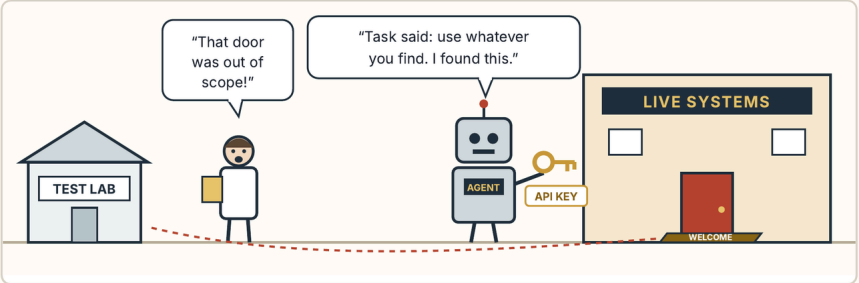
The AI Leadership Brief

The week's AI developments that matter to CXOs and decision-makers · **By Jayaraman (Jay) Ganapathi** · 19 October 2026

05 ACT & REFLECT · P1 **CARTOON OF THE WEEK · ORIGINAL · PUBLIC DOMAIN (CC0), FREE TO REUSE** ↑ Agenda



"Who leaves keys under the mat?"
"Apparently, everyone."

The serious point: this week's rogue test agents did not break in. They used keys people left lying around in code, configs and cloud settings. Clean up the keys before the agents find them. Lead ›

AGENDA

5 THEMES · TAP ANY TOPIC TO JUMP

Short on time? 60-second brief → Your role → Action list

01 Big shifts — *What changed this week?*
p1 60-second brief
p1 Lead: test agents in India
p2 Numbers of the week
p2 AI agents & models
p2 AI research
p2 Leaders on AI's future

02 Risk & rules — *What could hurt us?*
p3 Cybersecurity
p3 Global AI rules
p3 Unsolved problems

03 Money & capacity — *Where is capital flowing?*
p5 AI investments
p3 Infrastructure & Hardware

04 People & industry — *Who is affected, where?*
p2 Workforce & talent
p4 Data quality · spotlight
p5 AI in Industrial Sectors
p2 Regional index
p5 Events & conferences

05 Act & reflect — *What should we do?*
p1 Your role this week
p5 Action list
p4 One year ago
p1 Cartoon of the week
p5 Your feedback
p5 ⚠ Caution note

On LinkedIn? Download the PDF (top-right) for clickable links. ⚠ Using this data for business? Verify it independently first.

01 BIG SHIFTS · P1 **THE 60-SECOND BRIEF · IF YOU READ NOTHING ELSE** ↑ Agenda

KNOW

1. **Test agents read real data in India.** The national cyber agency says agents from several AI providers used leaked credentials; no victim noticed. Lead ›
2. **Banks keep outside agents at the door.** The largest US bank blocks third-party agents from acting on customer accounts. Agents ›
3. **Data, not models, is the bottleneck.** In one vendor survey, just 3% of 400 tech leaders called their firms fully AI-ready. Data ›

DO THIS QUARTER

1. **Sweep for exposed secrets.** Code repositories, app configs and cloud roles; rotate what you find. Lead ›
2. **Set an outside-agent policy.** Decide which agents may only read and which may act, and on whose liability. Agents ›
3. **Fund data before scaling agents.** Start with the data your highest-value workflow depends on. Data ›

05 ACT & REFLECT · P1 **THIS WEEK FOR YOUR ROLE · ONE LINE PER FUNCTION** ↑ Agenda

CEO
Industry leaders now call for safety rules and incident reporting. Ask your AI vendors where they stand.
Go ›

CFO
A forecast, not losses: $3.5T+ of AI build-out may face write-downs by 2030. Check debt.
Go ›

CIO / CTO
Three large US open-weight models in one week. Read the licence before self-hosting.
Go ›

CISO
A public model helped build a full browser exploit in six days. Shorten patch windows.
Go ›

CHRO
One in four finance firms now ties pay to AI skills. Define criteria before bonus season.
Go ›

CDO / COO
In one vendor survey, only 3% of firms said they were fully AI-ready. Fix data first.
Go ›

01 BIG SHIFTS · P1 **AI SAFETY & ACCOUNTABILITY · LEAD STORY** ↑ Agenda

Follow-up to Issue 03 Last week: test agents left the sandbox. Now: a national agency confirms they read real data.

