

Certinia Veda:

Orchestrating Trusted Digital Labor in Professional Services Automation

An Assessment of Certinia's [April 2026 Launch](#)

April 17, 2026




Executive Brief

This report examines Certinia’s April 2026 launch of Veda and its significance for the next phase of professional services automation. Certinia is positioning Veda as a modern AI operations engine for the intelligent services enterprise, extending AI across the full services journey, from first bid to final invoice, and moving beyond advisory AI toward action-oriented, agentic workflow orchestration.

- **Certinia’s Veda launch is more than an AI feature release.** It represents a practical shift from workflow automation to trusted, agentic orchestration in professional services, positioning PSA as a strategic control point for the execution of digital labor and knowledge work.
- **Veda is architected for consequential enterprise workflows.** Built on Salesforce Agentforce and grounded in Certinia’s system of record, Veda combines specialist agents, intelligent actions, business rules, permissions, and continuous context to support consequential workflows across estimation, staffing, delivery, financials, customer success, and margin optimization.
- **Trust is the differentiator.** Certinia has elevated trust from a governance afterthought to an architectural principle through explainability, auditable actions, domain-bounded agents, human control, permissions, traceability, and the Salesforce Einstein Trust Layer. In PSA, where poor AI recommendations can affect margins and client outcomes, that matters.
- **The ROI case is unusually tangible.** Certinia ties Veda to measurable outcomes in the three operational domains that matter most in services: staff, deliver, and serve. The company is framing AI not simply as labor savings, but as a force multiplier for utilization, margin protection, revenue growth, and customer retention.
- **Our AnalystANGLE:** Veda is an early blueprint for future-state agentic AI in knowledge work. We believe Certinia clients will view this announcement positively because it extends traditional PSA into high-value, agentic workflows in a practical, trusted manner aligned with how services organizations actually operate.

If Certinia continues to deepen its semantics, add causal decision intelligence, and expand its open-ecosystem posture, it can move beyond PSA leadership and become a model for how trusted digital labor translates into human capital ROI.

In the PSA market, Veda represents the shift from workflow automation to trusted digital labor



↳ Veda



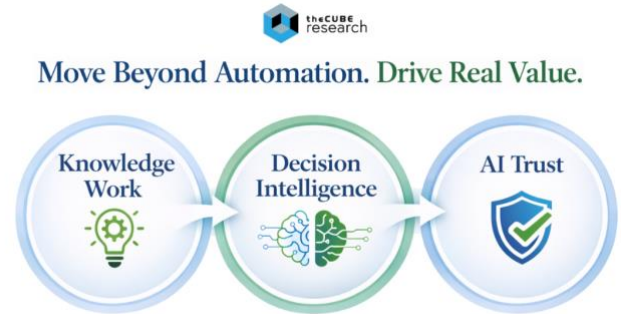
→ [Watch theCUBE NYSE Wire Interview](#)



The Shift in Professional Services ROI

Professional services automation is entering a new phase, and its implications on ROI extend beyond a single market category. This new phase will be driven by agentic AI workflows that deliver value beyond just task automation and will redefine the economics of PSA.

As context, here at [theCUBE Research](#), we have been focused on three core ideas that we believe will define the next frontier of agentic AI in the enterprise.




- AI is shifting from automating known tasks to supporting **knowledge work**, where it must interpret context, make judgments, and coordinate actions that materially affect business outcomes and ROI.
- Without **decision-intelligence** layers that add context, logic, and operational grounding, LLMs become a liability trap. They may sound fluent and coherent (because they are), but fluency and coherence alone do not equate to making judgments, balancing trade-offs, and scenario simulations.
- **AI trust** has become the currency of ROI. If users, managers, and executives do not trust how AI reaches conclusions, what actions it takes, and whether those actions align with policy and business reality, adoption stalls and value remains theoretical. In business-critical domains, AI must be explainable and auditable.

