LUMO INSIGHTS

Thursday, 16 July 2026

AI agents storm the enterprise, spur new guardrails

Executive Summary

Over the past 48 hours, developments from industry leaders and startups signal that autonomous AI agents are rapidly moving into real business operations. From front-office customer service to backend data systems, companies are not only adopting AI-driven agents but also racing to put new governance and trust frameworks in place. Meanwhile, open-source players are achieving unprecedented scale, challenging the notion that only Big Tech can deliver enterprise-grade AI autonomy.

AI agents reach the front office

On July 15, global consultancy PwC announced new "agentic" customer engagement and service solutions developed in partnership with OpenAI. These AI agents are designed to handle front-office interactions across the customer journey, from marketing and sales to contact-centre support (www.pwc.com [1]). By embedding OpenAI's advanced models into PwC's domain-specific playbooks, the solutions aim to deliver more personalised, seamless customer experiences at scale.

PwC's move is one of the first major efforts by a Big Four firm to productise generative AI for client operations. It signals that consultancies are now bringing AI out of the lab and into packaged services for enterprises. The expectation is that companies can deploy these ready-made AI agents to modernise their customer service and sales workflows much faster than developing in-house capabilities from scratch. The AI can handle routine inquiries or transactions with human-like conversations and even take action on behalf of customers, allowing human employees to focus on complex, high-value tasks that require empathy or expert judgment (www.prnewswire.com [2]).

For executives, this development lowers the barrier to adopting AI in customer-facing roles. Rather than waiting to hire AI teams or build custom solutions, organisations can leverage partnerships with firms like PwC to jump-start their AI transformation in the front office. That said, leaders should still insist on seeing real performance metrics - such as reductions in call handling times or improved customer satisfaction rates - achieved on their own data before fully scaling up these solutions . And given the reputational stakes of customer interactions, companies must maintain "human-in-the-loop" oversight for high-risk or sensitive cases to ensure AI-augmented services remain compliant and trustworthy .

References:

[1] www.pwc.com — https://www.pwc.com/us/en/about-us/newsroom/press-releases/pwc-openai-agentic-contact-service-solutions.html#:~:text=its%20agentic%20front%20office%20approach,service%20operations%2C%20improve%20service%20experiences

[2] www.prnewswire.com — https://www.prnewswire.com/news-releases/pwc-to-help-organizations-transform-agentic-customer-engagement-and-service-with-openai-302826711.html#:~:text=operating%20models%20can%20limit%20their,interactions%20paired%20with%20greater%20operational

One Agent, many tasks

In the legal and contracting arena, software provider IntelAgree has introduced a glimpse of how AI agents can take on highly specialised workflows. On July 15, the company launched “Saige Assist: Agent,” a general-purpose AI assistant for contract management now in private beta (www.morningstar.com [1]). Unlike typical contract tools that handle only one step of the process, this single agent can understand a company’s entire contracts portfolio by learning from its clause library and negotiation history (www.morningstar.com [2]). It can answer detailed questions about contractual terms, compare different versions of an agreement, suggest edits or redlines based on the organisation’s own standards, and even generate dashboards or reports on contract data - all within the contract management system (www.morningstar.com [3]). In other words, tasks that used to require multiple software tools (or significant human effort) are consolidated into one AI-driven “colleague.”

This all-in-one approach to knowledge work is a potential step-change in efficiency. For example, instead of separately using search tools, document comparison software, and manual editing, a legal operations team could rely on a single AI agent to do it all in a fraction of the time. By reasoning across all relevant internal data, the agent can ensure that a draft contract revision aligns with the company’s playbook and past deals, rather than a generic template (www.morningstar.com [4]). This reduces the need for numerous point solutions and manual cross-checks, potentially speeding up contract cycles and reducing errors.

The emergence of such multi-capable agents suggests a broader trend: enterprise AI assistants will increasingly handle end-to-end processes in specific domains. Leaders should identify areas - from contract management to finance or HR workflows - where a single AI agent with access to the right data could replace many narrow tools. The lesson from early adopters is to start small and safe: pilot these agents on low-risk, repetitive tasks (for instance, standard contract renewals) and verify that their outputs align with your policies and quality standards before scaling up. It’s also critical to implement guardrails: require that AI-proposed changes go through human approval and maintain detailed audit logs of every edit and decision an agent makes. With careful oversight, one well-trained agent could eventually streamline what used to be dozens of disconnected manual processes.

