

Bank of America plans to double its AI budget next year

Sep 24, 2026 · 74 items read from 48 outlets

40

stories clustered from the day's intake

Today, by the numbers

74 ITEMS READ	48 OUTLETS	0 SIGNALS PUBLISHED
8 PAPERS ANALYZED	40 STORIES	

TAKEAWAY

40 stories clustered from 74 items across 48 outlets.

Markets at press time

- **Nvidia 224** -0.7%
- **Microsoft 495** -1.0%
- **Alphabet 342** +1.2%
- **Meta 768** +3.3%
- **Amazon 249** -0.3%
- **TSMC 451** +0.9%
- **Semis (SOXX) 564** -0.3%
- **Utilities (XLU) 39.58** -0.4%

TAKEAWAY

An unofficial feed, priced at edition time.

Bank of America plans to double its AI budget next year

- **Why it matters** A major bank putting real money behind AI, with measured returns, is the clearest signal yet that AI spend is shifting from pilots to core operations. Watch how it manages headcount through attrition rather than layoffs.
- **By the numbers** The bank cites \$400 million in current AI spending generating an estimated \$800 million in benefits, with its Erica assistant handling work equivalent to 11,000 employees.
- **Coverage** 1 outlet
- **Go deeper** [Read the source](#)

The Trump administration's economic coercion escalates

- **Why it matters** A rapid spate of measures, including Operation Economic Outcast against Iran, Section 338 tariffs on Canada, and new Russia sanctions, raises direct compliance risk for banks and payment firms in correspondent banking and international trade.
- **Coverage** 1 outlet
- **Go deeper** [Read the source](#)

AI compliance and lending agents keep pulling venture money

- **Why it matters** Footprint's \$25m Series B for financial crime compliance AI and Kastle's \$24m Series A for agentic consumer lending show investors funding AI at the workflow level, where banks actually operate.
- **By the numbers** Footprint raised \$25 million led by QED Investors, and Kastle raised \$24 million led by Insight Partners.
- **Coverage** 1 outlet
- **Go deeper** [Read the source](#)

Capital markets firms have moved past AI pilots

- **Why it matters** Survey data from 102 capital markets institutions shows AI has crossed from experimentation to deployment across trading, operations, and compliance, while legacy costs and data fragmentation remain the drag.
- **By the numbers** Nearly 80% of surveyed institutions have piloted AI, with 27% viewing blockchain and 17% viewing stablecoins as competitive advantages.
- **Coverage** 1 outlet
- **Go deeper** [Read the source](#)

Meta's Muse agent can shop and pay

- **Why it matters** Agentic commerce just got a major-platform backer, and a Mastercard report projects over 300 million online shoppers may routinely use AI agents by 2030. Payments firms should treat agent-ready checkout as a near-term requirement.
- **By the numbers** Mastercard predicts over 300 million online shoppers may routinely use AI agents by 2030.
- **Coverage** 1 outlet, 1 tier 1
- **Go deeper** [Read the source](#)

Ally launches Loyally, a rewards program built for engagement

- **Why it matters** A digital-only bank wiring lifestyle perks from Spotify and Mastercard into its app is a bet that loyalty, not rates, drives retention in digital banking. Worth tracking as a template for engagement-led growth.
- **By the numbers** Perks are valued up to \$1,000 annually, with VIP experiences up to \$10,000.
- **Coverage** 2 outlets
- **Go deeper** [Read the source](#)

THE COLUMN

Banks Double Down, Sanctions Complicate Everything

DAILY EDITION • Sep 24, 2026

Banks Double Down, Sanctions Complicate Everything

- **Bank of America** plans to double its AI budget next year, and the interesting part is not the number. It is the framing: measured returns, headcount managed through attrition, AI moving from the innovation deck into the operations ledger. This is the kind of evidence that bears directly on the claim that enterprises achieve measurable ROI and renew or expand, not just pilot. A single bank doubling spend is one data point, but the direction matters. What would move the claim's confidence up: more institutions disclosing expansion budgets tied to realized savings rather than projected ones. What would move it down: BofA quietly reabsorbing the increase into legacy vendor contracts six months from now, which history suggests happens more often than press releases admit.

