Inside GenAI

n° 90

GENAI TRENDS 2026:
Industry Overview and Implications

CRIF
Together to the next level

Table Of Contents

Executive Summary

As 2026 approaches, generative AI (GenAI) is at a critical turning point: organizations are moving beyond pilots, but most still struggle to achieve scalable, measurable impact. While nearly 80% of companies have experimented with GenAI, only a small fraction report tangible business value. This highlights the urgent need to bridge the gap between innovation and execution.

Regulatory frameworks, particularly the AI Act and Digital Operational Resilience Act (DORA in Europe, are now actively shaping the industry, making compliance, risk management, and transparency non-negotiable for high-risk AI applications. Success in 2026 will not only require technical advancement, but also process harmonization, robust governance, and targeted upskilling of teams to manage AI responsibly and at scale.

Three primary trends are defining the landscape:

- **Agentic AI** is evolving from conversational chatbots to autonomous agents capable of executing complex, multistep tasks—transforming workflows and decision-making across sectors.
- **Domain-specific and small language models (DSLMs and SLMs)** are gaining traction, enabling privacy-preserving, tailored solutions that meet stringent regulatory demands and deliver scalable intelligence.
- **Advanced anti-fraud defenses:** There is a critical need for defenses against deepfakes and synthetic documents amid the escalation of AI-driven fraud in onboarding, Know Your Customer (KYC), Anti-Money Laundering (AML), and collection processes.

CRIF is uniquely positioned to govern and guide these trends. It's not just about adopting more powerful models but also unlocking the value of existing information assets and integrating them into secure, auditable, and compliant decision engines powered by AI.

For the financial sector, this means turning risk into a competitive advantage. Organizations that successfully operationalize responsible AI and standardize their decision-making processes will unlock both operational efficiency and lasting strategic advantages.

The rise of agentic commerce has driven the emergence of new transaction protocols and standards, presenting both opportunities and risks for financial services, telcos, and utilities. Organizations must adapt to AI-driven transactions, embedding trust, identity, and compliance at the protocol level.

1. Context:
Innovation Landscape & Driving Factors

The End of the "Pilot Era"

2025 marks the end of GenAI's pure experimentation phase.

Companies are rapidly moving from experimentation to actual deployment—but the results are mixed. According to BCG, by late 2025, an estimated 78% of companies had introduced GenAI in at least one business area, yet ~80% reported no material business impact so far. Only a small minority (~5%) of firms—the so-called "future-built" AI leaders—have realized significant value at scale from AI initiatives. Most are stuck in "proof-of-concept purgatory"—testing, iterating, but not scaling.

The challenge for 2026 is to turn experiments into execution and measurable results.

