The Emerging Pillars of AI for Business in 2026

Practical Patterns of Success — Where Businesses and Builders Are Finding Real Traction

Version: 1.0
Date: May 2026 | Initial Release
Status: Public Release — Free to share with attribution
Author: Paul Hollen — Founder, RAG9 / Framework Business Solutions LLC
Publisher: RAG9.com | RAG9Labs.com | paul@rag9.com

Executive Summary

2025 was the year the enterprise AI experiment ran headlong into reality. With the four major hyperscalers — Alphabet, Amazon, Meta, and Microsoft — collectively spending approximately **$400 billion** on AI infrastructure, the expectation of transformational returns was enormous. The actual outcomes told a different story.

MIT Media Lab's Project NANDA published its landmark *"GenAI Divide: State of AI in Business 2025"* report in July 2025, drawing on 150 executive interviews, 350 employee surveys, and analysis of 300 public AI deployments. The headline finding was stark: **95% of enterprise generative AI pilots delivered no measurable P&L impact.** The enterprise AI market had not failed to innovate. It had largely failed to integrate.

But embedded in that same report was a more important finding. The **5% that succeeded were not lucky — they were deliberate.** Narrowly focused. Integrated from day one. Human oversight maintained. Outcomes measured in dollars and hours, not demos and slide decks.

Those successes clustered. From that clustering, four clear patterns emerged — patterns now demonstrating real traction, measurable investment, and commercial momentum in 2026. We call them the **Emerging Pillars of AI for Business.**

~$31B

AI cybersecurity market in 2025 (Mordor Intelligence)

~$7B

Agentic AI market in 2025, growing at 40–44% CAGR

74%

Companies planning agentic AI deployment within two years (Deloitte, 2026)

62%

Organizations experimenting with AI agents (McKinsey, 2025)

21%

Companies with mature governance for autonomous agents (Deloitte, 2026)

95%

Enterprise AI pilots with no measurable P&L; impact (MIT NANDA, 2025)

The Four Pillars — Ranked by Combined Score

Pillar	Speed	Growth	Universality	Capital	Total / 20
Security	5	5	5	4	**19 / 20**

Agentic Systems & Data Processing	5	5	4	4	**18 / 20**
AI Governance	4	5	4	5	**18 / 20**
Communications	4	4	5	4	**17 / 20**

1 | Introduction — From 2025 Experimentation to 2026 Maturity

Companies do not build AI. They build value using AI. That distinction matters more than most boardroom conversations acknowledge.

In 2025, the gap between those two ideas cost the enterprise market an extraordinary amount of time, capital, and credibility. The hyperscalers poured approximately $400 billion into AI infrastructure. Enterprise organizations poured billions more into pilots, tools, and transformation programs. And 95% of those enterprise programs produced no measurable return.

The MIT NANDA report was not a condemnation of AI. It was a diagnosis of integration failure. The technology worked. The organizational approach largely did not. Teams built broad when they should have gone narrow. They experimented when they should have integrated. They automated surface tasks when the real value was buried in back-office friction. They bought generic tools when the ROI lived in workflow-specific implementation.

> *"Purchasing AI tools from specialized vendors and building partnerships succeed about 67% of the time, while internal builds succeed only one-third as often."*
>
> — MIT NANDA GenAI Divide Report, July 2025

The 5% that broke through taught us something valuable. Not by accident, but by pattern. They chose specific problems over broad ambitions. They treated security and governance as prerequisites, not afterthoughts. They measured outcomes in dollars and hours, not demos and slide decks.

From those patterns, four areas of genuine traction have emerged in 2026. These are not predictions. They are observations — backed by market data, confirmed by independent research, and visible in the capital allocation decisions of organizations that have learned from 2025's expensive lessons.

2 | How We Evaluated the Pillars

Each pillar was evaluated using four criteria scored from 1 to 5, for a maximum of 20 points. The criteria were selected to reflect real-world commercial viability rather than theoretical potential.

