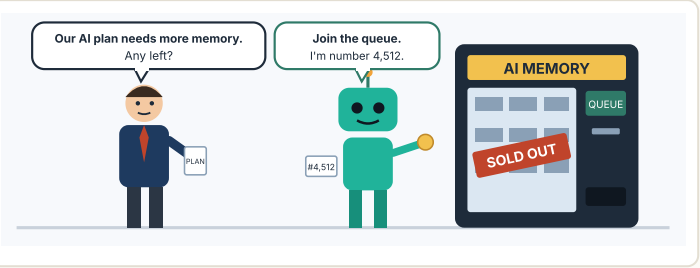
The AI Leadership Brief

Edition 06 · 26 Oct–1 Nov 2026 · By Jayaraman (Jay) Ganapathi, People Prime Worldwide · For CXOs and senior leaders

05 PEOPLE & ACTION · P1 **CARTOON OF THE WEEK** ↑ Agenda



"This week: SK hynix says memory supply is not keeping pace — and even the robots are queuing for AI memory."

AGENDA

5 THEMES · TAP TO JUMP

Short on time? 60-second summary → Your role → Leadership actions

01 Growth & revenue
Where is AI paying?
p1 60-second summary
p2 Lead: results & supply
p6 Key numbers
p6 Leaders said

02 Cost & output
What gets cheaper?
p2 New LLM releases or security patch releases
p3 AI research

03 Risk & compliance
What could hurt us?
p3 Vendor & security risk
p4 Unsolved problems
p4 Laws & guidance

04 Capital & markets
Where is money going?
p5 Investments ($100M+)
p4 Infrastructure
p5 Regional index
p6 Upcoming events

05 People & action
What do we do?
p1 Your role
p1 Actions
p5 Operating models
p5 Workforce
p6 One year ago
p6 Caution note

New section (p2). Every item now carries a **Business lens** (fact from the source unless marked "Our summary") and **Ask your leadership team** (editor's questions, not claims). Data quality section skipped (no qualifying item). Verify before business use.

01 GROWTH & REVENUE · P1 **60-SECOND SUMMARY** ↑ Agenda

1. **Demand is in the numbers.** SK hynix revenue +328% (prelim.); Intel Data Center and AI +78% (not attributed to AI by Intel); Capgemini's 2026 forecast raised, partly via acquisitions. ›
2. **Supply is on chipmakers' agenda.** SK hynix expects memory demand to exceed supply through 2030; Intel's CEO says the priority is to "unlock additional supply". ›
3. **One maker cut a model price.** OpenAI priced GPT-6.1 Terra 50% below its predecessor (maker's claim); Google is offering Gemini 4 Carbon first to "a subset of cybersecurity defenders". ›
4. **Vendor risk is real.** A cloud provider's hosted model returned 155 responses containing other users' data; provider says it is fixed. ›
5. **Four large rounds.** Four rounds of $100M+ totalling $1.27B (our sum of announced amounts): chips, cloud security, service agents and analytics. ›

05 PEOPLE & ACTION · P1 **THIS WEEK FOR YOUR ROLE** ↑ Agenda

CEO
Capgemini forecasts 2026 growth of +9.0–9.5%, about 5 points from acquisitions. ›

CFO
GPT-6.1 Terra standard API: $1 / $5 per million tokens, 50% below its predecessor (maker's claim). ›

CIO / CTO
SK hynix expects memory demand to exceed supply through 2030 (company outlook). ›

CISO / CRO
Provider's account: it knew of a cache-size mismatch before launch; 155 responses exposed other users' data. ›

CHRO
SAP India: 2 million people with AI skills by 2030; nearly 13% have taken part (company). ›

COO
Hitachi plans a 35,000-person digital company with physical AI as a focus from 1 Apr 2027. ›

