

# certinia

# Veda

## Certinia's AI engine for service operations.

Most AI tells you what happened.  
Veda changes what happens next.

**Generate Resource Summary**

Summary

Preston Mayo, a Consultant with 21 months of experience with our organization, has demonstrated exceptional expertise and versatility across multiple projects. Their extensive skill set and hands-on experience make them an ideal candidate for delivering outstanding results on forthcoming projects

**Staffing Agent**

Matteo Floyd is going to be away starting next month, help me reallocate his work

Here's a comparison of the first three candidates:

**Zoey Ewig**  
Role: Trainer  
Match Quality: 88.9%  
Skills Matched: PSA Training

**Leonel Lester**  
Role: Trainer  
Match Quality: 88.9%  
Skills Matched: PSA Training

Assign Zoey...

**Project Assistant Agent**

Hi! I can help you analyze projects and identify missing timecards.

Are there any missing timecards for this project?

Finding Timecards ...

Describe your task or ask a question...

**Generate Project Summary**

Overview

The United Partners - Multi Element Projects is currently in progress, having started on Feb 15, 2026, and is scheduled to end on July 27, 2026. The project is managed by Preston Mayo and has a duration of approximately 73 days.

Financials

Risks

Staffing

Deliverables

**Generate RAID Next Steps**

**Generate Exclusive Summary**

**Generate Client Facing Summary**

## THE PROBLEM WITH GENERIC AI

# Your AI is only as good as the context it has

Every major software vendor now claims AI. The question is not whether a solution has AI — it is whether that AI has the context, the rules, and the architecture to be trusted with your operations. Most AI tools available to professional services firms today fail on all three counts.

They sit outside the system of record, so they have to ask other systems for context before they can act — creating a **context gap** between what the AI knows and what the business needs. They apply probabilistic reasoning to financial outputs, which means they guess at your margins, your RevRec, and your billing. And they are general-purpose tools trained on generic data — not on the complex Services Math that makes professional services businesses work.

The result is a **trust deficit**. Leaders stop relying on AI-generated outputs and build shadow spreadsheets. The AI becomes a liability rather than an asset. The **manual tax** persists. And the Maintenance Spiral of keeping the Franken-stack running consumes the innovation budget that should be funding the shift to autonomous operations.

**Every AI vendor has generative capabilities. Only Veda knows when not to use them.**

## WHAT VEDA IS

# From reactive workflows to autonomous operations

Certinia is the industry's leading **Professional Services Automation (PSA)** system, unifying sales, delivery, finance, and customer success on a single record powered by a comprehensive AI engine that acts with certainty across the entire services journey.

Veda, Certinia's AI engine for services operations, combines specialist AI agents and intelligent actions grounded in decades of institutional memory to shift services organizations from reactive workflows to autonomous operations.

### **Most AI gives you a scribe.**

A scribe summarizes. It reads your project record and writes a nice update. It can tell you what happened. It cannot change what happens next. Scribe AI is passive — it waits to be asked, produces a document, and leaves the execution to your team.

Off-platform generalist AI tools are scribes. They sit outside the system, they lack the context to act, and they guess at the outputs that matter most to your business. *They talk about work. They do not act on it.*

### **Veda is an operator.**

An operator executes. Because Veda lives natively inside the PS Cloud project record, it can update RAID logs, reallocate resources, generate success plans, and calculate margins — within the rules and business logic of your specific organization. It automates the intelligence work so your people can focus on the judgment work that drives client value.

The question isn't whether you have AI. It's whether your AI **talks about work or acts on it.**

## HOW VEDA WORKS

# Rules-bound where your P&L depends on it. Generative where your team benefits from it.

Veda applies two types of AI reasoning depending on the job — and knowing which to use, and when, is what separates it from every generic alternative.

For financial governance — margins, RevRec, WIP, resource costs, billing — Veda uses **deterministic logic**: rules-bound, 100% auditable, zero probabilistic guessing. In finance, “close enough” is a strategic liability. Veda governs these outputs the same way your business rules govern them — with certainty.

For communications and insights — project summaries, client emails, RAID extraction, success plans — Veda uses generative capabilities grounded in **professional services institutional memory**: not open-ended probabilistic generation, but generation constrained and informed by decades of services-specific context built into the system. The output is fast, personalized, and grounded — not a random guess.

***The Simple Truth:** Veda is rules-bound where your P&L depends on it. Generative where your team benefits from it. Generic AI guesses at everything — including your financial outputs. Veda governs what matters and generates what helps.*

## WHERE VEDA LIVES

# Native architecture. No context gap. No integration debt.

Veda’s authority comes from where it lives. Not bolted on from outside. Not syncing data from a disconnected system. Built natively inside Certinia on Salesforce — the same platform where your customer data, project records, financial data, and resource information already live.

This is the native architecture advantage. When Veda needs to generate a success plan, it already has the opportunity data, the account history, the support cases, and the resource availability — without making a single API call to an outside system. When it reallocates a resource, it already knows the staffing rules, the project constraints, and the financial impact.

This is what it means to have zero context gap.