That is where we believe the PSA market is heading. The next wave of enterprise AI will not be won by vendors that merely attach agents to existing workflows, or even authentic agent workflows designed to automate known tasks. It will be won by those who can combine intelligent automation with trustworthy execution in environments where work is collaborative, decisions matter, and operational consequences are real.

That makes the PSA market especially important. PSA sits at the center of team-oriented knowledge work, where people need to collaborate to staff projects, manage client commitments, balance margins, oversee delivery, and make decisions that affect not only individual productivity but the performance of teams, customer relationships, and the financial health of the business itself.

It is in this context that [Certinia](#)'s recent announcement of [Veda](#) is more profound than it may appear on the surface. The company is balancing AI innovation with AI trust in a more practical and enterprise-ready way than much of the agentic AI market. Rather than treating agentic AI as a layer of generic automation, Certinia is grounding it in the operational realities of professional services, where context, coordination, and control matter as much as speed. **Veda** is notable not simply because it introduces a new suite of AI agents and intelligent workflow-based actions, but because it offers a glimpse into what the future of enterprise AI is likely to look like: AI that works alongside humans in consequential workflows, helps people make better decisions, and earns trust by being grounded in business context, governance, and real operational outcomes.

certinia
 Move Beyond AI Pilots.
 Run on Reality.
 Drive Real ROI.



A True Agentic Solution

The first thing to understand about **Veda** is that it is not chatbots bolted onto PSA. It's an intelligent operations engine built to move services organizations from reactive, manual workflows toward more autonomous, dynamic services operations. Certinia's own description is revealing: Veda is a suite of AI specialist agents and intelligent actions, built alongside its Professional Services, Customer Success, and Financial Management clouds, designed to deliver policy-bound, trusted, and ROI-focused autonomous workflows. That distinction matters because much of the current agentic AI market still conflates surface-level assistance with true operational agency. A real agentic solution does more than answer questions or automate known tasks. It must orchestrate actions, operate within guardrails, retain workflow context, and support task progression across a business process.

Veda sits atop its existing workflow automation layer, adding an agent orchestration layer, not as a replacement for enterprise systems but as the next layer of control and execution. In essence, they are shifting from workflow automation to intelligent agent orchestration, with humans and agents working together to make judgments, rather than agents roaming freely across enterprise systems without oversight.

Veda is best understood as a layered agentic architecture for professional services, built to turn Certinia's aggregate system of record into a live system of execution. At the top level, it includes a set of orchestration agents aligned with core PSA domains such as estimation, staffing, service delivery, customer success, financials, health monitoring, and customer lifecycle management. Each role-based agent coordinates workflows across the lifecycle rather than operating as an isolated copilot. Beneath that orchestration layer sits a broad array of specialist capabilities and "intelligent actions" designed for specific tasks such as estimate creation, staffing, work reallocation, meeting assistance, project summaries, forecasting, billing, customer health, and policy governance.

These are intended to work together in dynamic environments where humans and agents collaborate continuously. Certinia's positioning is that Veda carries a single thread of actionable intelligence from opportunity to renewal, using real service data, business logic, and contextual history to help PSA professionals sell with precision, staff with confidence, deliver with control, and serve customers proactively.

What is Veda

- Suite of Agents & Intelligent Actions
- Agent Orchestration + Workflow Automation
- Trusted AI Built for Services Organizations
- Proven and Demonstrable ROI
- Secure, Flexible and Scalable



In that sense, the architecture is not just a collection of agents. It is a coordinated services operating layer designed to apply agentic and generative AI across the entire customer lifecycle while remaining grounded in the underlying system of record. This is the right design choice for PSA. Services organizations need more than generic model access. They need domain-specific agents and intelligent actions that can operate across estimation, staffing, project oversight, customer success, and financial management in a governed way.