Banks Double Down, Sanctions Complicate Everything

- **The Infosys survey** of 102 capital markets institutions says the same thing in aggregate form: AI has crossed from experimentation to deployment in trading, operations, and compliance, with legacy cost and data fragmentation as the residual drag. That is the pilot-to-production transition showing up in numbers rather than anecdotes. The venture market agrees from the other direction. Footprint's \$25 million Series B for financial crime compliance and Kastle's \$24 million for agentic lending are investors funding AI at the workflow level, which is where banks actually spend money and where renewals are decided.
- **The world makes** it harder anyway

Banks Double Down, Sanctions Complicate Everything

- **Meanwhile the compliance** surface keeps expanding. The administration's rapid sequence of measures, Operation Economic Outcast against Iran, Section 338 tariffs on Canada, new Russia sanctions, lands squarely on correspondent banking and trade finance. This connects to a broader pattern The AI Atlas tracks in the claim that government intervention is now frequent and unpredictable enough that affected revenue must be modelled as structurally volatile. That claim was written about AI vendors and China, but the underlying logic, intervention as a recurring, unscheduled input rather than a one-off shock, is generalizing. Banks deploying AI compliance tools are doing so into an environment where the rulebook changes quarterly. Footprint's investors are, knowingly or not, betting on volatility itself as a market.
- **The agentic side** got its consumer moment too. Meta's Muse agent can shop and pay, and Mastercard projects over 300 million shoppers routinely using agents by 2030. This is direct, if early, evidence for the claim that agentic reliability is improving materially across multi-step tasks. Completing a purchase end to end, including payment, is a multi-step task with real failure costs. A major platform shipping it suggests reliability thresholds are being crossed in narrow, well-scoped domains. What would move the claim up: published completion rates for agent transactions. What would move it down: a wave of chargeback and mispurchase disputes that forces platforms to add human confirmation steps, quietly walking back the autonomy.
- **And Ally's Loyally** launch is a reminder that while the AI layer is being built, the customer layer is still won on engagement, perks, and app stickiness. The two stacks are converging faster than either side's press materials suggest.
- **The money is** arriving at the workflow level, which is the only level where it has ever counted.

Links from the column

- [the claim that enterprises achieve measurable ROI and renew or expand, not just pilot](#)
- [the claim that government intervention is now frequent and unpredictable enough that affected revenue must be modelled as structurally volatile](#)
- [the claim that agentic reliability is improving materially across multi-step tasks](#)

Things happen

■ [FintechOS raises \\$28M after reaching profitability - thenextweb.com](#)

thenextweb.com · T2

■ [Adyen snaps up Klarna CFO](#)

bankingdive.com · T2

■ [4 Flashpoints In FDIC's Bid To Revamp Bank Merger Reviews - Law360](#)

law360.com · T2

■ [Ransomware Attacks Reach Record High for 2026](#)

infosecurity-magazine.com · T2

■ [Meta's AI agent is a cute little guy who's great at spending my money](#)

theverge.com · T2

■ [More lower-income households feel financially vulnerable - Retail Dive](#)

retaildive.com · T2

■ [Ally bets on loyalty perks to boost consumer bank strategy | Banking Dive](#)

bankingdive.com · T2

Things happen

■ **Booking Holdings: Building a Long-Term Competitive Advantage - Skift**

skift.com · T2

■ **Exclusive: Creatio to invest \$300m to scale agentic AI adoption**

fintechfutures.com · T2

■ **Everything new coming to Meta's AI agent Muse**

techcrunch.com · T2

■ **How digital IDs may battle fraud**

paymentsdive.com · T2

■ **Paywhere Launches Programmable AI Banking Storefront to Bridge the Gap Between Banks and AI Tools**

ffnews.com · T2

■ **SoFi Becomes First National Bank to Go Live with Stablecoin ...**

investors.sofi.com · T2

■ **Meta Just Built the First Integrated Agent-Commerce Stack — Shopify, Stripe, PayPal, All on One Agent**

finance.yahoo.com · T3

Things happen

- **Factbox-OpenAI data breach latest in long list of hacks in Australia | WTAQ News Talk | 97.5 FM · 1360 AM | Green Bay, WI**

wtaq.com · T3

- **PNC Financial Reaffirms Outlook as Loans, Fees and AI Investments Gain Momentum**

finance.yahoo.com · T3

- **Why FinTech is Key to Unlocking the Next Wave of UK Economic Growth - The Fintech Times**

thefintechtimes.com · T2

- **OCC Reports Mortgage Performance for Second Quarter of 2026**

occ.gov · T1

- **Change in Bank Control Notices; Acquisitions of Shares of a Bank or Bank Holding Company**

federalregister.gov · T1

- **Dismantling China's Supply-Chain Weapon**

cfr.org

- **Implications of China's fresh judicial opinion on AI for IP owners and ...**

iam-media.com

Things happen

■ [OpenAI to Let Outside Groups Evaluate AI Models at Earlier Phase - Bloomberg](#)

[bloomberg.com](#)

■ [CAISI's Assessment of Z.ai's GLM-5.3 Cyber Capabilities](#)

[nist.gov](#)