Criterion	Measurement Approach
Speed of Evolution	Rate of capability improvement and transition from pilot to production deployment
Growth Rate	CAGR and absolute market size momentum, sourced from independent research firms

Universality	Applicability across business sizes, industries, and geographies
Capital Priority	Level of actual committed spending and urgency of investment decisions

Research Sources

Findings were cross-referenced against published research from McKinsey, Deloitte, Accenture, BCG, and independent market intelligence firms including Mordor Intelligence, Fortune Business Insights, and Grand View Research. All market size figures reflect 2025 published estimates. All CAGR figures reflect published forward projections. Any figure that could not be independently sourced was removed.

> Source discipline note: Every numeric claim in this paper is traceable to a named, dated source. Cross-referenced against McKinsey, Deloitte, Accenture, BCG, and independent market intelligence firms. If citing downstream, cite the original source — not this paper — for statistical claims.

3 | The Emerging Pillars of AI for Business in 2026

Pillar 1: Security 19 / 20

AI-powered defense systems that address both traditional cyber threats and the new attack surfaces created by AI itself.

Why This Pillar Emerged

Security earned the highest overall score for a simple reason: it is non-negotiable. Every organization adopting AI must address it. You cannot opt out. You cannot defer it. The cost of underinvesting in security is not a missed opportunity — it is an active liability.

The Threat Landscape Just Changed Permanently

On April 7, 2026, Anthropic announced **Project Glasswing** and its associated frontier model, **Claude Mythos Preview** — publicly confirming that AI had crossed a threshold in offensive security capability the industry had been anticipating and quietly dreading.

Mythos Preview, a general-purpose model not built specifically for security, autonomously identified thousands of previously unknown zero-day vulnerabilities across every major operating system and every major web browser — including flaws that had survived decades of human security review and millions of automated tests. In one documented case, it uncovered a 27-year-old vulnerability in OpenBSD. It developed working exploits on the first attempt in over 83% of cases.

Anthropic restricted Mythos from public release, forming Project Glasswing — a defensive consortium including Microsoft, Apple, Google, Amazon Web Services, JPMorgan Chase, NVIDIA, Cisco, and CrowdStrike — to find and patch vulnerabilities before threat actors could exploit them.

> The implication is unambiguous: if a general-purpose AI model can autonomously find and exploit zero-days at this scale, the security posture of any organization not actively using AI in its defense is measurably weaker than it was one year ago.

Also in May 2026, the **TanStack npm supply chain attack** — dubbed "Mini Shai-Hulud" — compromised 42 TanStack packages, Mistral AI's SDK suite, UiPath tooling, and over 170 npm and PyPI packages in a single coordinated campaign. OpenAI's own developer environment was among those impacted.

Key Technologies

- Behavioral anomaly detection — learns normal system patterns and identifies deviations in real time
- Automated incident response and threat containment
- AI-specific defenses against prompt injection, data poisoning, and model theft
- Predictive risk analysis and autonomous vulnerability discovery
- Supply chain integrity verification and dependency provenance tracking

Market Momentum

The AI cybersecurity market reached approximately **$30–35 billion in 2025**, growing at a **22–24% CAGR** (Mordor Intelligence, Grand View Research). This spending is largely non-discretionary — a baseline cost of AI adoption.

> Security is not one of four pillars. It is the foundation the other three rest on. You cannot scale agentic systems without securing them. Security earns the top score because it is the prerequisite for everything else.

Pillar 2: Agentic Systems & Data Processing 18 / 20

Autonomous agents that move beyond generating insights to delivering action — reasoning, planning, and executing across real business workflows.

Why This Pillar Emerged

One of the clearest lessons from 2025 was that generating insights is not the same as creating value. The projects that delivered measurable ROI moved from passive tools — tools that answer questions — to active systems that complete tasks. This is the shift from AI that informs to AI that acts.

What the Research Shows

McKinsey's *State of AI in 2025* survey, conducted across 1,993 participants in 105 countries, found that **62% of organizations are at least experimenting with AI agents.** In no single business function does more than 10% report having fully scaled agentic deployment. The technology is ready. The organizational infrastructure is still being built.

Deloitte's *2026 State of AI in the Enterprise* — surveying 3,235 senior leaders across 24 countries — found that **74% of companies plan to use agentic AI at least moderately within two years.** Only 23% are doing so today.