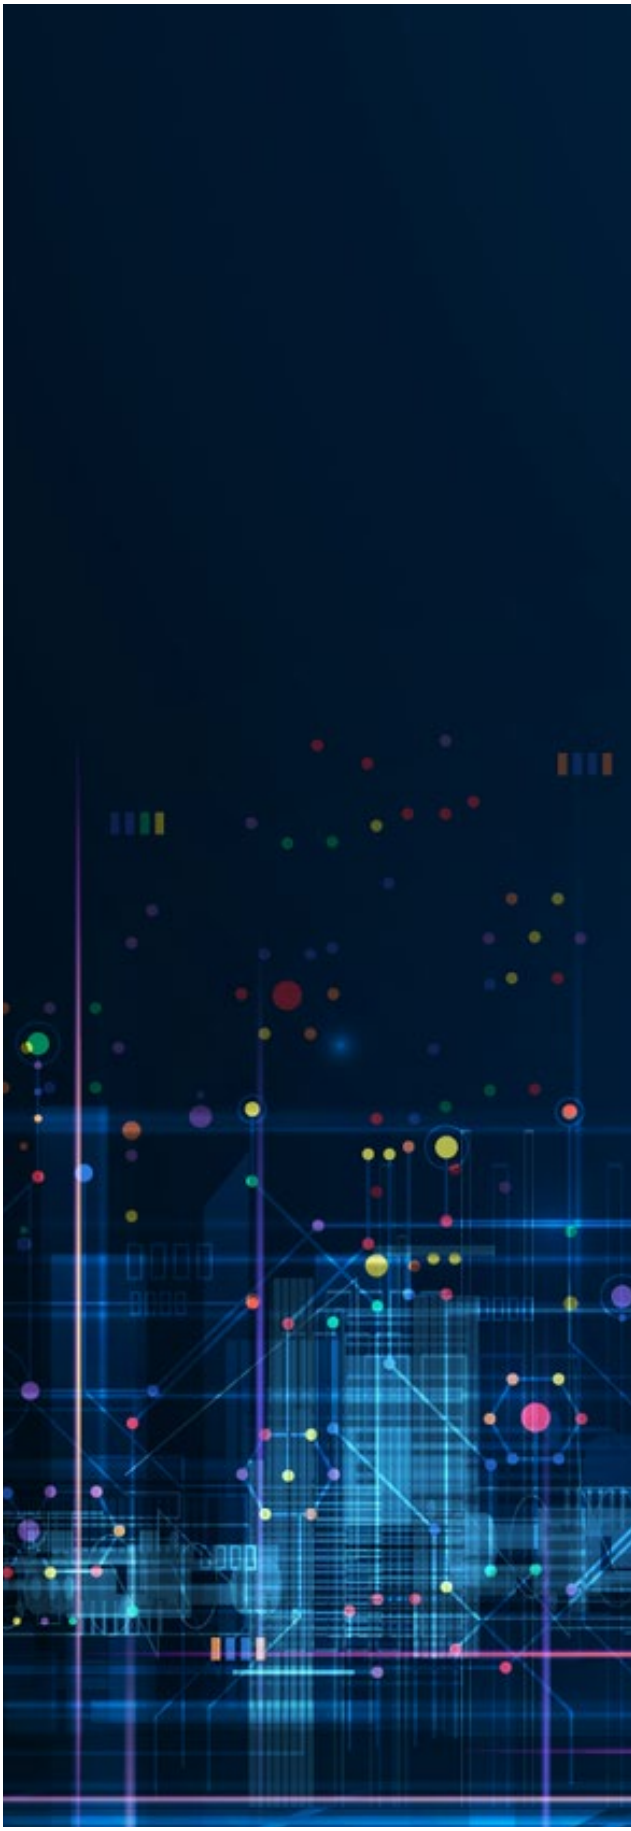

What's Driving (and Blocking) Progress?

- **Technological innovation:** Models are becoming more powerful, cheaper, and more accessible. Massive investments are fueling the next wave, but implementation complexity is growing.
- **Competitive pressure (fear of missing out):** GenAI is now a boardroom topic, not just an IT experiment. The perceived risk is no longer "investing too much", but "being left behind".
- **The human and process factors:** Technology is racing ahead, but people and processes are lagging behind. There's a real skills gap: most teams lack the expertise to move from pilot to production, or to manage AI risks at scale.

Regulation as the Blueprint for Responsible Innovation

Regulation deserves a separate mention as a driver of innovation—shaping corporate strategies and profoundly influencing the EU GenAI landscape.

As GenAI matures and starts delivering real value, regulators are moving from a wait-and-see approach to active enforcement. The EU AI Act and DORA are pushing firms to formalize AI use. These laws require risk management, transparency, and controls for high-risk areas such as lending and identity verification, requiring immediate investments in governance and compliance. The AI Act and DORA are reshaping boardroom agendas, forcing organizations to treat responsible AI not as a checkbox exercise but as a core business strategy.

Here's what that means in practice:

- **Operationalizing compliance:** The AI Act and DORA will require robust, dynamic frameworks for risk management, transparency, and quality control—especially in high-stakes domains.
- **Responsible AI as a market differentiator:** Firms that embed explainability, fairness, privacy, and human oversight into their AI systems will win trust and unlock new opportunities.
- **Privacy-preserving, domain-specific models:** SLMs and DSLMs will become the backbone of compliant, high-value AI—enabling secure, tailored solutions that meet both business and regulatory requirements.
- **Continuous governance:** Model documentation, lineage tracking, bias testing, and post-market monitoring will be routine, supported by dedicated teams and advanced tooling.

Agility will not be found in avoiding compliance but in using regulatory clarity to move faster, turning responsible AI into a growth engine.

Why This Matters for Businesses

The winners in 2026 will be those who close the gap between GenAI hype and business value. Success won't be measured by the number of POCs but by the ability to scale what works and integrate AI into critical decisioning workflows (e.g., credit origination, fraud management, etc.).

Regulatory readiness isn't optional. Clients need partners who deliver AI that's not just powerful, but provably safe, ethical and compliant. Upskilling teams and redesigning processes are just as important as the technology itself.