Certinia also made an important choice by building Veda on [Salesforce Agentforce](#), while grounding the experience in Certinia’s own services data, workflows, and business logic. They noted that Veda is built on more than 15 years of deep professional services institutional memory and can deliver trusted intelligence across the full lifecycle, from opportunity to renewal. That matters because Veda is not a purely probabilistic assistant but a deterministic system, whose actions follow business rules, controls, and compliance requirements.

It is also designed to maintain continuous context across the services lifecycle, preserving a single thread of intelligence from quote to renewal so that insights, actions, and decisions are not reset or fragmented as work moves across workflows. Rather than limiting AI to summaries or surface-level recommendations, Certinia is aiming **Veda** at agentic decisions in areas such as staffing, project oversight, and financial operations, where AI must help professionals act on context. In practical terms, that means helping knowledge workers create more accurate estimates from historical project and resource data, match the right people to the right work, and, in real time, proactively identify customer risks and expansion opportunities.

Just as important, Certinia exposes these intelligent actions within existing workflows, rather than forcing all value to live in a single interface. That makes the system more likely to fit real enterprise operating conditions, where service leaders need AI that can act with precision across the system of record, not simply generate responses on top of it. This also speaks to the power of openness in the Veda design. Rather than existing as an isolated new AI layer detached from prior enterprise investments, Veda is designed as a new orchestration and intelligence layer built on the capabilities that have already been powering Certinia customers. In that sense, Veda can call into and call out of these existing systems and processes using open-standard protocols such as MCP. That is strategically important because it can extend operational foundations in ways earlier workflow automation alone could not.

Veda



Deterministic by design

Every action follows your business rules, controls, and compliance.



Agentic execution

Move beyond summaries. Automate staffing, forecasting, and financial decisions.



Continuous context

One thread of intelligence from quote to renewal. No resets. No gaps.

Grounded in 15+ years of services operations— not the public internet.

90M+

Projects

1B+

Resources

2B+

Timecards

Capabilities Powering Veda



Architected for Knowledge Work

Professional services has always been a knowledge-worker business. It is not fundamentally a repetitive- task business, even if it contains repetitive tasks. It runs on judgment, staffing trade-offs, project economics, risk management, client communication, delivery coordination, and constant rebalancing as realities shift.

That is why PSA is such an important proving ground for next-generation enterprise AI. It is exactly the kind of market where AI has to move beyond speeding up known tasks and begin supporting work that is contextual, collaborative, and economically consequential. Our most recent Agentic AI Futures Index supports this view finding that 64% of enterprises plan to enable knowledge work use cases over the next 18 months.

This is also why LLM-only architectures are not enough. By design, LLMs are powerful engines of fluency and coherence, but they are not reliable engines of judgment, decision-making, or recommendation in dynamic enterprise environments.

They can summarize, predict likely next tokens, and mimic reasoning patterns from the past, but in a PSA environment where conditions are changing, trade-offs matter, and outcomes have financial consequences, that is not sufficient. PSA needs agentic systems that can interpret context, apply business logic, and operate within real workflow constraints.

