certinia

AI Readiness

Published: March 2026

Introduction

This white paper provides key considerations for deploying Artificial Intelligence (AI) capabilities within your Certinia environment. Thoughtful planning of the AI rollout increases the likelihood of success through higher adoption, better utilization, and achievement of targeted outcomes. AI commonly offers benefits such as greater efficiency and accuracy. This document outlines the 'what', 'why', and 'how' necessary to prepare your organization for this transformational step.

Strategic Readiness

Identify Challenges AI Can Solve

To ensure strategic alignment, our AI implementation will target high-frequency, data-rich processes where probabilistic logic offers a clear advantage over traditional "if-then" automation, specifically focusing on bottlenecks in the services delivery lifecycle. We define solvable problems as those characterized by human-intensive summarization, complex pattern matching (e.g., predictive resource staffing), and volatile forecasting (e.g., cash flow and revenue predictability).

To maintain a manageable scope, we will categorize initiatives into three areas: **Predictive**, **Generative**, and **Agentic** AI.

By anchoring our scope to these predefined value buckets and auditing the underlying data health of our Salesforce and Certinia objects, we will move from broad experimentation to a disciplined roadmap of outcomes that offer measurable Return on Investment (ROI).

Establish Measurable Success Criteria and KPIs

To successfully transition AI from a localized pilot to a core business driver, organizations must adopt a rigorous, value-based measurement framework. We define success through two primary lenses: **Operational Velocity** and **Decision Accuracy**. Operational Velocity measures the acceleration of the "time-to-insight" lifecycle, specifically quantifying the reduction in manual administrative overhead achieved through AI-augmented workflows.

Complementary to this, Decision Accuracy evaluates the precision of AI-driven outputs, measuring how consistently automated recommendations align with intended strategic outcomes.

By anchoring these criteria to specific Key Performance Indicators (KPIs) — such as **Resource Efficiency** and **Forecast Variance** — organizations can move toward a predictive operational model. This allows leadership to narrow the gap between planned and actual outcomes by leveraging analytics to mitigate risks before they materialize. Ultimately, by benchmarking these metrics against established pre-AI baselines, the transition shifts from anecdotal experimentation to a quantitative validation of AI's impact on bottom-line profitability and long-term scaling capacity.

Suggested KPI Tracking Matrix

You might include a table like this to show stakeholders exactly what success looks like:

Metric Category	Current Baseline (Pre-AI)	Target AI Metric	Strategic Impact
Resource Efficiency	5 hours to staff a project	< 1 hour via Agentic AI	Faster project kick-offs; higher utilization.
Forecast Accuracy	+/- 15% Variance	+/- 5% Variance	Predictable cash flow and investor confidence.
Management Productivity	40% time on admin tasks	20% time on admin tasks	More time to focus on higher value initiatives.
Project Health	Reactive (Manual audits)	Proactive (AI Risk Alerts)	Lower delays and fewer budget overruns.

Infrastructure

Cloud Platform

Certinia applications and AI are built on and are native to the Salesforce Platform, a web-based cloud platform. The Salesforce Platform uses several different persistence technologies, including a database virtualization technology designed by Salesforce to support multi-tenancy and customization at the tenant level.

As a cloud based solution, Certinia requires no client side installed software other than a web browser. Our solution operates in either Windows or Mac OS based environments. There is no separate hardware, database or application software required.

For mobility, Certinia utilizes the Salesforce mobile application, which is compatible with both Apple iOS and Android smartphones.

Certinia solutions are architected as **100% Salesforce-native applications**. By residing natively within the Salesforce environment, Certinia utilizes advanced **database virtualization and multi-tenant persistence technologies** designed to support enterprise-scale customization while maintaining strict data isolation.

This cloud-native approach ensures a **Zero-Footprint Deployment**: Certinia requires no proprietary hardware, local application servers, or client-side software beyond a standard web browser. This architecture IT governance, as the solution is universally compatible with both Windows and macOS environments. To support a mobile-first workforce, Certinia extends seamlessly into the **Salesforce Mobile Application**, providing full functionality and secure access across Apple iOS and Android devices without the need for additional middleware or third-party integrations.

Scalable Types of AI

Certinia delivers a comprehensive suite of embedded AI capabilities, spanning the full maturity spectrum from **Predictive** and **Generative** to **Agentic AI**. While these native features provide immediate value upon deployment, the platform's architecture is designed

for deep extensibility. Beyond standard functional offerings, organizations can leverage low-code configuration tools to tailor AI logic to their unique operational requirements. This allows customers to expand AI utility into bespoke, company-defined use cases—ensuring that the intelligence layer evolves in tandem with specific business complexities and strategic goals.

