



NUCLEUS
RESEARCH

Nucleus Top Ten predictions for 2027

ANALYST

Ian Campbell, Trevor White, Evelyn McMullen, Charles Brennan, Alexander Wurm, Cameron Marsh, Duncan Van Kouteren, Charlotte Belke, Marisa Modjeski

The Bottom Line

Nucleus's predictions for 2027 show AI moving from experimentation to accountability. As agents take on real work, the question is no longer what they can do but how much authority organizations are willing to give them, and who is responsible when they get it wrong. ROI will rise with the level of decision-making organizations trust AI to handle, and "managing digital labor" will become an explicit part of employees' responsibilities. With more authority comes the need for more control. Enterprises will want contractual ownership of the tribal knowledge that forward-deployed engineers turn into reusable agents. Finance and HR will need shared permissions and audit trails before agents can touch workforce data. Recruiting vendors will need fraud detection that is explainable enough to hold up to legal scrutiny. The economics are catching up too: finance leaders are asking whether the juice is worth the GPU, and enterprises want model portability rather than lock-in. Under the hood, agents still need conventional predictive models to supply the numbers, and the unbundled data stack is coming back together to keep up with the pipelines AI demands. Meanwhile, the application interface is giving way to headless experiences run by agents, and store managers are trading dashboards for one-tap approvals.

Enterprise tribal knowledge is the new IP

The rise of forward-deployed engineering is placing technology vendors directly within customer environments to solve complex implementation challenges and identify new applications for agentic AI. By working directly with customers, engineers gain visibility into processes, decision criteria, exception handling, and institutional expertise that differentiates an organization from its competitors. As vendors translate this knowledge into reusable workflows, agents, and industry-specific capabilities, ownership and licensing rights will become increasingly important in enterprise technology contracts and AI governance. Enterprises will need contractual clarity on whether customer-specific workflows, models, and agent logic remain their property or can be reused and commercialized by vendors for other customers, including direct competitors. Nucleus expects vendors that establish clear IP boundaries and preserve customer ownership of proprietary workflows, models, and agent logic to gain a competitive advantage as enterprises seek to scale AI without relinquishing control of the intellectual property that differentiates their businesses.

Clear IP boundaries and customer ownership of proprietary AI workflows will become competitive differentiators for technology vendors.

AI vs. AI in Recruiting

Within two years, AI-driven candidate fraud detection will go from a competitive differentiator to a baseline requirement for any ATS. Candidate fraud has grown more sophisticated and now reaches many stages of the hiring funnel: AI-generated resumes, agents that mass-apply to roles the candidate has little intention of pursuing, and deepfakes in interviews. Together, these create compounding problems for recruiters and hiring teams, especially those already handling high applicant volumes. Many ATS vendors have tackled the problem head-on, but dedicated fraud detection across the whole pipeline is still rare enough to set a vendor apart. With fraud showing no signs of slowing, that advantage won't last. That detection carries its own risk. A fraud flag that removes a candidate is an automated employment decision, and those decisions face growing scrutiny. NYC requires annual bias audits of automated hiring tools, Illinois and California hold employers liable for discriminatory AI outcomes, and Colorado's new law takes effect in January 2027. Meanwhile, Mobley v. Workday is advancing on the theory that a vendor whose software rejects applicants can be held liable. A model that mistakes an accent, a disability, or a bad video connection for fraud exposes both the employer and vendor. The vendors that win will offer detection that is auditable, explainable, and keeps a human in the loop before any candidate is rejected.

The vendors that win will offer detection that is auditable, explainable, and keeps a human in the loop before any candidate is rejected.

Ranked by Your Promise

Retailers have long competed for shoppers who browse, compare, and occasionally forgive a late package. The next shopper does none of that. AI shopping agents can now find a product, check availability, and place an order directly, and order management system (OMS) vendors are already adding support for open agentic commerce protocols. Nucleus predicts that in 2027, delivery promise accuracy will move from an operations metric to a revenue metric, as agents route orders toward retailers whose stock and delivery dates hold up. The reason is memory. A person who gets a package two days late may still come back next month, while an agent tracking thousands of orders will simply stop recommending the merchant that broke its promise. That makes the OMS, the system deciding what is available to promise and from where, the deciding factor in who wins the order. Vendors that calculate commitments from live carrier performance and distributed inventory, and can prove that accuracy at peak, will move into revenue conversations, while those quoting static lead times will watch orders route elsewhere without ever learning why. Under promise and over deliver was always good advice, and in 2027, it becomes an algorithm.