WHERE	HOW	WHY
<b>Native Architecture</b> Where Veda lives	<b>Modern Architecture</b> How Veda is built	<b>Deterministic Scaffolding</b> Why it can be trusted
<p>Veda is built inside Certinia on Salesforce — full context from first bid to final invoice, no syncing required. There is no context gap because there is no integration. Your data never leaves your governed environment, and Veda never has to ask another system for context before it acts.</p>	<p>Veda combines workflow automation for rules-based tasks with Agentic autonomy for dynamic reasoning decisions. It doesn't just automate — it reasons. Workflow automation handles the deterministic plumbing. Agentic autonomy orchestrates the decisions that require judgment about context, timing, and priority.</p>	<p>A hardened architectural layer ensures Veda reasons strictly within your auditable services math and business rules. This is not learned behavior — it is logic-bound authority. Financial outputs are governed, not generated. Every RevRec calculation, every margin figure, every resource decision follows the strict rules you've established.</p>

**WHAT VEDA DOES**

## Specialist AI agents and intelligent actions. Purpose-built for services operations.

Veda delivers a new generation of specialist AI agents and intelligent actions across PS Cloud. These are not generic AI features. Each agent is purpose-built and domain-specific — hardened for the exact jobs that consume the most time in a services operation and require the most precision to get right.

Collectively, they reclaim 20+ hours per month per Project Manager — shifting your experts from data janitorial work to the judgment work that drives client value and protects margins.

Agent / Action	For	What Veda Does — and Why It Matters
Project Summary	Project Managers & Leaders	A Franken-stack of disconnected tools creates a context gap that prevents AI from being effective and drains the innovation budget.
RAID Next Steps	Project Managers	Extracts the top risks, issues, actions, assumptions, decisions, and dependencies from the RAID log in seconds, with direct links and the ability to create new items without leaving the view.

Agent / Action	For	What Veda Does — and Why It Matters
Project Tasks Summary	Project Managers	Reads the entire WBS and surfaces what's completed, overdue, blocked, and needs attention — across both Gantt and agile/kanban views — with recommended next focus areas.
Executive Summary	Project Managers	Generates a concise internal narrative of the most important project facts for stakeholder communication. Configurable via Certinia's AI Prompt Configuration Builder.
Client Facing Summary	Project Managers	Drafts a proactive client update email grounded in real project data, highlighting key risks and actions taken, and automatically filtering out internal-only data like project margin.
Resource Summary	Resource Managers	One-click AI profile of a resource's experience, secondary skills, and clients served — so staffing decisions are faster and better informed.
Project Assistant Agent	Project Managers	Monitors milestones, tasks, and timecards, flagging overdue items, scheduling mismatches, and missing entries before they become problems. Distinguishes billable from non-billable milestones and flags those starting or ending soon.
Staffing Agent	Resource Managers	Identifies all assignments of a resource being replaced, executes best-match searches using role, skills, region, and adjusted timeframes, and proposes a complete reallocation plan in minutes.

WHO IT'S FOR

# The AI trust question. Answered for every buyer.

Every stakeholder in a Veda conversation has a different version of the same underlying question: can I trust this AI with my business? Here is how Veda answers it for each of them.

CFO / Finance Leader	PS / Delivery Leader	CIO / IT Leader
<b>THE AI TRUST QUESTION</b>	<b>THE AI TRUST QUESTION</b>	<b>THE AI TRUST QUESTION</b>
"Can I trust AI-generated financial outputs?"	"Will AI actually reduce my team's workload?"	"Will this create another integration to maintain?"
<b>WHAT VEDA DELIVERS</b>	<b>WHAT VEDA DELIVERS</b>	<b>WHAT VEDA DELIVERS</b>
Veda applies deterministic logic to every financial output — margins, RevRec, billing, resource costs. 100% auditable, no probabilistic guessing. Your P&L is governed, not generated.	Veda's specialist AI agents handle the intelligence work — project summaries, RAID extraction, client communications, resource reallocation — reclaiming 20+ hours per month per Project Manager.	Native Architecture means Veda lives inside Certinia on Salesforce — no syncing, no middleware, no integration risk. A closed-loop system where your institutional data never leaves your governed environment.
<b>THE OUTCOME</b>	<b>THE OUTCOME</b>	<b>THE OUTCOME</b>
The only AI architecture a CFO can rely on to run the P&L.	From data janitors to outcome orchestrators.	Zero context gap. Zero integration debt.

PROOF

# Governed Intelligence. Verified Outcomes.

These results were not produced by AI that guessed at the right answer. They were produced by Veda — governing what needed to be governed, generating what needed to be generated, always within the business rules of the organization.

<b>Salesforce</b>	1% forecast variance managing 20 million billable hours globally — deterministic logic governing the world's largest professional services operation.
<b>Hyland</b>	695% ROI and \$4.1M in total benefits. Veda's intelligence work automated the Manual Tax that had been draining margins for years.
<b>Five9</b>	\$2.5M in incremental annual revenue by shrinking the dead zone between sales and services — Veda generating success plans the moment deals closed.
<b>BeyondTrust</b>	20% higher global utilization and a 10% increase in project margins through Veda-powered resource orchestration.

# Universal access. Continuous innovation. Predictable ROI.

A single Veda subscription provides a transparent consumption model that pairs per-user fees with usage-based scaling for a predictable path to ROI.

## No feature gates

One Veda subscription gives you full suite access. As new specialist AI agents and intelligent actions are released, users are automatically entitled access. You never have to buy your way to the latest capability.

## Deployment flexibility

Veda integrates seamlessly. There is no rip-and-replace and no forced migration. Your existing investment remains intact while you layer on additional capabilities at your own pace.

## Consumption-based scaling

Per-user monthly access fees plus usage-based scaling ensure that your cost and the value Veda delivers are permanently aligned. You pay for what you use. You grow as the value grows.

## Purpose-built and domain-specific

Veda is not a generic AI tool adapted for services. It is built for services operations, grounded in professional services institutional memory, and hardened by deterministic scaffolding. The architecture is the advantage — and it is included in every subscription.

Ready to move from reactive to autonomous?

See Veda in action at [certinia.com](https://www.certinia.com)