Predictive AI

Predictive AI leverages advanced **Machine Learning (ML) algorithms** to synthesize historical data residing within Salesforce and Certinia, identifying complex correlations that remain invisible to traditional reporting. While standard analytics provide a retrospective view of performance, Predictive AI — delivered via **CRM Analytics** — evaluates thousands of historical variables, such as project burn rates and early-warning indicators of budget variance. By applying these learned patterns to real-time operational data, the system generates high-confidence probability scores. This shifts the organizational posture from reactive to proactive: rather than reviewing a project failure in arrears, leadership can identify a "70% risk of budget overrun" weeks in advance. Ultimately, this transforms static data into an **active guidance system**, optimizing cash flow and resource demand forecasting with mathematical precision.

Generative AI

Generative AI, delivered through the Agentforce platform, moves beyond simple data analysis to become an active creator within your organization. Generative AI leverages Large Language Models (LLMs) that have been securely "grounded" in your unique Salesforce and Certinia data. This means the AI doesn't just provide generic answers; it understands the specific context of your projects, customers, and contracts. For example, instead of a project manager spending hours drafting a status report or a customer success plan, Agentforce can instantly generate a first draft by "reading" the project's recent milestones, risks, and financial health. By automating the administrative "heavy lifting," Generative AI allows your team to spend less time on manual documentation and more time on high-value strategic delivery.

There are a multitude of options in selecting and using an LLM. Options include available "off-the-shelf" models to being able to bring your own proprietary model. Listed below are some options, but not all, to consider:

Agentforce LLM Ecosystem Overview as of 2026

Category	Model Provider / Name	Core Strength & Purpose	Typical Use Case
Salesforce-Managed (Default Out-of-the-box)	**OpenAI** (GPT-4o / GPT-4o mini)	Versatile, high-speed reasoning and natural language drafting.	Summarizing project status or drafting customer emails.
	Anthropic (Claude 3.5 / Claude 4)	High security; runs within Salesforce's VPC for financial-grade privacy.	Analyzing sensitive Certinia ledger data or contract terms.
Salesforce In-House (Specialized)	**xGen-Code / CodeGen**	Deeply integrated with Salesforce architecture and metadata.	Generating Apex code, building Flows, and technical automation.
Bring Your Own LLM (Einstein Studio)	**External Providers** (Google Gemini, Amazon Bedrock, Azure OpenAI)	Leverages existing corporate AI investments and custom-trained models.	Connecting to proprietary data lakes or specialized industry models.

Agentic AI

Agentic AI represents AI maturity within the Salesforce and Certinia ecosystem. Unlike previous AI models that wait for a human to type a prompt (reactive), Agentic AI is **goal-oriented and proactive**. It uses a **Reasoning Engine** to think through complex, multi-step business objectives. While Generative AI uses an LLM to write a response, Agentic AI uses the LLM to **reason** through a problem and decide which tools it needs to

use to reach a goal.

Every step in agent reasoning—intent classification, topic selection, action execution, and observation—is logged with timestamped metadata. Administrators can inspect the full reasoning chain for any interaction, providing a complete audit trail from user input to agent output.

For example, instead of just drafting an email, an Agentic AI "Staffing Agent" can identify a resource shortage, search the global talent pool in Certinia for a match, verify project margins, and independently book the consultant — only notifying a human once the task is complete. It essentially functions as **digital labor**, capable of perceiving changes in your data, like a project going over budget, and triggering a series of cross-system actions to resolve the issue without being asked.

Agentic AI allows your organization to scale operations without a proportional increase in administrative headcount. Selecting an LLM to use Agentic AI will also be required.

Data Readiness

Data readiness is the primary predictor of AI performance. Because AI models do not "think" independently but rather identify patterns within existing datasets, the integrity of your Salesforce and Certinia records directly dictates the quality of the AI's output. High data readiness means your information is **accurate, complete, and unified**. If project logs are missing, or if customer names are duplicated across systems, the AI will provide "**hallucinated**" or misleading insights — a phenomenon known as "Garbage In, Garbage Out." To achieve readiness, we must move beyond siloed data and focus on a "Single Source of Truth." This involves auditing our core objects for "data debt" (outdated or empty fields) and ensuring our data architecture is clean enough for the AI to understand the context of our business.

Ultimately, a sophisticated AI agent is only as intelligent as the data it is permitted to process. Our primary strategic objective is to treat our data as a **high-value corporate asset**, continuously groomed for machine consumption. By ensuring our data architecture is clean and contextually rich, we provide the AI with the "ground truth" necessary to drive

meaningful business outcomes.

Data Residency and Compliance

When implementing AI, especially within a platform like Certinia that handles sensitive financial and project data, **Data Residency and Compliance** are the legal guardrails that ensure your innovation doesn't create a liability. It is the practice of ensuring that your data stays where it belongs and is processed only under strict regulatory standards.