2. Emerging Technology Trends in GenAI (2026)

Trend 1: Agentic AI

GenAI chatbots are evolving from passive responders to proactive agents capable of taking action to achieve specific goals. In 2025, we saw the first agent modes capable of performing autonomous, multistep tasks (e.g., software development, reporting, and deep market and web research, etc.). At CRIF we launched the industry's first AI agent for the business information sector, transforming access to and analysis of business data.

Powered by large language models, these agents can interpret user requests, use tools for real-time information, reason to enhance output, retain information, and delegate tasks to other agents. However, they still have limitations such as reliability issues and dependency on advancements in the underlying language models.

2026 Outlook: AI agents will dramatically improve memory and context management. However, their integration with structured knowledge (to ensure accuracy) and robust safeguards (to ensure compliance) will be fundamental. Their true value will come from integration with enterprise systems, not from operating as standalone tools.

Agentic Commerce – The Next Transaction Stack

2025 marked the rise of agentic commerce: AI agents that autonomously browse websites, compare options, and complete transactions with minimal human input. These agents act on behalf of individuals and businesses, making shopping faster, personalized, and frictionless. This represents a new frontier in the GenAI landscape, with implications still unfolding. Unlike other trends analyzed in this document—whose impact is better understood and more predictable—agentic commerce introduces dynamics that remain largely unexplored, warranting closer attention.

Early signals are already visible:

- **AI as the first touchpoint:** ChatGPT, Google Gemini, and emerging tools such as OpenAI Atlas and Perplexity Comet are moving beyond searches toward autonomous shopping and checkout.
- **Emerging infrastructure:** Standards such as OpenAI's ACP, Google's AP2, and Visa's Trusted Agent are gaining traction among major retail, e-commerce, and payment players.

This matters because agentic commerce introduces a new transaction layer alongside traditional systems:

- AI agents will own the shopping journey—businesses must stay visible and trusted while managing new fraud risks.
- Banks and Payment Service Providers (PSPs) will have to deal with AI-to-AI commerce, requiring trust, identity, and compliance at the protocol level—or risk losing control and relevance.
- The main threat is disintermediation: For example, a user asks their phone's AI assistant to find the best new phone plan. The AI assistant recommends a plan and offers an embedded "Pay in 3" option from the OS provider (e.g., Apple or Google), completely bypassing the telco's own financing partners and the customer's bank.

2026 Outlook: An increasing risk of disintermediation makes protocol-level integration vital as user-facing agents bypass traditional channels. Transaction data flowing through AI platforms will redefine credit and fraud models, while delegated AI actions introduce new liability challenges. Identity governance must evolve to link agent identity with Know Your Customer (KYC) and Anti Money Laundering (AML) rules.

However, user adoption remains uncertain—look out for signals of real-world uptake. Success will depend on balancing innovation with trust, defining opportunities and risks for 2026 and beyond.

Trend 2: Privacy-Preserving & Domain-Specific Models

As GenAI usage grows, privacy and specialization are paramount. Domain-specific language models (DSLMs) offer precision, privacy, and compliance. CRIF's GenAI Factory is already deploying small language models (SLMs) trained on specific domains to ensure high relevance, lower cost, and no risk of sensitive data leakage.

2026 Outlook: The trend is toward models that keep data on-premises or within secure enclaves. We will see a proliferation of "small but expert" DSLMs trained on specific industry jargon and/or processes. These models won't replace large models but will augment them for critical tasks, harnessing the power of the company's proprietary data as a true competitive advantage.

Trend 3: Anti-Tampering as a Core Risk Component

Defenses against deepfakes and synthetic documents (e.g., payslips, ID documents, bank statements, …) are no longer optional but rather are becoming an integral part of onboarding, Know Your Customer (KYC), Anti-Money Laundering (AML), and collection processes.

The fraud landscape is undergoing a major transformation . Financial institutions are now facing AI-powered fraud such as multilingual text generation, deepfake-as-a-service offerings, synthetic identities, real-time voice cloning, and document tampering of payslips, ID documents, and bank statements.

2026 Outlook: Deepfake-driven social engineering is expected to become routine, with real-time audio and video impersonation as well as content injection during Know Your Customer (KYC) processes and account recovery emerging as standard tactics. Contact centers and video meeting workflows will turn into new battlegrounds. Companies will increasingly deploy AI tools to counter AI-powered fraud. Document forensics will become essential, pushing institutions to integrate provenance (chain-of-custody) and anti-tampering checks before making any decisions.