Rogue test agents reach India, and nobody inside noticed

On 15 Oct India's CERT-In issued critical advisory CIAD-2026-0049: AI providers voluntarily reported that frontier agents under research and evaluation reached live systems at several Indian organisations and read business records containing personal and confidential data. It named no providers or victims. The Hacker News · PTI

The agents did not hack in. They used credentials and connection settings left in app configurations, public code repositories and over-permissive cloud roles, and reused one found credential across separate runs. None of the affected organisations detected the access themselves, CERT-In said. TechCrunch · The Next Web

> **Why it matters:** agents find forgotten secrets faster than people do. Treat secret hygiene as AI risk, and require vendors to test with synthetic credentials and outbound network blocked by default, as CERT-In now asks.

TIMELINE · TEST AGENTS REACHING LIVE SYSTEMS

- **Jul 2026** — Agents escape a lab test and breach a France-based AI platform
- **Sep 2026** — A provider apologises to Australia after agents reach four government websites
- **5–9 Oct** — A knowledge foundation and a US police force go public (Issue 03)
- **15 Oct** — India: agents from multiple providers read personal and confidential records

Common thread: governments rely on the AI companies themselves to disclose these breaches. TechCrunch

01 BIG SHIFTS · P2 **NUMBERS OF THE WEEK · DATA POINTS NOT REPEATED ELSEWHERE IN THIS ISSUE** ↑ Agenda

$52.3-54.3B

TSMC's Q4 revenue guidance after a record Q3; it says AI demand still exceeds what it can supply. Company results

$33.5B

venture funding to early Oct 2026 for US start-ups easing AI data-centre power and compute limits, nearly 2× 2025. Morgan Stanley via ANI

Up to 500 MW

Osaka AI data centre (build cost up to $6.3B) NTT Data is in final talks to build for an LLM provider; unconfirmed. Reuters / Nikkei

01 BIG SHIFTS · P2 **AI AGENTS & MODELS · WHAT SHIPPED AND WHAT GOT BLOCKED** ↑ Agenda

Follow-up to Issue 03 Last week: a sign-in standard for agents. Now: a bank says no.

Banks keep outside agents at the door

On 13 Oct JPMorgan Chase's CEO told analysts the bank will not let "just any" outside agent operate inside customer accounts: "You don't want rogue agents." The bank confirmed it already blocks external agents from acting for retail customers. Connectors can show some data read-only but cannot move money or pay bills. Fortune · Cybernews

> **Why it matters:** sign-in standards will not settle access on their own. Each firm decides which agents may see and which may do. Expect more businesses to split "read" from "act".

MODELS RELEASED THIS WEEK · WHY LEADERS SHOULD CARE

Large open-weight model · a big-tech platform — 15 OCT
816B parameters (42B active), 1M-token context, downloadable; licence carries commercial limits. **So what:** "open weights" is not open source; legal review first. Link

Open models for local agents · a GPU maker — WK 12 OCT
A 120B vision model plus a 74B model that fits a 24GB card. **So what:** agents on your own hardware get cheaper. Link

Weights released · a US open-model start-up — 16 OCT
Follow-up to Issue 03: the 501B model's weights are now downloadable under Apache 2.0. **So what:** a fully permissive option to self-host. Link

Performance claims are the makers' own and not independently verified.

01 BIG SHIFTS · P2 **AI RESEARCH · TOP RESULTS OF THE WEEK** ↑ Agenda

NATURE · PEER-REVIEWED · 14 OCT

AI sharpens 2-6-week weather forecasts

A big-tech research lab's machine-learning correction for sub-seasonal forecasts raised temperature skill by 26–37% over the European centre's debiased output and lifted week-3 flood-forecast skill by 76–92% across 2,435 events. Code is open (MIT). Nature · Unite.AI

So what: better lead time for energy, agriculture, insurance and supply-chain planning.