As you insert AI into your core operations, maintaining the integrity and location of your data is of utmost importance. In a global regulatory environment governed by frameworks like GDPR and industry-specific financial standards, we must ensure that our use of AI does not inadvertently leak sensitive information across borders or into public models. All data in transit is encrypted via TLS, and data at rest via Salesforce Shield Platform Encryption where enabled.

By leveraging the **Einstein Trust Layer**, we ensure a Zero-Data Retention policy: while our data is used to ground AI responses, it is never stored by external LLM providers or used to train their global models. The Trust Layer controls — zero data retention, PII masking, toxicity detection, and audit trail — form the core governance layer for all AI interactions.

Furthermore, our strategy prioritizes **Regional Compliance**, ensuring that data processing occurs within the specific geographic boundaries required by our contracts and local laws. By establishing these clear digital boundaries and automated auditing trails, we can confidently deploy AI agents that respect our privacy commitments while delivering high-value insights, effectively turning compliance from a hurdle into a competitive advantage.

In addition to system controls there will also be a Human-in-the-loop control. Agents escalate to a human when they encounter queries outside defined topics, responses that trigger toxicity thresholds, or configurable uncertainty boundaries. Full conversation context is preserved on handoff. Administrators define escalation behavior declaratively — specifying which topics require human approval, which actions are restricted, and what fallback behavior applies.

Connection of Data

The Salesforce platform provides a flexible foundation for AI deployment, capable of utilizing both native and external data sets. To maximize the intelligence of your models, off-platform data can be seamlessly integrated through several strategic pathways:

- Salesforce Data Cloud: The premier solution for high-scale data unification.
- Data Lake Integration: Direct connectivity to existing enterprise data repositories.
- iPaaS & API Frameworks: Point-to-point connections for specific, real-time data requirements.

Technical Reference: For a comprehensive technical breakdown of these connectivity methods, please refer to the *Certinia Integration White Paper*.

When deploying **Agentforce AI**, it is a best practice to enable the **Data Cloud Connection**. While the full Data Cloud suite is not a prerequisite for basic functionality, activating this connection is essential for the advanced management, monitoring, and observability of agent interactions.

Data Quality

To achieve true AI readiness, organizations must look beyond surface-level data cleanliness and evaluate the **representativeness and integrity** of their datasets. While data quality — defined by accuracy, completeness, and consistency — is the mandatory foundation for reliable AI, it is only the first step. To prevent flawed forecasts and operational errors, we must also mitigate the systemic risk of **Data Bias**.

Data bias occurs when historical datasets reflect past human prejudices or systemic gaps rather than objective reality. Because AI models identify patterns in the past to predict the future, they risk creating a "feedback loop" that codifies previous inefficiencies.

- **The Risk:** If historical staffing data in Certinia favors specific regions or lacks diversity, an AI agent may inadvertently prioritize those same patterns.
- **The Consequence:** This leads to skewed resource recommendations and unfair performance benchmarks that reinforce — rather than correct — past imbalances.

Our deployment strategy is anchored in the **Salesforce Trusted AI Principles**, ensuring

that every model remains:

- **Responsible & Accountable:** Maintaining human oversight of AI-driven outcomes.
- **Transparent:** Providing visibility into the "why" behind AI reasoning.
- **Empowering & Inclusive:** Augmenting human potential without exclusion.

Technical safeguards are integrated into the architecture to defend against these risks. We leverage the **Salesforce Einstein Trust Layer** for active toxicity detection and output filtering. By implementing rigorous data grooming and fairness audits within our Certinia and Salesforce environments, we ensure that AI-driven decisions are equitable, turning data into a strategic asset rather than a historical liability.

Change Management

Technical readiness is only half of the equation. A successful transition to an AI-augmented organization is as much a **cultural evolution** as it is a digital one. To realize the full potential of Certinia's AI capabilities, an organization must ensure its human capital is both engaged and empowered.

Effective change management for AI diverges from traditional software implementations. It requires a strategic shift toward **workforce evolution**, where AI is not viewed as a replacement for staff, but as a sophisticated tool for enhancement. Success is measured not just by technical integration, but by **AI Literacy**—the baseline competency required for teams to effectively:

- **Prompt:** Interface with models to extract high-value insights.
- **Audit:** Critically evaluate AI outputs for accuracy and bias.
- **Collaborate:** Treat digital agents as "force multipliers" that handle administrative burdens.

In this new landscape, employees transition from manual executors to **strategic orchestrators**. By fostering a culture of continuous learning and transparent communication, organizations can mitigate "automation anxiety" and instead build a workforce that views AI as a collaborative peer. Ultimately, the most competitive firms will be those that align their technological roadmap with a robust strategy for human enablement.

