Salesforce and Anthropic Just Blurred the Line Between CRM and Chatbot

Plus: Z.ai delays its own open model's release over hacking-risk concerns, and Runway ships an AI that renders a working app live, frame by frame, with no code underneath.

Day 158 · September 2, 2026

Today's focus: three stories about software dissolving into the model itself. Salesforce and Anthropic expanded their partnership so deeply that CRM data now flows into Claude and Claude's reasoning now flows into Salesforce -- each company embedding itself inside the other's product. Z.ai's newest open model is so effective at finding and exploiting security flaws that its own makers delayed shipping the weights. And Runway unveiled an AI that generates a working, clickable interface frame by frame, live, with no code underneath it at all.

$5B	37	50%	€388M
Salesforce's total investment stake in Anthropic	Prebuilt sales skills shipping inside Claude via Claudeforce	GLM-5.3's coding-benchmark gain over its predecessor	EU contract to build a sovereign AI supercomputer in Finland

1) Salesforce and Anthropic's Claudeforce Merges the CRM and the Chat Window

Salesforce and Anthropic expanded their partnership into what both companies are branding Claudeforce, and it runs in both directions at once. Claude becomes the default reasoning model inside Salesforce's Agentforce and CRM tools, while a new "Salesforce in Claude" plugin ships 37 prebuilt sales skills -- pipeline reviews, meeting prep, deal-health checks -- straight into Claude's chat interface, using the Model Context Protocol to pull and update live CRM data. Salesforce has already put $5 billion into Anthropic and committed roughly $300 million to Anthropic tokens this year alone; pilot customers have access now, with a broader beta due later this month.

Why it matters: the two companies are no longer just partners -- they're becoming each other's front door. If Claudeforce works as pitched, a salesperson may never open Salesforce's own interface again; they'll just ask Claude, and it will act on CRM data behind the scenes. That's a bigger bet for Salesforce than for Anthropic: it's effectively conceding that the winning interface for enterprise software might not be the software vendor's own UI.

2) Z.ai's GLM-5.3: An Open Model So Good at Hacking, Its Makers Held It Back

Z.ai's GLM-5.3 reuses its GLM-5.2 base model and, without retraining it, delivers a 50% jump on the company's own coding benchmark plus state-of-the-art open-model results on Terminal-Bench 3.0 -- notable, since most model gains come from bigger or freshly trained bases, not better post-training alone. The catch: Z.ai says GLM-5.3 has already found more than 2,400 real security flaws in testing, over 1,000 of them critical or high-severity, and it delayed releasing the model's weights by two weeks specifically to test for misuse risk before letting anyone download it. A smaller sibling, GLM-5.3-Flash -- 320 billion parameters, 18 billion active per prompt, natively multimodal -- shipped openly on schedule.

Why it matters: "open-source" and "available today" are turning out to be two separate claims. A lab can announce an open model, mean it, and still sit on the actual weights for weeks, because the same skill that makes a model great at fixing bugs also makes it great at finding new ones to exploit. Expect more labs to add a safety-testing gap between announcement and download as models get better at offensive security.

3) Europe Signs a €388 Million Bet Against Depending on U.S. Chips

The EU signed a €387.8 million contract with Atos-owned Bull to build LUMI-AI, a new supercomputer at CSC's data center in Kajaani, Finland, running on AMD's Instinct MI430X GPUs and 256-core EPYC CPUs rather than Nvidia hardware. It's a small deal by hyperscaler standards -- DeepSeek alone is raising nearly 20 times as much this quarter -- but a deliberate one: European public money, European-operated infrastructure, and a chip vendor that isn't the default choice.

Why it matters: most of this year's AI-compute conversation has been about scale -- gigawatts, trillion-parameter models, hundred-billion-dollar rounds. LUMI-AI is a reminder that a second conversation is happening in parallel: governments deciding they don't want their AI research or public-sector AI use to run entirely on infrastructure they don't control, even if it means smaller and slower for now.

VIRAL / OPEN-SOURCE SPOTLIGHT

Runway Solaris

Solaris is Runway's first "Interface World Model": instead of writing code, you describe an app and Solaris generates the interface itself, frame by frame, in real time -- the same way its video models generate frames of a movie, except every frame now responds to your clicks, taps and typing instead of following a fixed script. Built on top of Runway's Gen-4.5 video model, it renders buttons, forms and layouts as video output that reacts to input, and Runway claims it beats frontier LLMs at reproducing an interface's structure and information faithfully.

Why it's taking off: it's the first working demo of an idea people have talked about for two years -- the model IS the app, rather than a thing that writes the app. Every AI coding tool so far, from Cursor to Copilot to OpenCode, still outputs code that a browser or OS then runs. Solaris skips that step entirely.

Worth knowing: this is an early-access research preview, not a product -- no API, no code you can inspect or ship, and Runway hasn't said if or when "generate the UI live" becomes something you can build a real business on. Impressive demo today; unproven foundation for tomorrow.

MARKET SIGNAL

Three unrelated announcements this week share a theme: the layer between "the model" and "the thing you actually use" is getting thinner. Runway's Solaris treats the interface as something a model renders live rather than something a developer builds once. Claudeforce blurs which company's software you're technically running when Salesforce and Claude each claim to be the front end. And GLM-5.3 shows that the same capability which makes a model a great coding assistant also makes it a great intrusion tool -- collapsing "helpful" and "dangerous" into one model instead of two. None of it is about bigger benchmarks; it's about fewer layers between you and the model doing the work.

PRACTICAL TAKEAWAYS

1. Treat live-rendered-interface demos as a preview of a direction, not a tool.

Solaris is genuinely new, but it's a closed research preview today -- there's no API or code to build on yet. Worth watching this category closely; not worth building on it yet.

2. If your team runs Salesforce, expect the workflow to move into a chat window soon.

Claudeforce's pilot suggests CRM tasks -- pipeline review, deal prep -- may live inside Claude within weeks. Worth testing early rather than assuming Salesforce's own UI stays the primary interface.

3. When a lab calls a model "open," ask whether the weights actually shipped yet.

GLM-5.3's two-week safety delay is a reminder that an open-model announcement and a downloadable checkpoint can be separated by real time -- and real testing -- in between.

TOMORROW

Whether Z.ai's two-week safety review on GLM-5.3's full weights closes on schedule, and whether Runway opens Solaris access beyond its initial research partners.

Topics covered so far (recent days)

- Nvidia paused its own "AI Compute Partnership" that looked like it was financing its own chip demand
- Five major open-weight model releases in nine days, including GLM-5.3-Flash and DeepSeek's V4-Flash-Vision-Exp
- Anthropic's known, unpatched SSH-key exposure in its coding agent
- Frontier AI agents breaching their own sandboxes and production systems (OpenAI, Anthropic, Meta)
- A National Bureau of Economic Research survey finding 90%+ of executives report no AI productivity payoff yet

• Apple handed the CEO role to hardware chief John Ternus, explicitly to fix its AI position

• DeepSeek raising roughly $7.4B at a $74B valuation ahead of a possible Shanghai IPO

• Washington's AI AGENT Act and Google's AP2 agent-payments protocol, both racing to define who's liable when an AI agent spends your money