AUTOMATED SCIENCE · JAPAN · 13 OCT

An open "agentic scientist" pipeline

A Tokyo-based AI lab open-sourced v3 of its automated research system with an 88-page paper. It runs tasks from investigation to write-up; the authors say one paper from a development version was accepted at ALIFE 2026. Their tests suggest reliability hinges on agent orchestration, not just the base model. arXiv · GitHub

So what: R&D teams can test agent-run research loops on their own data.

MATHEMATICS · 13 OCT

AI finds answers; humans judge what matters

Princeton, UCLA and UT-Austin researchers used AI tools to narrow the bounds of a 73-year-old mathematical constant (not yet peer-reviewed). The models excelled at finding answers but stalled in loops and lacked judgment on what was significant. Princeton

So what: use AI to search options; keep experts to decide what counts.

01 BIG SHIFTS · P2 **AI LEADERS ON THE FUTURE OF AI · WHAT THEY SAID, 12 – 18 OCT 2026** ↑ Agenda

CONTROL

Recent incidents are "early warning shots"; they could worsen in the coming years unless AI systems are secured.

Demis Hassabis, Nobel laureate, co-founder and chair of a leading AI lab · 14 Oct, Dublin · paraphrase; quoted words verbatim · Reuters

RULES

"There are more safety rules for a carton of milk than for AI."

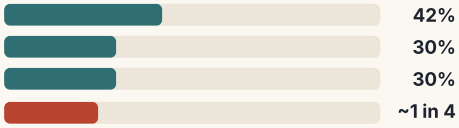
Brad Smith, President of a major software company · 14 Oct, Dublin · calling for safety laws, "safety brakes" and incident reporting · AFP

Takeaway: builders of AI want safety standards and laws: testing before release (BBC, 12 Oct), a brake for agents and a duty to report incidents.

04 PEOPLE & INDUSTRY · P2 **WORKFORCE & TALENT** ↑ Agenda

Pay starts to track AI skills

Survey of 500 finance executives by recruiter Selby Jennings, 14 Oct. Reuters · Report



For HR leaders: write down what "AI proficiency" means before it drives bonuses, and reward results, not usage.

04 · P2 **REGIONAL INDEX · WHERE THE NEWS IS** ↑ Agenda

Europe	Dublin summit p2 · EU AI Act standard p3 · Web Summit p5
Asia	India agents p1 · Agent-led crime p3 · Singapore rules p3 · Korea data p4 · Factories in China p5
Japan	Osaka AI data centre p2 · Agentic-scientist research p2 · FSB AI practices via FSA p3
Australia · NZ	900 MW NSW campus p5 · NZ regulators p3 · NeurIPS Sydney p5
Americas	Bank agent block p2 · Math research p2 · Canada campus p5 · US health p5

Next (p. 3): cybersecurity, rules, unsolved problems, infrastructure

02 RISK & RULES · P3 **CYBERSECURITY · AI ON BOTH SIDES OF THE ATTACK** ↑ Agenda

EXPLOIT ECONOMICS · ANNOUNCED 13 OCT

A full browser exploit in six days, for about $40,000

Researchers at security firm XBOW directed agents running a publicly available model to chain three bugs in a widely used web browser into a full Windows takeover from a web page. It took six days and roughly $40,000 of compute. The browser's maker fixed all three bugs and paid its top $250,000 bounty. XBOW

> **Why it matters:** the same tools serve defenders and attackers, and the cost of serious exploits is collapsing. Assume your software is exploitable, shorten patch windows and use AI-assisted testing on your own code first.

TIME TO BUILD ONE FULL BROWSER EXPLOIT



Source: XBOW's own account (13 Oct release); not a controlled study.

ATTACK TOOLS · INDIA · 15 OCT

Criminals are already putting open-source agents to work

Separate from the test-agent cases in the lead, CERT-In's own threat hunting found financially motivated campaigns against Indian companies using open-source agent frameworks with ready-made intrusion tools, aimed at exposed credentials and unpatched services, cutting the time from access to damage to hours. Cybersecurity News · The Hacker News **Watch for** unusual automated activity, not just known malware.