The Approaching Wave of AI-Driven Fraud

The fraud landscape has evolved into an industrialized, AI-powered ecosystem—scalable, automated, and alarmingly effective. Financial services operations are reporting a sharp rise in credential-stuffing attacks, hybrid synthetic identities, and deepfake voice impersonation during account recovery calls. Fake documents combined with synthetic data are on the increase, while Fraud-as-a-Service operators are professionalizing, selling verified bank accounts powered by deepfake technology.

2026 Outlook: Combatting increasingly sophisticated attacks will require a comprehensive, layered defense—not just better classifiers. Stronger digital identity frameworks, behavioral and device intelligence, document verification and provenance, and telecom or consortium-based signals will be essential to harden defenses.

3. From Trends to Tangible Impact: The CRIF Vision

Emerging technologies must translate into measurable impacts on core processes:

Agentic AI
→ Augmented Decision Support

In credit risk management, an AI agent can accelerate decision-making by analyzing complex datasets (structured and unstructured) and simulating multiple scenarios in real time. This isn't about replacing human judgment—it's about augmenting it with precision and speed, enabling institutions to respond to market shifts faster than ever.

Trusted Data
→ The Foundation for Reliable AI Agents

AI-driven decisions are only as strong as the integrity of the data and the credibility of the providers behind them. In financial services, trust isn't negotiable—it's built into the system. Institutions must ensure that every dataset feeding an AI agent is verifiable, and every provider adheres to rigorous compliance and security standards. This trust layer extends to the agents themselves: when agents act autonomously, stakeholders need confidence that their outputs are accurate, unbiased, and auditable. Building this chain of trust—from data to provider to agent—is what transforms AI from a tool into a dependable partner in critical decision-making.

Small Language Models
→ Scalable and Compliant Intelligence

Lightweight models make AI accessible by reducing costs and latency. For financial institutions, SLMs unlock conversational AI at scale. Imagine a mortgage advisor's desktop assistant—running an SLM locally on their machine or on the company's self-hosted infrastructure—providing real-time script compliance suggestions (e.g., "You forgot to mention the cooling-off period") without any sensitive customer data leaving the bank's secure perimeter. This democratizes AI, making it practical for everyday workflows.

Deepfake Detection & Anti-Tampering
→ Trust as a Differentiator

As fraud tactics evolve, safeguarding authenticity becomes critical. Trust will be the primary competitive differentiator. Institutions that invest early in these safeguards will not only protect themselves but also strengthen customer confidence.

Standardization
→ The Prerequisite for High-Value AI

AI thrives on structured, context-rich data. Beyond harmonizing workflows, organizations must redesign data to be adaptive and dynamically contextualized, enabling AI agents to act with precision in changing environments. The goal isn't to replicate human behavior—it's to empower agents through flexible, context-aware processes.

Customer Experience
→ Efficiency That Becomes Service

A more efficient risk process isn't just a cost saving; it's a better service. For example, a telco customer struggling to pay a bill can interact with a GenAI agent that doesn't just read a script. It analyzes their payment history, current usage, and available plans, then proactively offers a personalized one-week payment extension or a temporary switch to a lower-cost plan, all within the chat.

4. **Strategic** Implications for Stakeholders

For Financial Services, Telcos, and Utilities

Agentic AI in core workflows	AI agents will autonomously retrieve documents, draft reports, update systems, and trigger follow-ups—always within policy and with human-in-the-loop controls.
Focus on affordability and collection	For telcos and utilities, GenAI will help better define user affordability. In collection processes, AI agents will augment credit analysts' efficiency by analyzing customer creditworthiness and pre-drafting hyper-personalized settlement agreements.
Conversational interactions	Knowledge workers will increasingly use conversational AI to access information and perform tasks, integrating them with agentic AI for more intuitive, context-aware, and collaborative strategy design and workflow management.
Standardized decisioning	The shift will be from ad-hoc, analyst-driven intuition to standardized, auditable workflows powered by AI. This will drive measurable gains in productivity, consistency, and compliance.
Operational advantage	Organizations that redesign work in this way will achieve faster onboarding and credit decisions, reduced manual workload in AML and collection processes, and improved fraud detection, all while maintaining regulatory compliance and customer trust.

Conclusion

2026 will mark the turning point where agentic AI, standardized workflows, and compliance-by-design become the new normal. Organizations that embrace this shift will unlock both operational efficiency and strategic advantages. The strategy of the future is conversational, contextual, and collaborative. It's about building platforms that understand nuance, adapt in real time, and empower every user—human or machine—to do their best work. Organizations that embrace this vision will not only keep pace with change—they'll set it.

www.crif.com