05 PEOPLE & ACTION · P1 **LEADERSHIP ACTIONS THIS WEEK** ↑ Agenda

1. **Measure value, not activity.** Ask each unit for one AI use case with a revenue, cost or cycle-time baseline. ›
2. **Lock in capacity.** Check that cloud and hardware contracts cover the AI capacity your plans assume. ›
3. **Re-run rejected cases.** Re-price AI business cases turned down on cost at this week's lower token prices (per OpenAI). ›
4. **Tighten vendor terms.** Add data separation, incident notice and audit rights to AI vendor contracts. ›
5. **Gate agent actions.** Require a check before AI agents move money, data or customer records. ›
6. **Confirm patching.** Ask IT whether self-hosted AI serving software (vLLM, SGLang) is on the fixed versions. ›
7. **Set a skills target.** Publish an AI-skills goal with progress reporting, as SAP India does. ›
8. **India operations:** review MeitY's DEPA draft on AI training with partner data; it invites recommendations. ›

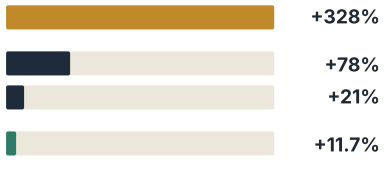
01 GROWTH & REVENUE · P2 **LEAD: DEMAND SHOWS UP IN THE RESULTS** ↑ Agenda

COMPANY RESULTS · 28–30 OCT 2026 · SOUTH KOREA, NORTH AMERICA, EUROPE

A memory maker, a processor maker and an IT-services firm report growth; SK hynix says supply is not keeping pace and Intel's CEO prioritises more supply

COMPANY · DATE	WHAT IT REPORTED	DEMAND AND SUPPLY, IN THE COMPANY'S WORDS	SOURCE
SK hynix 28 Oct · memory chips	3Q26 revenue 104.6 trillion won, up 328% year on year; operating margin 80%. Preliminary and unaudited.	Supply "is not keeping pace with the growth of the market"; expects "the structure in which demand exceeds supply to persist through 2030" (company outlook, not an independent forecast).	R1
Intel 29 Oct · processors	3Q26 revenue $16.5B, up 21%; Data Center and AI segment $7.3B, up 78%. Q4 guidance $15.0–16.0B (company forecast).	Growth led by the Data Center and AI segment (our reading: the segment added about $3.2B, more than the group's total increase). CEO Lip-Bu Tan: "Our immediate priority is to unlock additional supply to meet extraordinary demand for our Xeon family." (per company release). The release does not attribute the growth to AI.	R24
Capgemini 30 Oct · IT services	Q3 revenue €6,073M, up 11.7% at constant currency; bookings €5,401M, up 3.4%. 2026 growth outlook raised to +9.0% to +9.5% (constant currency; company forecast). Growth includes acquisitions (notably WNS and Cloud4C), which the company says had a "material contribution to growth" in North America, UK & Ireland and Asia-Pacific; it estimates about 5 points of 2026 growth is inorganic and gives no Q3 organic figure.	CEO Aiman Ezzat: "Clients are now focused on deploying AI at enterprise scale" (per company release). The release cites clients "transforming their business operations with agentic AI".	R25

YEAR-ON-YEAR GROWTH REPORTED THIS WEEK



Company figures; measures differ (total vs segment; reported vs constant currency). Bars scaled to +328%.

BUSINESS LENS · FROM THE SOURCES

SK hynix says supply "is not keeping pace" and expects demand to exceed it through 2030; Intel's CEO says the priority is to "unlock additional supply"; and Capgemini's CEO says clients are now focused on deploying AI at enterprise scale.

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTIONS

- Have we secured the compute our AI plans assume, at a known price?
- Which of our AI use cases has a measured revenue or cost result we could report?
- Are we still in pilots while clients of our suppliers scale?

