolve with it.

Auditability:

Every consequential AI decision needs an audit trail: what prompt was used, what data was retrieved, what output was generated, and who reviewed it.

What We Covered — And Why It Connects

Six themes. One system. The decisions that determine where we land on the AI curve.

1 A–F Strategy & Foundation

The decisions before the first line of code.

Architecture, data strategy, culture, and build vs buy — these gates determine everything downstream.

If we skip these, every deployment becomes a workaround.

2 G–L Deployment & Operations

How we close the Valley of Death between pilot and production.

Governance, human-in-loop design, integration depth, and KPI baselines — this is where most AI programs stall.

Without P&L ownership and a pre-AI baseline, we cannot prove value.

3 M–R ROI & Business Value

The receipt the board is waiting for.

Model selection on our own data, hybrid operating model, quality evaluation as OpEx — not afterthought.

The organizations that compound AI returns treat ROI as engineering, not finance.

4 S–V Security, Risk & Compliance

The work that makes everything else defensible.

Shadow AI audit first. Zero-trust architecture. Vendor contracts with teeth. EU AI Act compliance — now.

Security deferred is a liability that compounds faster than AI ROI.

5 W–X People & Organization

AI is a technology problem until it isn't.

Workforce trust — not just training. A steering committee with CFO co-sponsorship. The 90-day accountability rule.

The organizations stalling in the middle are almost always stalling here.

6 Y–Z The Horizon

Our data moat is the only durable advantage.

Models commoditize. Our proprietary data doesn't. Zero-trust AI as AI agents start taking actions, not just generating text.

The window to build the moat is open right now — and it closes as the curve steepens.

The CIO's 180-Day AI Enablement Action Plan

Sequence matters — do these in order

Days 1–60

Assess & Govern

1. Audit shadow AI — what tools are our people using today?
2. Classify our data into AI tiers: public, internal, confidential, restricted
3. Establish AI Steering Committee with cross-functional membership
4. Define our data governance policy for AI workloads
5. Identify our top 5 use case candidates using a scoring matrix

Days 61–120

Build Foundation

1. Select our foundation model platform — run a bake-off on our actual data
2. Define our AI architecture: cloud tier, private tier, data access controls
3. Launch AI literacy program — Tier 1 (awareness) for all staff
4. Deploy AI champions network — identify and activate in each department
5. Start our first use case — pick the one with the clearest baseline metric

Days 120–180

Deliver & Measure

1. Go live with first use case — measure against pre-AI baseline on day 1
2. Present first ROI data point to CFO — actual numbers, not projections
3. Launch Tier 2 AI literacy training for power users and builders
4. Review AI vendor contracts — add data ownership and IP clauses
5. Publish 90-day results to leadership — visible wins drive the next wave

The Complete A–Z Reference

Take a photo — our enterprise AI enablement map

A Architecture — decide early: cloud, private, or hybrid

B Build vs Buy — build our moat, buy the commodity

C Culture Change — the people problem is harder than the tech problem

D Data Strategy — our data is the only thing competitors can't copy

E Ethics & Bias — we are accountable for our AI's outputs

F Foundation Models — evaluate on our data, not vendor benchmarks

G Governance — who approves, who owns, who is accountable when it fails

H Human-in-Loop — design the review step before deployment, not after an incident

I Integration — AI in the workflow beats AI beside the workflow every time

J Job Redesign — map every role before go-live, not after the backlash

K KPIs — baseline, AI metric, and business outcome. All three. Missing one = no proof

L Legal & Compliance — EU AI Act, GDPR, and sector regulations apply now

M Model Selection — run a bake-off on our actual data, not the benchmark leaderboard

N Network & Infra — AI is hungry; audit capacity before we scale

O Operating Model — hybrid CoE: central standards, decentralized deployment

P Pilots to Production — a pilot IT owns is a science project; a P&L owner makes it a product

Q Quality & Evaluation — AI drifts; build systematic evaluation into operating costs

R ROI — measure baseline before deployment; receipts beat projections every time

S Security & Risk — shadow AI audit first, policy second, controls third

T Talent — 3 tiers: users, builders, governors. Train all three differently

U Use Case Selection — high volume, low judgment, high consequence for error

V Vendor Strategy — always maintain the ability to leave

W Workforce Readiness — address fear with specifics, not reassurance

X X-Functional Ownership — IT enables, business unit owns, steering committee governs

Y Your Data Moat — models commoditize; our proprietary data doesn't

Z Zero Trust AI — verify AI outputs; apply least-privilege to AI agents

"AI is not a technology initiative.
It is a business transformation initiative
that happens to be enabled by technology."

Thank you.

GS Jha | CIO & CISO, QuantumScape

Let's continue the conversation • LinkedIn: /in/gsjha

Open Discussion | Q&A

What is the one letter on the A-Z map we are going to act on first?