Delivery promise accuracy will move from an operations metric to a revenue metric, as agents route orders toward retailers whose stock and delivery dates hold up.

Agents Need a Calculator

Every data science and machine learning (DSML) platform Nucleus evaluated this year shipped agents, and the marketing suggested LLMs would take over the work predictive models used to do. They will not. LLMs reason over language, and they remain unreliable at the numerical prediction that runs much of the enterprise, from demand forecasts and credit scores to fraud flags. Nucleus predicts that in 2027, the most valuable work agents do on DSML platforms will be calling conventional predictive models rather than replacing them, and the model registry will become a catalog of tools that agents across the business can use. Model Context Protocol (MCP) already gives agents a standard way to find and call those tools, and the pattern is visible in receivables credit recommendations and OMS returns disposition, where the agent runs the workflow, and a trained model supplies the number. That raises the stakes on fundamentals, because an agent calling a drifting model passes its errors along with full confidence. Vendors that expose validated, monitored models as agent tools will find new demand for decades of predictive work, while those chasing agent builders at the expense of traditional machine learning will lose the workloads that matter most. Even the smartest agent in the room still needs a calculator.

The most valuable work agents do on DSML platforms will be calling conventional predictive models rather than replacing them.

WFPA Plays Well With Others

Workforce planning used to be a once-a-year exercise built on a headcount spreadsheet and a lot of email. That is changing as human capital management (HCM), workforce management (WFM), finance, and enterprise resource planning (ERP) systems share more data, and as application programming interfaces (APIs) and the Model Context Protocol (MCP) make it easier for AI agents to work across them. Several planning vendors released MCP servers or agent connections this year. Nucleus predicts that in 2027, control over what AI agents can see and do will overtake new AI features as the deciding factor when finance teams buy workforce planning and analytics software. Finance usually owns the workforce plan, but HR owns the employee data behind it, and that split is where deals will stall. Several finance teams Nucleus spoke with this year still keep sensitive employee data out of their planning tools entirely over privacy concerns, and an agent asking nicely won't change that. Once an agent can pull payroll and HR records into a headcount scenario, the financial planning and analysis (FP&A) leader has to show HR who can see salary data, which system holds the official number, and a record of every change an agent proposed and who approved it. Vendors that give finance and HR one set of permissions and a clear audit trail of agent actions will become the hub other systems connect to. Walled gardens that make customers export data to plan will get left out of the conversation. Every planning system wants to be the single source of truth, and in 2027, the winners will be able to prove where their truth came from.

Control over what AI agents can see and do will overtake new AI features as the deciding factor when finance teams buy workforce planning and analytics software.

Store Managers are promoted to Approver

Store managers have spent the last decade as the most expensive integration layer in retail. Every system sends them something, whether it's a CSV export that assumes a free afternoon and a knack for pivot tables, a dashboard announcing that turnover is up without saying what to do about it, or a forty-page setup guide for the new holiday display. The manager reads it, decodes it, and turns it into a schedule change or a list of who does what. Nucleus predicts that in 2027, AI will take that translation job off their plate. Workforce management (WFM) systems will hand managers a specific decision to approve instead of a chart to interpret. Frontline task management, the software that tells store teams what to do each shift, will send corporate tasks straight to the associates on shift, broken into steps and matched to who has the skills.

The shift is happening now because scheduling, task, and HR data finally sit close enough together for AI to connect them. AI can spot that turnover is concentrated among associates scheduled outside their preferences and recommend the one hire that fixes it. It can also read the display guide and give each associate a short checklist with photo verification at the end. The hands-off store vendors keep demoing, but it still won't arrive in 2027. Managers are responsible for the schedule and the shelf, so the AI customers actually turn on will leave them the final call. Vendors that tie WFM and task data together and offer recommendations a manager can approve with one tap will win the next round of deals. Those still shipping another dashboard will find out how little patience a store manager has left. After a decade as the store's integration layer, managers are ready for a promotion.

WFM systems will hand managers a specific decision to approve instead of a chart to interpret.

The \$10,000 Coffee Maker

Nucleus predicts that 2027 will be the year companies realize they have been using flamethrowers to light birthday candles. The first wave of enterprise AI rewarded experimentation: if AI could automate a task, someone tried to automate it. Economics came later. As usage scales, organizations will discover that not every five-minute task deserves a frontier model, an agent orchestration layer, three integrations, and a governance committee. There is little financial genius in spending \$10,000 to build an espresso machine that saves employees a two-minute walk to Dunkin'. In 2027, finance leaders will increasingly force AI projects to answer a painfully simple question: Is the juice worth the GPU?