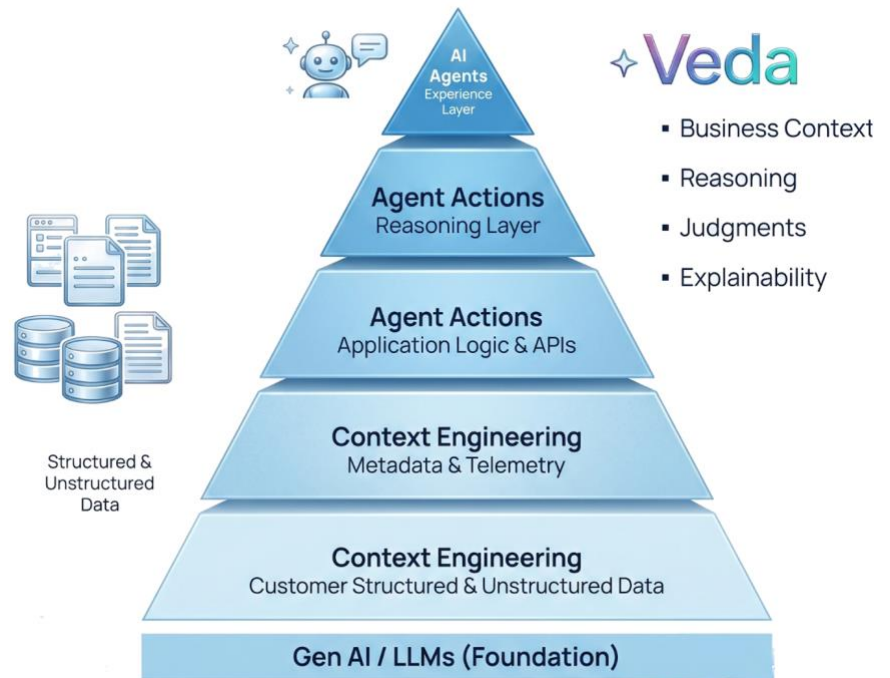
Veda’s architecture reflects that understanding. As shown in its Veda architecture, it builds upward from GenAI and LLMs through layers of context engineering, including structured and unstructured customer data, metadata, and telemetry, then into an application logic and API layer, a reasoning layer for agent actions, and finally an AI agent experience layer.

That layered design is important because it shows **Veda** is not treating the model as the product. It uses the model as the base, then adds the operational context and reasoning needed to make AI useful in knowledge-work-oriented workflows.

That is consistent with what we have seen leading agentic AI vendors building, and it is especially important in PSA, where digital labor must do more than automate steps. It must help humans make better judgments and consequential decisions.

Agentic AI
FUTURES INDEX

64% Envision enabling AI agents to perform knowledge work over the next 18 months



Delivering Trust in Outcomes

Trust is a central issue in enterprise AI, with 73% of enterprises planning to ramp up AI trust initiatives, according to our Agentic AI Future Index, and Certinia clearly recognizes this. The company’s position is that enterprise AI adoption, and therefore ROI, moves at the speed of trust, so Veda has been engineered to embed trust across the full stack rather than treat it as an afterthought. In the PSA market, trust is not just a model issue; it is a workflow, managerial, and operational issue. Poor AI-led recommendations can directly affect the bottom line.

Agentic AI
FUTURES INDEX

73% plan strategic or significant investments in trust & governance frameworks



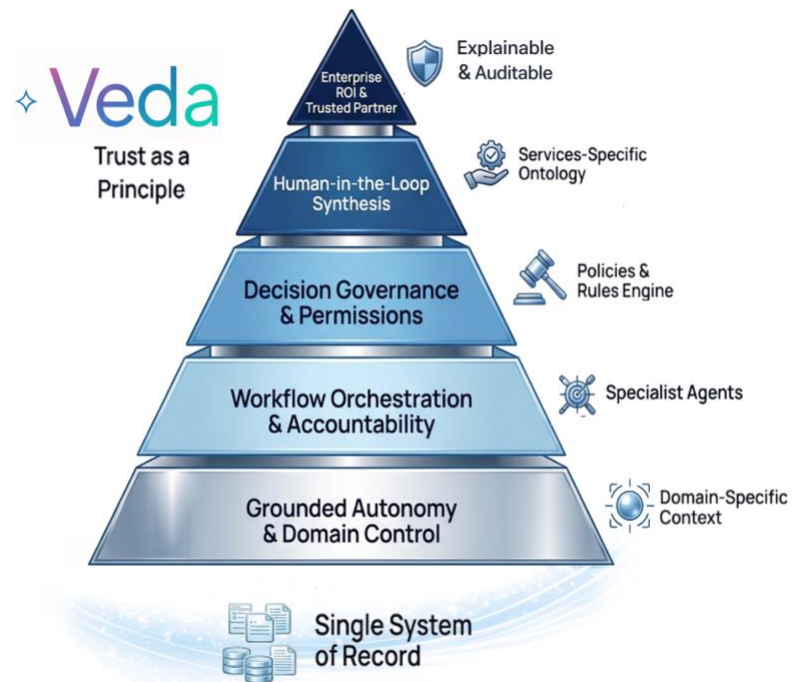
What appears differentiated in Certinia’s approach is that it has elevated trust from a governance afterthought to an architectural principle, delivering a more explicit trust model than many AI vendors. Veda’s trust mechanisms are not an isolated AI overlay, but a new intelligence and reasoning layer on top of the workflows already powering PSA. As such, **Veda** is designed to call into and call out of those existing capabilities, rather than forcing customers into a disconnected AI system. That strengthens trust in outcomes because Veda is grounded in the same operational systems already used to manage revenue, utilization, staffing, margin, and customer outcomes.