References:

[1] www.morningstar.com — https://www.morningstar.com/news/pr-newswire/20260715fl05407/intelagree-launches-saige-assist-agent-one-ai-agent-for-your-entire-contract-portfolio#:~:text=Assist%3A%20Agent%2C%20a%20single%2C%20general,clause%20library%2C%20playbooks%2C%20and%20negotiation

[2] www.morningstar.com — https://www.morningstar.com/news/pr-newswire/20260715fl05407/intelagree-launches-saige-assist-agent-one-ai-agent-for-your-entire-contract-portfolio#:~:text=narrow%2C%20typically%20a%20feature%20attached,generative%20AI%20offering%20can%20write

[3] www.morningstar.com — https://www.morningstar.com/news/pr-newswire/20260715fl05407/intelagree-launches-saige-assist-agent-one-ai-agent-for-your-entire-contract-portfolio#:~:text=library%2C%20playbooks%2C%20and%20negotiation%20history%2C,And%20unlike

[4] www.morningstar.com — https://www.morningstar.com/news/pr-newswire/20260715fl05407/intelagree-launches-saige-assist-agent-one-ai-agent-for-your-entire-contract-portfolio#:~:text=narrow%2C%20typically%20a%20feature%20attached,generative%20AI%20offering%20can%20write

From no-code to pro-code: integrating AI agents

The past two days have also seen major strides in making AI agent development accessible to both non-technical users and professional developers. Oracle announced a new AI-native builder for its Fusion Cloud Applications that invites software developers to create "agentic" applications using pro-code tools (siliconangle.com [1]). This extends Oracle’s existing AI Agent Studio - which previously catered to business users with low-code interfaces - to now support full-fledged coding via Visual Studio, Git, and even AI coding assistants like OpenAI’s Codex and Anthropic’s Claude

(siliconangle.com [2]). The key benefit is that these autonomous workflows can be developed and run *within* the secure environment of Oracle's core ERP, HR, and CRM systems, rather than as disconnected bots operating on the sidelines. By running natively inside the enterprise software (with access to live business data, roles, and approvals), Oracle's agentic apps come with identity, security, and audit controls baked in (siliconangle.com [3]). This built-in governance aims to eliminate one of the biggest hurdles in enterprise AI adoption - moving from prototype to production - by avoiding the need to retrofit compliance after the fact (www.oracle.com [4]).

Meanwhile, at the other end of the spectrum, startups are simplifying agent integration for non-developers. For example, Frigade this week launched a no-code "Skills" toolkit that lets product managers embed an AI assistant directly into their software without writing any code (www.prnewswire.com [5]). Unlike a typical chatbot that merely provides information, Frigade's in-app agent can actually perform tasks for users - a customer can ask the software's assistant to generate a report, change an account setting, or help configure an onboarding process, and the agent will carry out the action on the user's behalf (www.prnewswire.com [6]). This kind of automation can reduce support tickets and improve user satisfaction by letting people accomplish tasks through natural conversation instead of navigating menus (www.prnewswire.com [7]).

Both trends - empowering skilled developers and enabling no-coders - point to an accelerating democratisation of AI within enterprise products and processes. Whether through enterprise software giants like Oracle or agile startups like Frigade, it's becoming easier to weave AI agents into the fabric of day-to-day operations. For senior leaders, this means more teams (from IT to line-of-business) can begin experimenting with agent-driven improvements. The role of leadership is to encourage these innovations while ensuring that any AI integrated into core workflows adheres to your company's security, data governance, and quality standards. As a starting point, leaders might sponsor a small-scale pilot: for example, use Oracle's new toolkit to automate a routine back-office task in a controlled setting, or let a business unit try a no-code assistant for an internal process, and then rigorously evaluate the results.

References:

[1] siliconangle.com — https://siliconangle.com/2026/07/14/oracle-opens-fusion-agentic-applications-pro-code-developers-coding-agents/#:~:text=Oracle%20Corp,President%20of%20Applications%20Development%20Natalia

[2] siliconangle.com — https://siliconangle.com/2026/07/14/oracle-opens-fusion-agentic-applications-pro-code-developers-coding-agents/#:~:text=objective%2C%E2%80%9D%20added%20Rachelson,into%20a%20compatible%20coding%20agent

[3] siliconangle.com — https://siliconangle.com/2026/07/14/oracle-opens-fusion-agentic-applications-pro-code-developers-coding-agents/#:~:text=environment,would%20reason%20and%20transact%20against

[4] www.oracle.com — https://www.oracle.com/news/announcement/oracle-introduces-ai-native-builder-experience-2026-07-14/#:~:text=agents,The%20new

[5] www.prnewswire.com — https://www.prnewswire.com/news-releases/frigade-launches-skills-giving-any-product-an-ai-assistant-that-takes-action-for-users-no-code-required-302824744.html#:~:text=lets%20a%20company%27s%20AI%20assistant,the%20change%20in%20the%20product

[6] www.prnewswire.com — https://www.prnewswire.com/news-releases/frigade-launches-skills-giving-any-product-an-ai-assistant-that-takes-action-for-users-no-code-required-302824744.html#:~:text=Skills%20completes%20the%20request%20inside,the%20user%27s%20behalf%20and%20deflects