02 COST & OUTPUT · P2 **NEW LLM RELEASES OR SECURITY PATCH RELEASES** ↑ Agenda

DATE	TYPE	WHAT WAS RELEASED	COST, ACCESS AND RISK (FROM THE SOURCE)	SOURCES
26 Oct	MODEL	**OpenAI GPT-6.1 Terra**, a medium-sized model for computer use and writing; for Plus, Pro, Enterprise and Edu users in ChatGPT Work and Codex, and via the API.	**$1 input / $5 output per million tokens (standard API price)**, which the maker says is 50% below its predecessor; it also says tasks finish "over three times faster" than GPT-6 Astra and GPT-6.1 Sol. Maker's claims, as reported by BleepingComputer; OpenAI's page not opened.	Announcement[R2] · report[R2a]
28 Oct	MODEL RESTRICTED	**Google Gemini 4 Carbon**, Google's newest frontier model, for cybersecurity tasks and long-running engineering work.	**$2 / $10 per million tokens.** Offered first to "a subset of cybersecurity defenders" through a dedicated access programme while safeguards are tested (as reported by Reuters). Google's page not opened.	Announcement[R3] · Reuters[R3a]
26 Oct	SECURITY PATCH	**vLLM v0.32.0** and **SGLang v0.5.22+**, open-source engines for running AI models on your own servers.	*Our view:* relevant mainly where your teams or vendors self-host models. vLLM 0.32.0 fixes five advisories published 26 Oct (per CVE records and vulnerability databases; IDs in the Source Register); CISA's enrichment of one record lists no known exploitation (as of 5 Nov 2026). SGLang v0.5.22 or later is advised after the cache incident on p3.	CVE[R4] · OSV[R4a] · release[R4b] · report[R5]

BUSINESS LENS · FROM THE SOURCES

OpenAI prices GPT-6.1 Terra 50% below its predecessor (maker's claim). Google is offering Gemini 4 Carbon first to a subset of cybersecurity defenders.

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTIONS

- Which AI business cases did we reject on cost that should be re-run now?
- Do we know which AI serving software our vendors run, and who patches it?

02 COST & OUTPUT · P3

AI RESEARCH: RELIABILITY AND QUALITY CHECKS

↑ Agenda

PREPRINT · 28 OCT 2026 · AGENT RELIABILITY

Preprint: a “decision gate” checks what an AI agent is about to do before it acts

The authors' framework, JevAct, turns domain policies into checks that run before selected agent actions; a separate decision model scores the structured questions as probabilities, and flagged actions get targeted feedback for a bounded revision. No change to the agent's model or tools is needed. Tested with eleven LLM agents on τ²-bench Airline and Telecom tasks and a function-calling benchmark (BFCL), the authors report reliability gains, most consistent on telecom and BFCL; airline results were “more mixed”. Authors' own measurements; not peer reviewed.

Shoham I, David Y, Akrabi O, Bilgory E, Levy I, Marreed S, Zeltyn S, Shlomov S. “Think Before You Act: Decision Gate for LLM Agents.” arXiv preprint v1, 28 Oct 2026. arxiv.org/abs/2610.38620v1[R10] · DOI 10.48550/arXiv.2610.38620 (arXiv-issued)

BUSINESS LENS · FROM THE SOURCE

The tests use τ²-bench Airline and Telecom tasks and BFCL. The authors report gains were not uniform: airline results were “more mixed”.

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTIONS

- Before our agents act on customers, money or data, what stops a wrong action?
- Who owns the policies an agent must follow?

PREPRINT (MEDRXIV) · V2 POSTED 28 OCT 2026 · HEALTHCARE QUALITY ASSURANCE

Preprint: AI models caught most planted errors in clinical-trial reports, and flagged issues in reports without them

The authors inserted methodological errors into randomised clinical-trial reports and asked 15 LLMs (four open-weight) to review them. On 58 trials published in 2026, major-error detection ranged from 58% to 99%, with mean review times of 1.4 to 16.3 minutes. Models “still flagged 14% to 29% of the error types in control reports without inserted errors”.

Limits: the human comparison (27% mean detection) comes from an earlier study of three reports; scoring used two blinded LLM judges; four authors are affiliated with ResearchHub Foundation, “whose open-science platform … is implementing AI-assisted review”. Not peer reviewed.

Brian D, Pereira Hernández XI, Yin H, Koury J, Rust R. “Large language models detect methodological errors that human peer review misses in clinical trials.” medRxiv preprint v2, posted 28 Oct 2026; v1 date not confirmed. doi.org/10.64898/2026.10.25.26366455[R11]

ERROR-DETECTION SCORES, 58 TRIALS (2026)



Ranges as reported (0–100%). Control flags: share of error types flagged in control reports without inserted errors. 15 models reviewed 3 trials; 12 reviewed the other 55.