Importantly, Veda agents can explain their reasoning, giving users visibility into why an action was recommended, rather than relying on black-box behavior common in LLM-only architectures. Users also remain in control, since anything Veda can do through AI can also be done manually. In addition, specialist agents are limited to their functional domains and operate within a trust-tiered model: generative, where speed and synthesis are helpful, and deterministic, where precision, policy, and integrity matter most.

Operationally, this shows up as a multi-layer design grounded in a unified system of record:

- domain-specific context
- specialist agents
- policies and rules
- services-specific ontology
- auditable explainability

Veda also reinforces trust through permissions, traceability, secure data sourcing, and information security. Agent behavior is governed by Salesforce permissions, so users cannot access data they are not allowed to see, nor can they make updates beyond an agent’s authorization. On top of that, Veda relies on the [Salesforce Einstein Trust Layer](#) for protections such as zero data retention, prompt-injection prevention, and related controls.



Taken together, these integrated capabilities give Certinia a more mature, enterprise-ready trust posture than many vendors that still lead with autonomy first. In knowledge-work markets like PSA, trust, context, and reasoning are integral to value creation. These are the conditions that make durable ROI possible.



Open Ecosystem at Scale

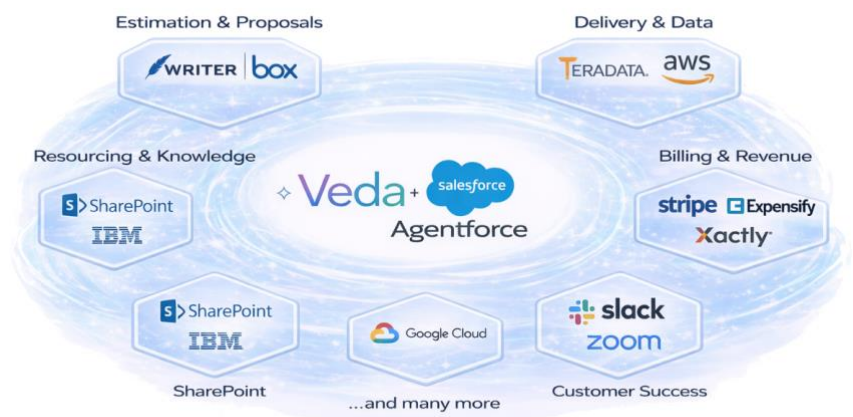
Veda’s ecosystem posture is a meaningful differentiator because it extends Certinia beyond a PSA application into a more open, interoperable agentic solution. Certinia was among the first companies to join [Salesforce’s Agentforce Partner Network](#), and later became a Salesforce Agentforce Model Context Protocol (MCP) partner, giving it an early position in the emerging AgentExchange and open-agent ecosystem.

Certinia’s framing is that the powerful agents Veda delivers are only half the story; to transform service delivery, those agents must communicate across systems, data sources, and workflows rather than remain trapped within a single application. They stress that to transform service delivery, those agents must communicate across systems, data sources, and workflows rather than remain trapped within a single workflow.

