Kill switches, examiners and a sovereign alliance

CrowdStrike ships a kill switch for AI agents, the US OCC tells bank examiners to judge AI by the cost of failure, Bull and Schwarz Digits team up on sovereign AI, and Capgemini measures how far multi-agent AI has spread in finance.

The theme of the past weeks, controlling agents rather than just deploying them, moved from vendors' labs into products and supervision. A major security vendor now sells a way to stop every agent at once, and the main US bank supervisor gave its examiners a framework that puts autonomous payment agents in the highest risk tier. In Europe, a French and a German company agreed to build sovereign AI infrastructure together, and new research shows most banks and insurers are experimenting with multi-agent systems while few run them at scale.

AGENTS & TOOLING

CrowdStrike launches a kill switch for AI agents

At its Fal.Con Europe conference on 3 November, CrowdStrike introduced an AI kill switch that lets security teams cut agents off from AI models at the gateway in a single action, with stopping agents on endpoints described as coming later. It also announced detection that learns each agent's normal behaviour and flags deviations even when no policy is broken, though that feature is still at an early stage.

Why it matters: Supervisors increasingly expect firms to show they can stop an AI system quickly. A tested shutdown path is becoming as basic as a backup for any agent that touches live systems.

Our view: A central kill switch helps, but it works best alongside agents designed to fail safely. We build agent workflows so that stopping one mid-task leaves records consistent and the half-finished step clearly logged.

Source: CrowdStrike

Adobe puts its enterprise agent inside ChatGPT, Claude, Gemini and Copilot

Adobe announced on 4 November that CX Enterprise Coworker, its agent for marketing and customer-experience work, now runs inside ChatGPT, Claude, Gemini for business, Microsoft Copilot and Databricks Genie, with Amazon, IBM and Snowflake platforms to follow. In each assistant an administrator has to approve the app before staff can use it.

Why it matters: Business software is starting to reach staff through general-purpose assistants rather than its own screens. That moves part of the access control and audit trail into the assistant platform.

Our view: Before approving a vendor agent inside an assistant, check which data it can reach, where its actions are logged and whether those logs reach your own monitoring. The admin approval step is the right moment to ask.

Source: Adobe

ENTERPRISE ADOPTION

Capgemini: most banks and insurers test multi-agent AI, few run it at scale

A Capgemini study published on 5 November found that 59% of the banks and insurers surveyed are piloting or deploying multi-agent AI systems. Only about 10% have reached enterprise-wide orchestration, where agents work together across functions. The research frames measurable business returns as the main driver of the shift.

Why it matters: The gap between pilots and production is now the main challenge in financial services AI. Moving from one agent to many raises new questions about ownership, hand-offs and audit.

Our view: Orchestration fails most often at the hand-offs between agents. We define what each agent passes on, who owns the result and where a human signs off before connecting them into one workflow.

Source: Capgemini

REGULATION & GOVERNANCE

US OCC gives bank examiners an AI framework that rates risk by the cost of failure

On 4 November the US Office of the Comptroller of the Currency issued its first supervisory framework for examiners reviewing banks' use of AI. It tells examiners to judge an AI use by what a failure would cost, how reversible it is and how much human oversight applies, and places autonomous payment agents and real-time credit decisions in the highest tier. The OCC says the framework is a resource for examiners and does not create new requirements for banks.

Why it matters: Even without new rules, this shows what examiners will look at: delegated permissions, transaction limits, monitoring and a tested way to switch an agent off. European supervisors are asking similar questions.

Our view: Tiering by consequence is a practical way to decide where to start. We map each agent's actions to what an error would cost and put the strictest controls on the irreversible ones, such as payments.

Source: American Banker

SOVEREIGN AI & INFRASTRUCTURE

Bull and Schwarz Digits agree to build sovereign AI infrastructure together

France's Bull and Germany's Schwarz Digits, the IT arm of the group that owns Lidl, signed a memorandum of understanding on 3 November to develop AI and cloud services under European jurisdiction for regulated and strategic sectors. Schwarz Digits brings its STACKIT cloud, and Bull brings high-performance computing expertise and servers built at its factory in Angers. The agreement is preliminary and gives no financial terms or delivery dates.

Why it matters: European regulated firms have few providers that combine local ownership, local hardware and AI capacity. A Franco-German stack would widen the choice, once it exists as a product.

Our view: A memorandum is a direction, not a service. When we assess sovereign options, we look at what is available and contractually committed today, and design so a better European option can be adopted later.

Source: The Next Web

REGIONAL WATCH

CrowdStrike opens an in-country cloud in Saudi Arabia

CrowdStrike said on 4 November that its Falcon security platform now runs in-country in Saudi Arabia, with customer telemetry, workloads and management infrastructure hosted locally. The company says the setup is designed to fit Saudi cybersecurity and data protection rules and is available now. It is the first of the in-country clouds for Saudi Arabia, India and the UAE that CrowdStrike announced in January.

Why it matters: Security tools see a great deal of sensitive data. For banks and energy companies in the Kingdom, keeping that telemetry in-country removes a common data-residency objection.

Our view: Data residency applies to the tools around AI too, not just the model. When we design in-country AI deployments in the Gulf, we include logging, monitoring and security telemetry in the residency plan.

Source: CrowdStrike

WHAT WE'RE WATCHING

Microsoft has said its Saudi Arabia East cloud region will open to customers this month. The UK ICO's call for evidence on agentic AI closes on 20 November, and Qatar's Future AI Summit takes place in Doha on 23 November. In Europe, Google holds a sovereign AI summit in Brussels on 17 November and OVHcloud its AI summit in Paris on 19 November.

Want to apply this in your organisation?

We build production-grade AI agents for banking, insurance, manufacturing and energy. A working POC usually takes one to two weeks.

Read on agenticdynamic.com **Talk to us →**