Adoption and Champions

The most effective Champions are not necessarily the most technical; rather, they are the individuals with the deepest domain expertise and the highest level of trust among their peers. Their role is to translate AI capabilities into departmental wins. By equipping these Champions with early access and specialized training, we establish a scalable support structure that operates at the "speed of business."

This human-centric approach transforms employees from passive users into **active innovators**. By investing in a decentralized network of experts, the organization ensures that its AI transition is not a one-time event, but a long-term shift toward a confident, upskilled workforce. Ultimately, this cultural resilience is what translates a Certinia investment into a sustained competitive advantage.

AI Literacy and Enablement

AI readiness is predicated on a workforce that possesses the **AI Literacy** required to interpret, challenge, and refine machine-driven insights. In an agentic environment, enablement must go beyond traditional software training; it must focus on "Grounding"—the ability of human operators to anchor AI outputs in real-world business logic.

To extract maximum value from systems like **Agentforce**, teams must be equipped with three core competencies:

- **Contextual Grounding:** Ensuring AI insights align with specific departmental realities.
- **Risk Identification:** Recognizing "hallucinations" or logical gaps in machine-generated data.
- **Prompt Optimization:** Mastering the iterative process of directing AI to produce high-fidelity results.

A critical component of this enablement is the **Continuous Feedback Loop**. Daily employee interactions provide the essential "ground truth" data necessary to measure AI accuracy and relevance.

By establishing structured channels for users to report when an AI-generated project forecast lacks context or a staffing recommendation is misaligned, we transform every employee into a **system contributor**. This decentralized oversight ensures that the AI's evolution is directly informed by frontline expertise.

This dual approach of proactive education and structured data gathering ensures that as the AI becomes more sophisticated, the workforce becomes more capable. This creates a **self-reinforcing cycle of trust and excellence**, where human intuition and machine intelligence evolve in tandem to drive long-term organizational success.

AI Skill Maintenance

In a dynamic ecosystem like Salesforce and Certinia, AI maintenance is a dual responsibility. Organizations must continuously update technical guardrails while simultaneously refreshing the human skill sets required to oversee them. This ongoing synchronization ensures the system remains a secure, high-performing asset.

To maintain safety and compliance, we implement automated guardrails that function as **Agent Policy**. These boundaries ensure that AI agents operate strictly within predefined parameters, preventing them from exceeding their authority or deviating from corporate compliance standards.

Key technical maintenance includes:

- **Dynamic Data Masking:** Continuously updating rules within the **Einstein Trust Layer** to protect PII (Personally Identifiable Information).
- **Permission Set Audits:** Regularly refining access levels to ensure agents only "see" and "act" on data appropriate for their specific role.
- **Logic Calibration:** Adjusting the agent’s reasoning paths as business workflows evolve.

AI skills should be treated as a **perishable commodity**. As Large Language Models (LLMs) are updated and Agentforce capabilities expand, the workforce must engage in iterative learning cycles. This ensures teams can master new prompting techniques and understand the shifting reasoning logic of the latest models.

By treating guardrails as dynamic boundaries and skills as evolving assets, we protect the

organization from **AI Drift**—the gradual loss of alignment between AI outputs and business objectives.

This rigorous approach to maintenance ensures that our digital assistants remain safe, accurate, and perfectly aligned with our evolving business goals. In the Agentic era, the most successful organizations will be those that prioritize **operational endurance** as much as initial deployment.

Further Information

Please visit: https://www.certinia.com/trust/ and https://www.certinia.com/privacy/ for more information on our policies and procedures.

Disclaimer

This document is in response to a request for information and discussion purposes only. It does not contain any binding commitments.
Certinia reserves the right to amend or modify this information upon receipt of additional information following its submission. We further advise that, due to the rapidly evolving nature of our solution, as well as our technology, security and business environments, the information contained within our response is subject to change.
Upon completion of the appropriate due diligence should Certinia be selected as your vendor of choice, we will sign Certinia's form of Master Subscription Agreement, and Consulting Services Addendum if applicable, with reasonable changes to be negotiated in good faith. Those agreements will contain the definitive terms and conditions governing our relationship. Any legal terms or conditions in your request or related documents that are not expressly included in the definitive, signed agreement will be void and without effect.

Statement of Confidentiality

This document contains proprietary and confidential trade secrets and information of Certinia and is provided on the express condition that you maintain this information as confidential and not disclose its contents to any third party except as may be required by law or regulation, in which case adequate notice shall first be given to Certinia, Inc. to permit any objection to such disclosure. This information may be disclosed to employees, attorneys, or accountants for the purposes of evaluating a business relationship with Certinia.