> *"Agentic AI is poised for growth — close to three-quarters of companies planning to deploy it within two years. Yet only 21% report having a mature model for agent governance."*
>
> — Deloitte State of AI in the Enterprise, 2026

Real-World Success Patterns

- Bank of America deployed targeted agentic workflows dramatically reducing analyst research time
- Lumen Technologies compressed multi-step processes from hours to minutes
- EchoStar saved tens of thousands of work hours through production agentic applications
- Startups building narrowly-focused agentic tools on single pain points hit $0 to $20M ARR within a year (MIT NANDA)

Market Momentum

The agentic AI market was valued at approximately **$7–7.5 billion in 2025**, projected to $9–10.9 billion in 2026, with CAGRs of **40–44%** through 2030 (Mordor Intelligence, Fortune Business Insights, Precedence Research). This is the fastest-growing segment in the enterprise AI landscape.

> Builder Opportunity: The highest-value opportunity is vertical specificity. Generic agents are commoditizing rapidly. Agents with deep domain knowledge, pre-built integrations to industry-specific systems, and clear ROI metrics are building durable differentiation.

Pillar 3: AI Governance 18 / 20

The tools, frameworks, and organizational practices that build trust, ensure compliance, and make AI outputs reliably actionable — not just impressively fluent.

Why This Pillar Emerged

Today's AI systems all share one characteristic: they sound confident. A hallucinated answer delivered with the same tone and fluency as a correct one is indistinguishable to the reader — until a business decision made on that answer produces a consequence. Governance is what separates AI tools useful for drafting emails from AI systems trusted with consequential decisions.

> The governance gap: Deloitte's 2026 report found that only 21% of companies have a mature governance model for autonomous agents — while 74% plan to deploy them within two years. That gap is not a planning oversight. It is a ticking liability.

What the Research Shows

The Deloitte report identified data privacy and security (73%), legal and IP compliance (50%), and governance capabilities and oversight (46%) as top AI risk concerns. Regulatory pressure is accelerating the timeline: EU AI Act enforcement began February 2025, with fines up to EUR 35 million or 7% of global turnover for non-compliance. Insurance carriers now link premium discounts to certified governance

frameworks.

Market Momentum

The AI governance market is early-stage but growing at **28–51% CAGR** through 2030 (Mordor Intelligence, MarketsandMarkets). This market is in its early exponential phase — under-built relative to the urgency of the problem.

> Governance is what turns AI from a productivity tool into a business system. Organizations that build governance in parallel with AI deployments — rather than retrofitting after an incident — will scale faster, with less risk, and with greater board confidence.

Pillar 4: Communications 17 / 20

Modern conversational AI systems replacing the most universally hated interface in business — the legacy IVR — with natural, intelligent, action-capable experiences.

Why This Pillar Emerged

We have all been there. You call your bank. A robotic voice greets you with seven options, none matching your problem. That interface has persisted for decades because the cost of replacing it exceeded the political will to do so. That calculus has changed.

Gartner projects conversational AI will **reduce contact center labor costs by $80 billion in 2026.** That is the kind of number that ends budget debates.

Real-World Success Patterns

Heathrow Airport's Hallie AI is the most cited success story in this space — a natural-language voice agent integrated with airport operations that handles real traveler queries at scale without traditional IVR frustration. ServiceNow and EY implementations demonstrate that conversational AI connected to enterprise systems can handle a significant percentage of inbound requests without human escalation.

Market Momentum

The conversational AI market was valued at approximately **$13.6–19.2 billion in 2025**, projected to $17–21 billion in 2026. CAGR projections range **20–26% through 2030** (Research and Markets, Fortune Business Insights). This pillar scored highest on Universality — every business with a customer-facing communication channel has a use case.

> Conversational AI is the most accessible on-ramp to AI adoption. The ROI is visible, the comparison point (the existing experience) is poor, and the implementation path is well-defined. The most practical first step for organizations that haven't yet found their AI footing.

4 | Runner-Ups & Honorable Mentions

Three additional areas scored well but fell short of the core four on universality, capital priority, or market maturity.