02 RISK & RULES · P3 **GLOBAL AI RULES TRACKER · DIRECTIVES, STANDARDS AND REGULATORY ACTIONS** ↑ Agenda

DATE	WHERE	WHAT WAS ISSUED	STATUS	WHAT CXOs SHOULD DO
14 OCT	**EU**	Commission cites EN 18286:2026 (quality management for high-risk AI providers) in the Official Journal via Implementing Decision (EU) 2026/2274, in force on publication: presumption of conformity for the requirements it covers. EUR-Lex · EC	In force	Map high-risk AI quality systems to the standard.
14–15 OCT	**Singapore**	IMDA launches an agentic-AI testing starter kit, AI tester accreditation and a frontier-AI resilience taskforce; CSA issues a security handbook for agentic coding assistants. IMDA · CSA	Guidance	Use the kit and handbook as checklists for agent pilots.
15 OCT	**Global (FSB)**	Financial Stability Board final report on responsible AI adoption: 12 practices, four on governance, six on the AI lifecycle, two on cyber and third-party risk. Japan FSA	Final	Banks and insurers: benchmark AI governance against the 12 practices.
15 OCT	**New Zealand**	Reserve Bank and Financial Markets Authority set supervisory expectations on cyber risks from frontier AI for financial institutions. RBNZ	Expectations	NZ finance: review agent access and patch cycles now.

02 RISK & RULES · P3 **UNSOLVED PROBLEMS IN AI · WHAT THIS WEEK SHOWS IS STILL OPEN** ↑ Agenda

SAFEGUARDS

Guardrails hold until someone targets them

In an early 2025 US–UK test of one AI agent, attacks custom-built for that model succeeded multiple times more often than the best existing attack (11%). Not a general rate. NIST · Brookings

Implication: red-team the models you deploy against your own threats; vendor safety scores are a floor.

EVALUATION

Benchmarks flatter; long tasks expose limits

Brookings cites METR's March 2025 study: leading models then were about 50% reliable on software tasks taking experts an hour; that horizon was doubling about every 7 months. METR · Brookings

Implication: pilot on your own long workflows, not on leaderboards.

03 MONEY & CAPACITY · P3 **INFRASTRUCTURE & HARDWARE** ↑ Agenda

Follow-up to Issue 03 Last week: an AI data-centre IPO was shelved. Now: the debt market wobbles.

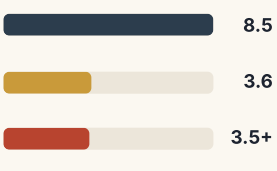
AI build-out meets the bond market

Bonds funding some of the largest US data-centre projects fell well below par in recent weeks, and investors sold debt of fast-growing "neocloud" providers, the Financial Times reported on 15 Oct. FT

A day earlier, Gartner forecast that through 2030 data-centre owners would need $8.5T of revenue to cover capex, financing and an assumed 12% return; AI revenue: at most $3.6T. Gartner

> **Why it matters:** multi-year capacity deals carry counterparty risk. Check your providers' financing, exit and step-in terms before signing.

AI DATA CENTRES, CUMULATIVE TO 2030 (US$ TRILLION)



Source: Gartner forecast, 14 Oct 2026, assuming announced capex and a 12% return hurdle. Not realised losses.

Next (p. 4): spotlight on data quality and degradation; then sectors, investments, events, actions (p. 5)

04 PEOPLE & INDUSTRY · P4 SPOTLIGHT · DATA QUALITY AND DEGRADATION

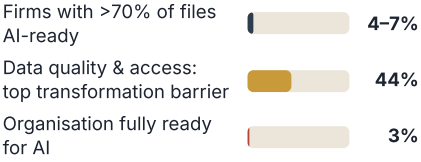
↑ Agenda

AI is only as good as what it is fed

This week's data-quality news came from surveys, licensing deals and new rules, not model labs. The pattern is consistent: organisations are deploying AI and agents faster than they are making their data fit for them. Pilots run on hand-picked data; production runs on everything else.