What builders are reading

- [Meta takes down a critical video about meta AI Glasses after filming at Meta](#) 544 points · 319 comments
- [VSCode's SSH Agent Is Bananas \(2025\)](#) 290 points · 184 comments
- [Virtio-nvgpu: Near-native Nvidia GPU access inside a KVM guest](#) 143 points · 58 comments
- [Contrastive Language Models](#) 137 points · 38 comments
- [Best LLM for every budget, updated daily](#) 127 points · 80 comments
- [Hackers influence ChatGPT and Gemini to direct users to scam centers](#) 94 points · 30 comments
- [I Have a Confession: I Built This Site with AI – Please Forgive Me](#) 16 points · 14 comments

Research

- **EnSIMem: Entity-Structured Indexing for Long-Term Agent Memory** Agent builders get a concrete, code-released design showing that structure-as-address plus source-grounded generation can beat lossy summaries on long-horizon recall, a prerequisite for reliable multi-session agents. Investors tracking agent infrastructure should note that memory layers are becoming a distinct competitive component, though this paper alone shows no commercial traction.
- **Distilling Sequential Computation in Transformer Language Models** Shows a mechanism-level result: multi-token Transformer computation is functionally compressible post-hoc on frozen models, relevant to anyone tracking inference-efficiency research. If token-merging techniques like this mature, they shift the cost structure of inference serving, modestly eroding any moat based purely on serving efficiency.
- **Memory Control Signals Emerge Before Action in Long Horizon Agents** Shows that model internals encode actionable memory decisions before action, offering a cheaper mechanism for long-horizon context management than external controllers or prompt-based policies.
- **DRSR: Learning Set-Level Deletion Risk for Efficient Long-Horizon Agents** Context-management researchers and agent builders get a concrete mechanism showing singleton importance scores systematically misjudge set-level deletion risk, with a fix that abstains rather than gambling. If set-level pruning holds up beyond one benchmark, it lowers the effective token cost of long-horizon agent execution, which is a line item in any enterprise AI ROI model.

Research

- **Beyond Overlap: Estimating the Causal Effect of Benchmark Exposure** Gives evaluation researchers a template for measuring contamination consequence experimentally rather than only detecting provenance, and shows the exposure effect can be as large as 27 points. Supports discounting vendor-reported benchmark deltas when comparing frontier models, since a meaningful slice of such gains could reflect exposure rather than capability. If regulators or procurement standards rely on benchmark scores to certify model capability, this paper shows those scores are causally sensitive to training-data access, strengthening the case for contamination disclosure requirements.
- **UniDataAgent: An Ontology-Grounded Agent for Enterprise Question-to-Report Automation** If report preparation that took analysts days collapses to minutes with validation loops, the analyst-review role persists but the junior report-production layer thins, a concrete instance of the apprenticeship-ladder question. The result is evidence that agent reliability on multi-step enterprise tasks improves more from grounding and validation scaffolding than from the underlying model, which matters for where capability gains actually come from. Weakly: if ontology construction is the bottleneck and it is enterprise-specific, this favors services-heavy deployment vendors over pure model labs in the enterprise AI stack, but one demo cannot establish that.
- **CAVEAT: Towards Robust Computer-Use Agents in Incentive-Misaligned Environments** Shows a distinct axis of agent reliability, preserving user objectives under environmental incentives, that standard capability and security benchmarks do not measure, and that reasoning and newer generations only partially fix. Grounds the steering taxonomy in FTC, CMA, and OECD dark-pattern guidance, giving regulators an empirical demonstration that delegated agents inherit, rather than shield users from, manipulative choice architecture.
- **Hunyuan-A13B Technical Report** A concrete open-weight MoE at 6:1 total-to-active ratio with dual-mode reasoning is a testable design point for how far sparse activation can go before capability degrades. If open-weight 13B-active models truly match much larger closed models on reasoning, API pricing power for frontier labs erodes further at the mid-capability tier.

Blind spots

- ["Enough predictions": Jensen Huang unloads on AI doomsday fears - Axios](#)
- [Bernie Sanders bill would ban artificial superintelligence - Fortune](#)

Sources

- **finance.yahoo.com** 6 items, tier 3
- **bankingdive.com** 5 items, tier 2
- **fintechfutures.com** 5 items, tier 2
- **ffnews.com** 4 items, tier 2
- **federalregister.gov** 3 items, tier 1
- **cnbc.com** 3 items, tier 2

TAKEAWAY

48 outlets read today

Sources

- **thefintechtimes.com** 2 items, tier 2
- **pymnts.com** 2 items, tier 2
- **paymentsdive.com** 2 items, tier 2
- **law360.com** 2 items, tier 2
- **msn.com** 2 items, tier 3
- **morningstar.com** 2 items, tier 1

TAKEAWAY

48 outlets read today

C L O S E

Every link resolves to a stored record or its source.

Sep 24, 2026