[7] www.prnewswire.com — https://www.prnewswire.com/news-releases/frigade-launches-skills-giving-any-product-an-ai-assistant-that-takes-action-for-users-no-code-required-302824744.html#:~:text=faster%20way%20to%20get%20things,building%20or%20maintaining%20a%20single

Closing the Governance Gap

Rapid adoption of autonomous agents is exposing a new weak point in enterprise readiness: governance and trust infrastructure. Recognising this, digital security firm Entrust announced an "Agentic AI Trust Accelerator" on July 14 - a collaborative programme to help companies establish rigorous identity verification, authorisation controls, and cryptographic proof for AI agents operating in business processes (finance.yahoo.com [1]). In essence, Entrust is bringing together a coalition of enterprises and tech partners to co-develop the standards and tools needed to ensure that when an

AI agent acts on your behalf, you know exactly *who* (or what) is behind it, what it's permitted to do, and what it actually did (finance.yahoo.com [2]). This focus on an identity-first approach to AI is timely, as organisations are eager to scale up autonomous workflows but need confidence that these "virtual employees" won't go rogue.

New research underscores why such an effort is urgently needed. In a recent IBM survey cited by Entrust, 77% of CIOs and CISOs said their AI adoption is already outpacing their ability to govern it, and 59% identified security and compliance worries as the top barrier to deploying AI agents at scale (finance.yahoo.com [3]). Traditional governance processes and IT controls simply haven't kept up with the speed and autonomy of modern AI systems. Experts note that as businesses shift from having humans in direct control to a 'human-on-the-loop' model of oversight, new frameworks are required to monitor and manage AI-driven decisions in real time (finance.yahoo.com [4]).

For business leaders, the takeaway is that investing in AI without simultaneously investing in governance is a recipe for trouble. As you begin to deploy AI agents, ensure you treat their "digital identity" and permissions with the same rigour as you would a human employee's credentials. Every autonomous agent should have its own unique credentials and defined scope of authority, with short-lived access keys and role-based permissions that can be tightly controlled and monitored. It's also critical to implement continuous verification and maintain auditable logs of all agent actions. By establishing this trust framework early - essentially an identity and audit layer for AI - organisations can confidently scale new autonomous solutions without losing oversight.

References:

[1] finance.yahoo.com — https://finance.yahoo.com/technology/ai/articles/entrust-launches-agentic-ai-trust-130000968.html#:~:text=The%20Accelerator%20program%20helps%20enterprises,announced%20the%20Agentic%20AI%20Trust
[2] finance.yahoo.com — https://finance.yahoo.com/technology/ai/articles/entrust-launches-agentic-ai-trust-130000968.html#:~:text=The%20Accelerator%20program%20helps%20enterprises,Entrust%2C%20a%20global
[3] finance.yahoo.com — https://finance.yahoo.com/technology/ai/articles/entrust-launches-agentic-ai-trust-130000968.html#:~:text=govern%20autonomous%20actions%20across%20systems%2C,human%20on%20the%20loop
[4] finance.yahoo.com — https://finance.yahoo.com/technology/ai/articles/entrust-launches-agentic-ai-trust-130000968.html#:~:text=and%20how%20its%20actions%20can,human%20on%20the%20loop

Open-Source agents usher new competition

Finally, a notable signal of the evolving AI agent landscape comes from the open-source community. Nous Research, the team behind the open-source agent "Hermes," is reportedly finalising a new funding round of over $75 million that would value the company at about $1.5 billion (techcrunch.com [1]). This is a remarkable milestone: it suggests that an independent, community-driven AI agent platform can achieve a "unicorn" valuation, underscoring the level of interest and trust it has garnered. Hermes launched as a free alternative to a popular closed-source agent and quickly amassed a huge following on GitHub (over 214,000 stars and 40,000 forks to date) from developers eager to build and share new capabilities (techcrunch.com [2]).

Hermes exemplifies how rapidly open-source AI tools can mature. The agent comes with a library of built-in "skills" (from web search and code writing to image analysis) and is designed to learn new skills from user interactions, reducing the need for manual updates (techcrunch.com [3]). Unlike proprietary systems that are locked behind corporate APIs, Hermes can be run by users locally on their own machines or on private servers, or accessed through Nous's optional cloud service for convenience (techcrunch.com [4]). This flexibility means businesses could potentially deploy autonomous agents that operate 24/7 on internal data, without handing the keys to a big tech vendor.

For senior executives, the rise of an open-source agent with such traction is both exciting and a cause for careful consideration. On one hand, it points to a future where innovation in AI is not limited to tech giants - your teams might leverage community-driven platforms like Hermes for faster prototyping and

customised solutions at lower cost. On the other hand, using open-source AI in an enterprise setting still requires rigorous due diligence. Leaders should ensure any open-source agent platform is thoroughly vetted for security, compliance, and reliability, just as with any vendor product. As a rule, ask for evidence that these tools have secure default configurations and robust governance mechanisms before entrusting them with sensitive business operations.