In 2027, finance leaders will increasingly force AI projects to answer a painfully simple question: is the juice worth the GPU?

The same math reshapes the vendor landscape. As foundation models become more capable and accessible through APIs, enterprises will evaluate them on cost, performance, security, and suitability for specific use cases, widening the divide between vendors that build proprietary models and tie customers to their own AI ecosystems and those that remain model-agnostic and compete through enterprise workflows, integrations, governance, and execution. Nucleus expects model portability and multimodel architectures to become standard enterprise requirements, reducing dependence on individual model providers and shifting differentiation toward the software and processes that operationalize AI. The winners will not be the organizations deploying AI everywhere. They will be the ones who know when a \$20 coffee maker works just fine.

Software Goes Headless

User experience drove differentiation on the internet and then on mobile, and with AI, it is being reimagined again around headless experiences, embeddability, and orchestration. Nucleus predicts that in 2027, the traditional application interface will begin to lose its position as the primary means by which users interact with enterprise software. Across ERP, supply chain planning, and manufacturing, vendors are building agentic layers that interpret business context, coordinate workflows, and execute actions across multiple applications, and SaaS vendors more broadly are exposing their capabilities as services that agents and applications call on demand through MCP endpoints. As these capabilities mature, users will spend less time navigating individual systems and more time interacting with an agent or orchestration layer that draws on ERP, planning, manufacturing, finance, and other systems behind the scenes.

Nucleus predicts that in 2027, the traditional application interface will begin to lose its position as the primary means by which users interact with enterprise software.

The business impact will come from reducing the friction between systems rather than making each application easier to use. Instead of moving between applications, reconciling information, and manually coordinating handoffs, users will initiate a process once and allow the agent layer to carry it across the relevant systems, with ROI coming from shorter decision cycles, fewer manual handoffs, less duplicated analysis, and greater process throughput. The strategic implication for vendors is equally important: owning the system of record no longer guarantees ownership of the user experience once every enterprise application is expected to serve agents as readily as it serves people. Vendors that still ask users to open one more application will lose share, and vendors that enable embedded, mobile, and headless experiences will gain it.

ROI Follows the Org Chart

Nucleus predicts that in 2027, enterprise AI maturity will be defined not by how many copilots or agents an organization has deployed, but by how much meaningful work it is prepared to delegate to them. Product roadmaps across ERP, supply chain planning, and manufacturing already show AI moving beyond research and recommendations into investigating exceptions, adjusting plans and schedules, creating transactions, and coordinating workflows, and vendors are building confidence thresholds, approvals, auditability, bounded authority, and escalation into these systems.

The next barrier to adoption will be organizational rather than technological: determining which decisions AI can make, under what conditions, and when a human must remain involved. The first wave of agentic AI focused almost entirely on what autonomous agents can do. The next will focus on who is responsible when they do it wrong.

That responsibility lands on people. As agents spread across customer service, sales, finance, development, and operations, employees will assign their work, review their performance, manage permissions, correct mistakes, and decide when an agent should be retrained or fired. Ironically, technology designed to reduce management overhead will create a new class of management, and Nucleus expects “managing digital labor” to appear in job descriptions, organizational structures, and performance evaluations by the end of 2027. Organizations that confine AI to summarization, research, and recommendations will see incremental productivity gains. Those that can safely delegate repeatable, high-volume decisions will compress cycle times and scale operations without adding headcount. In 2027, the AI maturity gap will become a delegation gap, and ROI will rise with the level of decision-making organizations trust AI to handle.

In 2027, the AI maturity gap will become a delegation gap, and ROI will follow the level of authority that organizations can confidently delegate to their systems.

The Data Stack Reunites

From 2019 to 2024, data integration vendors won by isolating a single step of the pipeline, and buyers assembled stacks from specialist tools. Nucleus expects 2027 to be the year the unbundled data stack comes back together. Development activity is booming, with AI-assisted coding pushing out more applications, models, and agents than data teams can wire up, and every one of them needs a pipeline behind it. As a result, customers are demanding consolidated integration, and the vendors gaining ground are the ones whose solution sets deliver the whole use case from source to curated data product. By consolidating onto one platform, customers retired point tools and the infrastructure under them, cutting ongoing infrastructure costs by 43 to 76 percent and total cost of ownership by up to 40 percent. In the proverbial tornado, buyers cling to simplicity and standards over radical optimization at any single step. Bundling alone does not win, and in 2027, customers will standardize on one platform per use case and retire the specialist stacks of the unbundled era.

Nucleus expects 2027 to be the year the unbundled data stack comes back together.