Area	Why It Matters	Score
AI Code Review & Verification	Governance layer for AI-generated code — enabling safe scale of developer productivity	**16 / 20**
Personal Assistants & Desktop Integration	High individual productivity gains; lower capital priority at enterprise level	**15 / 20**
Physical AI & Intelligent Automation	High sectoral growth in manufacturing, logistics, healthcare; lower universality	**14 / 20**

AI Code Review, Verification & Trust Tools — 16/20

This is one of the most practically urgent problems in the enterprise AI landscape. AI-assisted coding tools are now used by the majority of professional developers. The productivity gains are real. So is the risk: AI-generated code can be subtly flawed, architecturally unsound, or contain security vulnerabilities that pass automated testing but fail in production.

Specialized tools that provide architectural review, security scanning, and quality scoring tailored to LLM-generated output are addressing a gap that will only grow as AI code generation scales. This area scored lower on Universality — it is deeply relevant for software-heavy organizations but less applicable to businesses with minimal development operations. McKinsey and BCG both identify AI-generated code governance as a critical emerging need for organizations scaling developer productivity programs.

Personal Assistants & Desktop Integration — 15/20

Deeply embedded AI copilots in everyday productivity tools — Microsoft Copilot, Google Gemini integrations, and emerging independent alternatives — are showing strong individual adoption. The productivity gains for knowledge workers are well-documented.

This area scored lower on Capital Priority at the enterprise level, where many organizations view these as individual productivity tools rather than foundational infrastructure investments. That perception is beginning to shift.

Physical AI & Intelligent Automation — 14/20

Robotics, cobots, drones, autonomous vehicles, and AI-powered predictive maintenance represent the highest growth potential in specific sectors — manufacturing, logistics, healthcare, and defense.

Advances in surgical robotics, security automation, and industrial AI — particularly work powered by platforms like NVIDIA Cosmos — make this a vertical to watch closely. It scored lower on Universality because not every business has meaningful physical operations. For those that do, it may be the most strategically important pillar of all.

5 | How the Pillars Connect — and Why That Matters

These four pillars are not independent investment categories. They form a reinforcing system. Understanding their interdependencies is what separates organizations that build AI strategically from those that accumulate AI tools.

From	Connection
Communications	Generates rich real-time interaction data that becomes fuel for Agentic Systems. Every conversation is a training signal and a workflow trigger.
Agentic Systems	The intelligent core — turning communications into action, strengthening security through predictive monitoring, enforcing governance policies in real time.
Security	The foundation that protects all other pillars. Secures communication channels, validates agent behavior, and protects governance data from manipulation.
AI Governance	The trust layer across everything. Without governance, security has no audit trail, agents have no behavioral boundaries, and communications have no accountability.

Practical Implementation Guidance

- Start with one high-impact use case in one pillar rather than attempting broad transformation
- Prioritize integration with existing systems from day one — AI that cannot talk to your data is not useful
- Build security and governance in parallel with deployment, not as an afterthought
- Measure outcomes in business terms: time saved, resolution rates, cost reduction, revenue influence
- Treat AI as an augmentation layer that expands human capability — not a replacement layer that removes human judgment

Career and Skill Opportunities

According to LinkedIn's 2026 data, AI-related roles are among the fastest-growing job categories globally. The specific emerging titles tracking directly to these four pillars include:

- AI Agent Architect / Agentic AI Engineer
- AI Security & Risk Analyst
- AI Governance & Ethics Lead
- Conversational AI Designer / Voice Experience Architect
- MLOps / Agent Orchestration Engineer

McKinsey, Deloitte, and Accenture all report that professionals who combine deep domain expertise with AI-adjacent skills are seeing faster career progression and stronger compensation outcomes than those pursuing either dimension alone.

6 | Conclusion

2025 will be remembered as the year enterprise AI ran its most expensive experiment. Approximately $400 billion in hyperscaler infrastructure investment. Billions more in enterprise pilots. A 95% failure rate that generated extraordinary lessons.

The conversation about AI failure has been useful. But it is incomplete. The more important story is in the 5% — in the specific patterns of success that are now visible, repeatable, and investable.

The Emerging Pillars of AI for Business are not predictions about where AI might go. They are observations about where genuine traction already exists. Security, because AI adoption without AI defense is a liability. Agentic Systems, because insights without action are just expensive research. AI Governance, because confidence without accountability is a liability dressed as a feature. Communications, because every business talks to someone, and the technology to do that better has finally arrived.