BUSINESS LENS · FROM THE SOURCE

Mean review times were 1.4 to 16.3 minutes; in control reports without inserted errors, models still flagged 14% to 29% of error types.

ASK YOUR TEAM · EDITOR'S QUESTION

Where could AI pre-screening speed our quality or compliance reviews, and who reviews the flags it raises?

No peer-reviewed paper published in the week met this brief's relevance and sourcing rules; both items above are preprints.

03 RISK & COMPLIANCE · P3

VENDOR & SECURITY RISK

↑ Agenda

CLOUD AI SERVICE · REPORT DATED 26 OCT 2026 · PROVIDER'S OWN ACCOUNT

A large cloud provider says 155 responses from one hosted model contained other users' data

In the provider's incident report, a hosted coding model returned 155 responses containing data from other users' sessions between 27 Aug and 9 Sep 2026, all to very long prompts. The model was offered with a larger context window than its cache could hold; the provider says, “We were aware of this mismatch before launch.” A bug in the open-source SGLang engine then let those requests read leftover data from earlier requests. A researcher reported it on 8 Sep; the provider made the model private within about 40 minutes of declaring an incident, found “no evidence of deliberate exploitation”, and says no other hosted model was affected. The provider is not named because the news is negative.

Source: provider's incident report[R5]

BUSINESS LENS · FROM THE SOURCE

Per the provider, the exposure ran from 27 Aug to 9 Sep, was reported by a security researcher through its HackerOne programme, and it knew of the size mismatch before launch.

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTIONS

- Do our AI contracts cover data separation, incident notice and audit rights?
- Which sensitive workloads should run on dedicated rather than shared capacity?

03 RISK & COMPLIANCE · P4 **UNSOLVED PROBLEMS IN AI** ↑ Agenda

PREPRINT · 28 OCT 2026 · PROMPT INJECTION

Preprint: where a hidden instruction sits can change whether AI agents obey it

The authors built three agent testbeds (tool use, coding and daily workflows) and moved the same injected instruction between locations. They report that agents trust some parts of their environment more than others, which they call the "trust region effect". In one result, an agent powered by OpenAI's GPT-5.4 followed an unsafe instruction placed in a repository documentation file (README or TODO) 2.6 times more often on average than the same instruction placed in a code or data file. Not peer reviewed.

> Guo C, Liu X, Yang W, Feng M, Song J, Ma Q, Buractaon C, White C, Li B, Gao J. "Where You Inject Matters: The Trust Region Effect in LLM Agents." arXiv preprint v1, 28 Oct 2026. arxiv.org/abs/2610.38867[R8] · DOI 10.48550/arXiv.2610.38867

BUSINESS LENS · FROM THE SOURCE
In one result, an agent powered by GPT-5.4 followed the same unsafe instruction 2.6 times more often on average in a documentation file than in a code or data file.

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTION
Are our agents allowed to act on instructions they find in documents, emails or tickets?

03 RISK & COMPLIANCE · P4 **LAWS, RULES & GUIDANCE** ↑ Agenda

INDIA · LAUNCHED IN BENGALURU, REPORTED 31 OCT 2026 · DRAFT FRAMEWORK

MeitY releases a draft handbook for training AI on sensitive data without handing it over

Status: marked "Draft for Discussion"; a framework, not a law. **Scope:** the handbook on the Data Empowerment and Protection Architecture (DEPA) for AI Training, prepared with the iSPIRT Foundation and steered by MeitY, proposes that training happens where the data sits: in its main design, raw data stays with the training-data provider inside secure data environments and is not exposed to those using it for training. It names roles including training-data providers and consumers and data-exchange platforms. It calls itself a "living resource", invites "actionable recommendations" and gives no deadline. As reported by ETV Bharat, MeitY Secretary S Krishnan unveiled it at a roundtable.

Sources: MeitY handbook (PDF)[R12] · ETV Bharat[R12a]

BUSINESS LENS · FROM THE SOURCE
In the draft's main design, raw data stays with its provider and is not exposed to those training on it. The draft invites recommendations.

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTION
If we operate in India, would this model open data partnerships we cannot do today, and should we respond to the draft?