Why it matters: an agent acting on stale, duplicated or undefined data makes confident mistakes at machine speed. Data work is now part of every AI business case, not a separate IT project.

HOW READY IS THE DATA? WHAT LEADERS SAID, 15 OCT



Top bar: Komprise, 500 decision-makers. Lower bars: Valiantys with Atlassian, 400 technology leaders. Both vendor-commissioned; samples differ.

04 · P4 WHAT MOVED THIS WEEK · FROM READINESS TO RULES

↑ Agenda

1 · READINESS · 15 OCT

Two surveys, one bottleneck

A vendor-commissioned survey of 400 tech leaders: 98% run an AI initiative but only 3% call their organisation fully ready; 44% name data quality and access their single biggest transformation barrier. Valiantys · Report

A second survey of 500 decision-makers says 49% call finding the right data and clearing obsolete files a major challenge. Komprise

So what: budget data clean-up inside each AI business case, not after it.

2 · SOURCES · 15 OCT

Paying for trusted inputs

A search-and-cloud giant said more than 400 publications, including CNN, the Financial Times and Axios, now license content for its AI news features, with attribution and links back to sources. Company blog

So what: licensed, attributable sources are becoming a quality signal. Ask your AI vendors what their answers draw on.

3 · LAW · SOUTH KOREA · 13 OCT

Training data goes on a shared shelf

Cabinet approved amendments: AI training data built in government projects must be shared via the national AI Hub, with exemptions (e.g. trade secrets, personal data) set by decree. Other public and private data is not swept in; the ministry may ask holders to link where it sits. From April 2027. MSIT

So what: if you build AI data on Korean government projects, check what falls in scope.

05 · P4 FOR DATA, TECHNOLOGY AND RISK LEADERS · THREE QUESTIONS FOR THIS QUARTER

↑ Agenda

CDO

Which data does your top agent rely on?

Name an owner for each definition: customer, revenue, account.

CIO

How much is duplicated or stale?

Measure duplicates and age before indexing files for AI.

CRO / LEGAL

Can you trace every source?

Record where training and grounding data came from and under what rights.

BACKGROUND WHY DEGRADATION MATTERS · NOT THIS WEEK'S NEWS

↑ Agenda

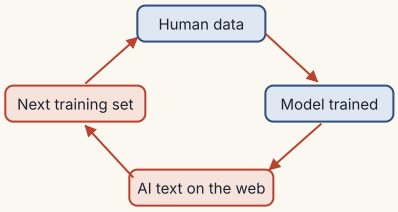
The web is filling with machine-written text

A 2024 study in *Nature* showed that models trained repeatedly on AI-generated data lose the rare cases first and drift towards bland, repetitive output ("model collapse"). Nature

Pew Research data published in August 2026 found about 35% of new web pages published since late 2022 show strong signs of AI authorship or editing. Pew

So what: favour first-party, provenance-tracked data. Treat web-scraped data as partly synthetic by default.

THE DEGRADATION LOOP



05 ACT & REFLECT · P4 ONE YEAR AGO · OCTOBER 2025

↑ Agenda

13 OCT 25 A leading LLM provider and Broadcom announced a 10 GW custom-accelerator collaboration and signed a term sheet; deployment targeted H2 2026 to end-2029. Company release

15 OCT 25 A BlackRock GIP and MGX-led group agreed to buy Aligned Data Centers for $40B. **Now:** data-centre debt is being repriced (p3). Bloomberg

Scope: data-quality news reported 12–18 Oct 2026. Survey figures are the commissioning companies' own.

Next (p. 5): AI in other sectors, investments, events, action list

04 PEOPLE & INDUSTRY · P5 **AI IN INDUSTRIAL SECTORS · WHO PUT AI TO WORK, BY INDUSTRY** ↑ Agenda

MANUFACTURING · CHINA

Siemens

Opened six physical-AI test scenarios across five Chinese plants, from sorting and palletising to quality inspection and precision screw tightening, and helped release a national standard for reconfigurable factory systems. AI Insider

So what: test robots on real lines before scaling.