References:
[1] techcrunch.com — https://techcrunch.com/2026/07/13/hermes-agent-maker-nous-research-in-talks-for-new-funding-at-1-5b-valuation/#:~:text=PDT%20%C2%B7%20July%2013%2C%202026,USV
[2] techcrunch.com — https://techcrunch.com/2026/07/13/hermes-agent-maker-nous-research-in-talks-for-new-funding-at-1-5b-valuation/#:~:text=Hermes%20has%20amassed%20a%20massive,any%20setting%20up%20on%20their
[3] techcrunch.com — https://techcrunch.com/2026/07/13/hermes-agent-maker-nous-research-in-talks-for-new-funding-at-1-5b-valuation/#:~:text=Capital%2C%20and%20Balaji%20Srinivasan%2C%20according,in%20apps%20like%20Telegram%20and
[4] techcrunch.com — https://techcrunch.com/2026/07/13/hermes-agent-maker-nous-research-in-talks-for-new-funding-at-1-5b-valuation/#:~:text=chat%20with%20these%20agents%20or,via%20various%20paid%20tiers%20ranging

Key Statistics

- 77% of CIOs and CISOs say AI adoption is already outpacing governance capabilities ([finance.yahoo.com](https://finance.yahoo.com/technology/ai/articles/entrust-launches-agentic-ai-trust-130000968.html#:~:text=and%20how%20its%20actions%20can,human%20on%20the%20loop)).
- 59% of those leaders cite security and compliance as top barriers to AI deployment ([finance.yahoo.com](https://finance.yahoo.com/technology/ai/articles/entrust-launches-agentic-ai-trust-130000968.html#:~:text=and%20how%20its%20actions%20can,human%20on%20the%20loop)).
- 61% of data and AI leaders report "silent failures" (undetected errors) in their data systems ([www.alation.com](https://www.alation.com/blog/introducing-aios-alation-intelligence-operating-system/#:~:text=61,it%20isn%27t%20a%20technology%20problem)).
- Only 18% of data and AI leaders have scaled AI across multiple teams ([www.alation.com](https://www.alation.com/blog/introducing-aios-alation-intelligence-operating-system/#:~:text=61,it%20isn%27t%20a%20technology%20problem)).
- Open-source AI agent Hermes has ~214,000 GitHub stars and nearly 40,000 forks to its name ([techcrunch.com](https://techcrunch.com/2026/07/13/hermes-agent-maker-nous-research-in-talks-for-new-funding-at-1-5b-valuation/#:~:text=Hermes%20has%20amassed%20a%20massive,any%20setting%20up%20on%20their)).

KEY TAKEAWAY

AI agents are moving from pilot projects to core operations; leaders should move quickly to explore emerging agent solutions (via partners or in-house) to drive efficiency, while strengthening data governance and security for safe scaling.

Sources

PwC to Help Organizations Transform Agentic Customer Engagement and Service with OpenAI
https://www.pwc.com/us/en/about-us/newsroom/press-releases/pwc-openai-agentic-contact-service-solutions.html

IntelAgree Launches Saige Assist: Agent, One AI Agent for Your Entire Contract Portfolio
https://aijourn.com/intelagree-launches-saige-assist-agent-one-ai-agent-for-your-entire-contract-portfolio/

Oracle Introduces AI-Native Builder Experience to Create and Run Agentic Applications in Oracle Fusion Applications
https://www.oracle.com/news/announcement/oracle-introduces-ai-native-builder-experience-2026-07-14/

Entrust Launches the Agentic AI Trust Accelerator to Help Enterprises Move AI Agents From Pilot to Production
https://finance.yahoo.com/technology/ai/articles/entrust-launches-agentic-ai-trust-130000968.html

Frigade Launches Skills, Giving Any Product an AI Assistant That Takes Action for Users, No Code Required

https://www.prnewswire.com/news-releases/frigade-launches-skills-giving-any-product-an-ai-assistant-that-takes-action-for-users-no-code-required-302824744.html

Alation Launches AIOS™: All-New Intelligence Operating System for Enterprise AI
https://finance.yahoo.com/technology/ai/articles/alation-launches-aios-intelligence-operating-130000869.html

Oracle opens Fusion Agentic Applications to pro-code developers and coding agents
https://siliconangle.com/2026/07/14/oracle-opens-fusion-agentic-applications-pro-code-developers-coding-agents/

Hermes agent maker Nous Research in talks for new funding at $1.5B valuation
https://techcrunch.com/2026/07/13/hermes-agent-maker-nous-research-in-talks-for-new-funding-at-1-5b-valuation/