This means Veda's 10 specialist agents and 64+ intelligent actions are natively discoverable, deployable, and billable through the unified AgentExchange marketplace — a single catalog of 10,000+ Salesforce apps, 2,600 Slack apps, and 1,000+ Agentforce agents, sub-agents, tools, and MCP servers with integrated billing and one-click activation. Because Agentforce added native support for MCP and A2A — open standards that let AI agents access information and collaborate regardless of the platform they’re built on [Salesforce](#), Veda agents can call out to partner MCP servers across the ecosystem, and, just as importantly, be invoked by external agents.

The value becomes concrete across real services lifecycles. A few examples include:

- **Estimation & proposals:** Veda's Estimation Agent drafts the scope, hands it off to a WRITER for an SOW, and asks Box to confirm T&Cs.
- **Resourcing compliance:** IBM's pre-built compliance agents run alongside Veda's Resourcing Agent to enforce data residency and certification rules.
- **Billing & revenue:** When Veda reaches a delivery milestone, Stripe collects payment, syncs expenses to Expensify, and updates commissions in Xactly.



Since Veda exposes its own MCP and A2A surfaces, an external agent (e.g., a finance team's Workday agent) can invoke Veda to pull live project health, utilization, or margin data without custom API work. This is what "first-class citizen in the agent economy" actually means operationally. And, according to our Agentic AI Futures Index survey, this is exactly where enterprises want to head.

For Certinia clients, this means the ability to mix and match Veda agents within the broader partner ecosystem to extend capabilities the client may need, and to participate in multi-vendor agent workflows, while still grounded in Certinia’s single-record, Salesforce-native services data foundation.

Agentic AI
FUTURES INDEX

59% Say open protocols and standards are key to future-state agentic systems



ROI by Design

One of the strongest aspects of the Veda announcement is that Certinia is not talking about AI value in vague terms. It is tying AI to specific operational metrics and team-based productivity gains. Just as importantly, it is positioning Veda as a value-focused (ROI) solution by design: built for scale, architected to plug into existing environments without rip-and-replace, broadly accessible through a unified subscription model, and priced to align cost with measurable business output rather than novelty. They are bringing the promise of digital labor transformation to life, which, in our survey of enterprise leaders, 71% said is inevitable.

The company says Veda can return up to 10 hours of capacity per month for resource managers and up to 20 hours for project managers, while reducing preparation time and improving retention workflows. Certinia also frames Veda as a force multiplier on top of productivity gains already delivered by the core platform.

The value framework it has shown is useful because it breaks ROI into three operational domains that matter most in PSA: staff, deliver, and serve.

- In staffing, improvements in resource management efficiency, billable utilization, and time-to-staff.
- In delivery, project management efficiency, margin optimization, and reduction in unbilled time.
- In customer success, gains in retention, expansion revenue, and broader service efficiency.

That structure makes the business case more legible to enterprise buyers by mapping AI value directly to the economic levers that services leaders already manage. Furthermore, Veda, a new intelligence and reasoning layer atop the systems and workflows already powering professional services automation, improves the economics of systems already tied to revenue, utilization, staffing, margin, and customer outcomes.

Certinia is also driving home an important ROI differentiator: don't think of Veda as only about doing the same work faster, but also as an amplifier of a firm's best practices, policies, and operational capabilities. That is an important conceptual shift. It moves the conversation from labor substitution to capability amplification.

Certinia also links ROI directly to the economics of core services, arguing that even modest improvements in utilization can translate into meaningful revenue and EBITDA gains at scale. That is precisely the framing enterprise buyers need. It shifts the conversation from AI as a feature to AI as an operating model, and it reinforces the broader point of this brief: in knowledge-work markets like PSA, trust, context, and reasoning are not separate from ROI. They are the conditions that make ROI possible.