For business leaders, these pillars are a practical orientation guide. The question is not whether to invest in AI. The question is where to concentrate effort to produce measurable returns. Start narrow. Integrate deeply. Measure honestly. Build security and governance in from the start. That is what the successful 5% did. It is not complicated — but it requires discipline that broad transformation programs rarely sustain.

For builders, these pillars are a commercial map. The platform layer is commoditizing rapidly. The integration layer — domain-specific agents, industry-tailored governance tools, vertical conversational systems, AI-native security tooling — is where durable value is being created right now.

And for anyone wondering whether the AI era is an opportunity or a threat: the research and real-world patterns show something genuinely encouraging. This is not a zero-sum game. Scientists and mathematicians continue pushing the boundaries. Teachers and trainers are in high demand. Marketers, creatives, and storytellers are finding powerful new tools. New roles are appearing every month. The companies and individuals who focus on these pillars — with discipline, integration, and honest measurement — will be best positioned to create lasting value.

> **In 2026, success belongs to those who integrate thoughtfully rather than transform chaotically. The future is not coming to replace us. It is inviting us to build it.**

Key Takeaways

- Security is the prerequisite pillar — non-negotiable, non-deferrable
- Agentic systems are the fastest-growing segment with the clearest ROI pattern
- The governance gap between deployment plans and actual oversight is 2026's biggest enterprise liability
- Conversational AI is the most accessible on-ramp for organizations starting their AI journey
- The 5% that succeeded in 2025 were deliberate, narrow, integrated, and measured — that playbook scales

Author's Note — On the Methodology of This Paper

This paper was not produced by a single AI prompt. It was built through a multi-model, human-directed research and editorial process — and that process is itself an example of the principles it describes.

The thesis came first: from reading the MIT NANDA report, cross-referencing it with Deloitte, McKinsey, Accenture, and BCG research, and identifying a pattern in where the 5% of successful AI deployments were actually clustering. The four pillars emerged from that analysis — not reverse-engineered from a desired conclusion.

Multiple AI systems — including Claude (Anthropic), Grok (xAI), and ChatGPT (OpenAI) — were used as research assistants, editorial consultants, and verification tools across different stages. Every market statistic, named report, company reference, and specific event cited was independently verified against primary or high-credibility secondary sources. Claims that could not be sourced were removed.

The human in this loop — the architect, the editorial authority, and the person responsible for every judgment call — is Paul Hollen. The AI systems contributed capability. The perspective, the framework, the sourcing discipline, and the final word on every sentence belong to a person with thirty years of data engineering experience and skin in the game.

> This paper is its own proof of concept. Human-directed, multi-model AI collaboration — with rigorous verification and clear human accountability — is precisely the pattern MIT research identified as the differentiator for the successful 5%. We tried to model that here.

If you find an error, a number that does not trace to a source, or a claim that does not hold up to scrutiny — that is the human's responsibility, not the AI's. Reach out: paul@rag9.com

About the Author

Paul Hollen is the founder of RAG9.com, Framework Business Solutions LLC, RAG9Labs.com, and VitaCraftAI.com. He brings thirty-plus years of data engineering and architecture experience to the applied AI space — spanning work in the Navy, electrical systems, and three decades building data infrastructure for businesses of every size.

He is currently building RAG9-Adam-Prime, a multi-tenant AI orchestration platform designed around the principles described in this paper: security-first, governance-integrated, human-in-the-loop, and built for real-world business application rather than demonstration.

He writes and publishes under the builder identity **DataBob** — because the best way to understand what AI can do for business is to build it, break it, and build it again.

rag9.com | rag9labs.com | paul@rag9.com

Key Sources

- MIT NANDA "GenAI Divide: State of AI in Business 2025" — July 2025
- McKinsey "The State of AI in 2025: Agents, Innovation, and Transformation" — November 2025
- Deloitte "State of AI in the Enterprise" — January 2026

- Mordor Intelligence, Grand View Research, Fortune Business Insights, Precedence Research — market sizing 2025–2026
- Anthropic Project Glasswing announcement — April 2026
- TanStack / Mini Shai-Hulud supply chain attack reporting — May 2026