04 CAPITAL & MARKETS · P4 **INFRASTRUCTURE & CAPACITY** ↑ Agenda

EUROPE · ROMANIA · 30 OCT 2026 · COMPANY PLANS

Runware plans an assembly line for factory-built AI data-centre pods

As reported by TechCrunch, Runware's chief executive says the company plans to scale its Romania factory to 30 MW of data-centre capacity per day by 2028 ("nearly 1 gigawatt per month"), with a second factory due early next year and US expansion to follow. The liquid-cooled pods need power and internet but no water connection. Runware aims to have pods online within 48 hours of delivery (installation and setup, not manufacturing) and cites $1 billion of committed interest covering more than 20 MW. Company plans and claims, not measured output.

Source: TechCrunch (Dominic-Madori Davis)[R6]

BUSINESS LENS · FROM THE SOURCE
Runware aims to bring pods online within 48 hours of delivery and to produce 30 MW of capacity per day by 2028 (company plans, not measured output).

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTION
Would modular, faster-to-deploy capacity change our build, rent or wait decision for AI compute?

04 CAPITAL & MARKETS · P5

INVESTMENTS ($100M AND ABOVE)

↑ Agenda

COMPANY	AMOUNT	SERIES · LEAD(S)	BUSINESS PROBLEM IT TARGETS	CONFIRMED BY
FuriosaAI	**$620M**	D · no lead named; largest investor Korea's National Growth Fund ($280M)	AI inference chips for data centres; plans 20,000 units in 2026 and 40,000+ in 2027	Company release[R13], 29 Oct 2026
Upwind	**$400M**	C · Bessemer Venture Partners	Cloud and AI security	Company release[R14] (not opened; Reuters copy[R14a]), 28 Oct 2026
Serval	**$150M**	C · Sequoia	AI agents for enterprise service management	Company release[R15] (not opened; Fortune[R15a]), 29 Oct 2026
Hex	**$100M**	Peak XV, 137 Ventures	AI analytics for people and agents	Company blog[R16] (not opened; DealStreetAsia[R16a]), 27 Oct 2026

Only rounds of $100M or more, announced 26 Oct–1 Nov 2026. Amounts and leads come from each company's announcement; where we could not open it, a bylined report quoting it is linked. Total $1.27B is our sum of the four announced amounts; the "business problem" column is our short description. Counsel note supports only (R16b). No valuations.

BUSINESS LENS · OUR SUMMARY
Inference chips, cloud and AI security, service-management agents and AI analytics.

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTIONS
In each category, are we buying, building or waiting, and are our current vendors keeping pace?

05 PEOPLE & ACTION · P5

OPERATING MODELS: AI IN INDUSTRIAL SECTORS

↑ Agenda

INDUSTRIAL · JAPAN · 28 OCT 2026 · ANNOUNCED PLAN

Hitachi to launch a 35,000-person digital services company with physical AI as a named focus

Hitachi plans to launch Hitachi Digital Inc. on 1 April 2027, uniting GlobalLogic and Hitachi Digital Services: about 35,000 professionals in 27 countries serving more than 880 active clients. Named focus areas are physical AI, data and intelligence, cloud and site reliability engineering, and enterprise applications. An announced plan; the company is not yet operating.

Source: Hitachi announcement[R7]

BUSINESS LENS · FROM THE SOURCE
Hitachi is uniting two digital units into one company of about 35,000 people serving more than 880 clients.

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTION
Is our AI capability run as a business with clients and targets, or spread across cost centres?

05 PEOPLE & ACTION · P5

WORKFORCE & TALENT

↑ Agenda

INDIA · MUMBAI · 28 OCT 2026 · COMPANY PROGRAMME

SAP India adds programmes toward 2 million people with AI skills by 2030

SAP announced more programmes toward equipping two million people in India with AI-ready skills by 2030, part of a global pledge to help 12 million. SAP India says nearly 13% of the two million have already taken part through certification, student education and community programmes. Target and progress figures are the company's own.

Source: SAP India News Center[R9]

SAP INDIA AI-SKILLING GOAL



Company-reported progress toward its own target; not independently verified.

ASK YOUR LEADERSHIP TEAM · EDITOR'S QUESTION
What share of our workforce has AI training, and is it tied to changed roles?