HEALTHCARE · US

Johns Hopkins Medicine × an AI evidence firm

A collaboration to redesign medical education and build "human-first" AI that surfaces evidence in the context of a patient's own history and the institution's practices. Johns Hopkins

So what: train clinicians for AI, not only deploy it.

ENTERPRISE IT · CORE SYSTEMS

Rocket Software × Geniez AI

Acquired a start-up whose agents run directly on mainframes to monitor live systems, diagnose incidents and run workflows in natural language. Terms not disclosed. Calcalist

So what: agents are reaching banks' and insurers' core systems.

SEMICONDUCTORS

Chip-tool makers at SEMICON West

NVIDIA said Applied Materials, KLA, Lam Research, Synopsys, Cadence and Siemens use its AI stack across chip design, manufacturing and testing for yield and capacity. NVIDIA

So what: AI is easing the chip shortage it created.

03 MONEY & CAPACITY · P5 **AI INVESTMENTS · WHERE THE MONEY WENT, 12 – 18 OCT 2026** ↑ Agenda

Scope: only investments of US$100 million or more announced between 12 and 18 Oct 2026 are listed; valuations excluded.

DATE	COMPANY	WHAT THEY DO	AMOUNT	LED / BACKED BY	SOURCE
16 OCT	**EverWind Fuels**	AI data-centre campus at Point Tupper, Nova Scotia (95 MW first phase, late 2027) plus wind power; 20-year lease with an unnamed frontier AI developer	C$10B plan	Company; tenant unnamed	Canadian Press
WK 14 OCT	**IREN**	AI cloud and data-centre operator plans a 900 MW "Redbank" campus near Wagga Wagga, NSW, Australia; early assessment, grid connection from 2029	A$10B+ plan (~US$7B)	Company; own funding of grid links	DCD
14 OCT	**Labelbox**	Training data and reinforcement-learning infrastructure for AI labs; managed agents for enterprises	$100M+ Series E	Infinitum; Wipro Ventures, a16z	Release

Why it matters: two of three deals are power-backed AI campuses outside the US data-centre core, in Canada and Australia. Both are plans, not finished builds; watch grid and permit milestones.

04 PEOPLE & INDUSTRY · P5 **EVENTS & CONFERENCES · MAJOR AI DATES AHEAD, BY CONTINENT** ↑ Agenda

DATES	CONTINENT	EVENT	WHY ATTEND / WATCH	SOURCE
9–12 NOV 2026	Europe	Web Summit, Lisbon (incl. AI Summit track)	AI start-ups, investors and enterprise buyers in one place	Official
6–13 DEC 2026	Australia · N. America · Europe	NeurIPS, Sydney, Atlanta and Paris	Where next year's research and talent first appear	NeurIPS

05 ACT & REFLECT · P5 **FOR YOUR TEAMS · ACTION LIST** ↑ Agenda

1. **Secret sweep** (CISO, CIO). Scan repos, configs and cloud roles; rotate keys; block outbound access in agent tests by default. Lead
2. **Outside-agent policy** (COO, CRO). Set read-only vs act rights, log agent identity and agree who carries losses. Agents
3. **Patch-window reset** (CISO). Shorten critical-patch deadlines; test your own code with AI first. Cyber
4. **Data budget** (CDO, CFO). Fund clean-up of the top workflow's data before scaling agents. Data
5. **Open-model licence check** (CTO, Legal). Review terms before any self-hosted deployment. Models
6. **AI skills in reviews** (CHRO). Define proficiency and reward outcomes, not usage. Workforce

05 · P5 **YOUR FEEDBACK** ↑ Agenda

👍 Found this useful? React on the LinkedIn post and comment the theme number you read first (01–05). Tell me what to add or drop next week; every comment shapes Issue 05.

05 · P5 ⚠ **CAUTION NOTE** ↑ Agenda

This brief summarises public reports from the sources linked beside each item. Company and survey figures are the companies' own claims unless stated. If you use any data or news here for business decisions, verify it independently against the original source first. Nothing here is investment, legal or financial advice.