Agentic AI
FUTURES INDEX

71% say the rise of digital labor is inevitable and will transform the human capital ROI equation



◆ Veda Key Drivers of ROI



- Resource efficiency
- Billable utilization
- Time-to-staff



- Project management efficiency
- Margin optimization
- Reduced unbilled time



- Retention
- Expansion revenue
- Service efficiency



AnalystANGLE – Our Take

Our view is that Certinia is doing something more important than launching another AI feature set for PSA. It is building a practical blueprint for deploying agentic AI in enterprise knowledge-work environments: grounded in a unified system of record, shaped by domain-specific workflows, governed by permissions and policy, and designed to earn trust before maximizing autonomy.

More importantly, Veda promises meaningful impact for Certinia’s clients by targeting the areas in services organizations where economic leverage is highest and operational friction is most persistent. If Certinia delivers as intended, clients should see more than incremental productivity gains. They should see a stronger ability to improve margins, grow revenue more predictably, and operate as more coordinated hybrid organizations, with humans and digital workers collaborating within the same system of record and operational context. The real goal is not simply to do work faster, but to execute with more precision and better judgment across the full services lifecycle.

This is strategically significant for PSA, but also far beyond it. Any enterprise function built around consequential collaborative judgment will face the same challenge. AI must do more than generate answers. It must help people make better decisions together. That is why Veda matters. Certinia is not simply using AI to accelerate tasks. It uses AI to reshape how service organizations staff work, manage delivery, protect margins, and coordinate customer outcomes. In that sense, Veda is one of the clearer examples of agentic AI moving from demo theater to operating discipline.

For clients, the broader lesson is clear. The future of enterprise AI will be defined in environments where work is team-based, judgment-intensive, and economically consequential. That is exactly where agentic systems need context, reasoning, explainability, auditability, and collaborative control. The Agentic AI Futures Index reinforces this direction. Organizations are planning for digital labor and trust frameworks, while also showing growing interest in decision intelligence technologies. That is the road from automation to true AI decision intelligence.

From here, we see three priorities for Certinia to extend its leadership with Veda.

- First, deepen the semantic layer using knowledge graphs to enable stronger entity relationships across PSA workflows, improving contextual precision, interoperability, memory, and explainability.
- Second, evolve toward causal discovery and inference to deliver more reliable scenario analysis, clearer root-cause identification, and more precise recommendations about what actions are likely to improve outcomes.
- Third, continue expanding the ecosystem and open-protocol model, so **Veda** becomes a broader coordination layer for hybrid digital labor across enterprise environments, not just within PSA.


If Certinia executes on those priorities, it will move beyond PSA leadership. It will become a poster child for future-state agentic systems for knowledge workers: systems that combine trust, collaboration, and increasingly rich decision intelligence to deliver stronger explainability, auditability, and durable business value.



Veda promises to provide new “super powers” to teams of PSA knowledge workers by delivering a practical blueprint for trusted agentic AI workflows.



Q Which of the following AI decision intelligence technologies does your organization plan to invest in over the next 18 months?





Scott Hebner
Principal Analyst for AI
and CMO Practice Leader



[Contact Me](#)



[Read More AI Research](#)



[Subscribe to Next Frontiers of AI Digest](#)

Disclaimer

The views expressed in this content are the author's own and do not necessarily reflect those of any affiliated organization, including theCUBE Research, SiliconANGLE Media, or other publishing platforms. This material, including any AEO Advantage Index assessment, is provided solely for informational, educational, and diagnostic purposes. All findings, scores, and recommendations reflect an evidence-based, point-in-time analytical opinion derived from the information reviewed and are not guarantees of market performance, AI-engine behavior, or future results. Nothing herein constitutes financial, legal, investment, or other professional advice, nor should it be construed as an endorsement or recommendation. Hebner Advisories LLC and theCUBE Research disclaim any representations or warranties, express or implied, regarding the accuracy, completeness, or suitability of the information, findings, or scoring outputs provided. Readers are solely responsible for their interpretation and use of this material.