04 CAPITAL & MARKETS · P5

REGIONAL INDEX

↑ Agenda

North America	Intel results · Upwind · Serval · Hex · Cross-tenant cache incident
Europe & UK	Capgemini results (France) · Runware pod factory (Romania)
India	MeitY DEPA draft · SAP India AI skilling
Japan	Hitachi Digital (physical AI)
Australia	No AI item met our sourcing and relevance rules this week.
South Korea	SK hynix results and outlook · FuriosaAI $620M
Global	New models and patches · Injection placement · Agent decision gate · AI review of trials

01 · P6 **LEADERS SAID** ↑ Agenda

"Capgemini delivered another solid quarter, with growth of +11.7% at constant exchange rates"

Aiman Ezzat, CEO, Capgemini · per company release, Paris, 30 Oct 2026 · release[R25]

"This Series D financing marks an important milestone for FuriosaAI and reflects the confidence of investors around the world in our technology and ability to execute."

June Paik, Co-Founder and CEO, FuriosaAI · per company release, 29 Oct 2026 · release[R13]

"Of the 2 million people we expect to train in AI skills before 2030, nearly 13% have already participated through certification, student education and community programmes."

Manish Prasad, President and Managing Director, SAP Indian Subcontinent · per company release, Mumbai, 28 Oct 2026 · release[R9]

All quotes come from company announcements. Intel's CEO quote appears in the lead (p2).

05 · P6 **ONE YEAR AGO** ↑ Agenda

Same week in 2025: 26 Oct–1 Nov 2025.

28 Oct 2025: A large online retailer said it would cut about 14,000 corporate jobs as it ramped up AI spending (per AP).

Associated Press: "[Company] cuts 14,000 corporate jobs as spending on artificial intelligence accelerates", 28 Oct 2025. AP copy on News 12[R21]

28 Oct 2025: NVIDIA agreed a $1.0B equity investment in Nokia alongside an AI-RAN partnership for mobile networks.

Nokia stock exchange release: "Inside Information: NVIDIA to make USD 1.0 billion equity investment in Nokia...", 28 Oct 2025. Open original[R22]

30 Oct 2025: OpenAI put Aardvark, an "agentic security researcher", into private beta; an archived copy of its page (Mar 2026) says the tool has since been renamed.

OpenAI launch post: "Introducing Aardvark", 30 Oct 2025. Archived original[R23]

01 · P6 **KEY NUMBERS** ↑ Agenda

+78%
Intel Data Center and AI revenue, 3Q26 year on year (company) ›

2030
SK hynix expects memory demand to exceed supply through 2030 (company outlook) ›

+9.0–9.5%
Capgemini's raised 2026 outlook, constant currency; about 5 points from acquisitions (company forecast) ›

50%
lower API price for OpenAI's GPT-6.1 Terra vs its predecessor (maker's claim) ›

155
responses with other users' data, one hosted model (provider's count) ›

$1.27B
across four rounds of $100M+ (our sum of announced amounts) ›

04 · P6 **UPCOMING EVENTS** ↑ Agenda

15–20 Nov 2026	SC26 (high-performance computing and AI infrastructure)[R18] · McCormick Place, Chicago
16–19 Nov 2026	MEDICA (medical technology and digital health)[R19] · Messe Düsseldorf (dates per JETRO listing[R19a])
9–10 Dec 2026	The AI Summit New York[R20] · Javits Center, New York (organiser site not opened; dates per event listings)

None of these events was listed in Edition 05.

05 PEOPLE & ACTION · P6 · CAUTION NOTE

Caution: verify before business use. This brief summarises the public sources linked beside each item. Forecasts, plans, preliminary results and maker's claims are the sources' own. **Business lens** lines restate the linked source unless marked "Our summary"; **Ask your leadership team** lines are the editor's questions, not facts or advice. Headlines, summaries, role tips and leadership actions are the editor's wording. Verify any item, figure or quote independently against the original source before using it for business decisions. Small red reference numbers (R1, R2...) point to the companion Source Register. ↑ Agenda

Found this useful? Give it a 👍 on LinkedIn and share it with a colleague.

AI model makers are named; the cloud provider in the incident item is not named because the news is negative. Every item links to its source.