



Mphasis

SILVERLINE

The Next Applied

**ACTIVATE YOUR DATA
TO POWER AI ANALYTICS:**

A Tableau Next Implementation Guide

Table of Contents

Intro	2
What Tableau Next Actually Does	3
Tableau Next Use Cases	4
Tableau Pulse vs. Tableau Next: Understanding the Difference	5
Structure Your Data to Optimize Tableau Next Implementation	6
The Role of the Center of Excellence	7
Three Steps to Prepare Your Data for AI Activation	8
The AI Activation Checklist	9

Most Salesforce organizations treat AI as something to switch on. The ones that succeed treat it as something to architect.

Salesforce's next-level AI analytics solution, **Tableau Next**, is built on its core platform and works best when the data beneath it is ready. Our Tableau Next Implementation Guide breaks down what Tableau Next actually does, explores the role data readiness plays in your implementation, and outlines an action plan for AI activation.

MIT research puts a precise number on what many practitioners already know: roughly 95% of AI pilot projects stall before they deliver real business impact. They impress in demos, fade after the first quarter, or collapse when scaled beyond the team that built them. The organizations that make it do not treat AI as an experiment. They treat it as infrastructure – embedded into the systems where decisions are actually made, not running alongside them in a parallel tool.

The MIT research identifies a counterintuitive differentiator: the successful 5% embrace friction rather than avoid it. Friction — redesigning workflows, retraining teams, rebuilding feedback loops — is not the obstacle to AI success. It is the mechanism of it. Organizations that skip this work get AI that generates output nobody acts on. Organizations that do this work get AI that changes how decisions are made.

What follows is what it means specifically for Salesforce organizations through the lens of analytics. Salesforce has made a deliberate architectural choice: Tableau Next is its AI platform for analytics. Not a bolt-on feature, not a third-party integration, but the native answer to the question of how Salesforce organizations turn governed data into intelligent action.

What Tableau Next Actually Does

Tableau Next is an agentic analytics platform built natively on Salesforce, powered by Agentforce, connected to Data 360, and designed to operate on the governed semantic layer your organization built during data readiness. The difference between Tableau Next and any analytics tool you have used before is not the interface. It is the posture. Tableau Next does not wait to be asked. It monitors, reasons, surfaces, and acts.

Tableau Next's two primary AI agents, Concierge and Data Pro, work within Salesforce workflows:

- Concierge handles conversational analytics: a Sales leader asks a complex multi-part question about deal risk across a territory and gets a reasoned answer, not a report request.
- Data Pro handles deep analytical work: probing datasets, identifying patterns, surfacing what the data implies rather than just what it shows.

Neither agent is a search box. Both require that the underlying data is clean, defined, and semantically governed.

Tableau Next Use Cases



Pipeline Health and Deal Risk (Sales)

A Sales VP opens Salesforce on a Monday morning. Tableau Next has already surfaced what matters: three deals that have gone quite longer than the historical close average, one opportunity with no economic buyer activity in three weeks, two forecast categories that have shifted twice this quarter.

No report was requested. No analyst was pulled in. The VP does not start the week reviewing data. They start acting on intelligence.



Proactive Case Escalation (Service)

A Service director finds out about SLA risk before a breach happens. Tableau Next monitors open cases against resolution time, account tier, complexity, and workload continuously. When a trajectory looks like a miss, it recommends an action inside the Salesforce workflow: reassign, escalate, or engage the customer before the clock runs out. The metric that shifts is not just resolution time. It is customer trust, which is not recoverable once it is lost.



Forecast Variance Detection (Finance and RevOps)

End-of-quarter forecast reconciliation is one of the most painful recurring processes in any Salesforce organization. Nobody fully trusts either number. Tableau Next changes this by monitoring pipeline signals throughout the quarter and surfacing variance indicators as they emerge, not just at submission.

When a segment tracks above or below its historical close rate, Tableau Next surfaces which deals are driving it, what the signals look like, and what comparable deals did in prior quarters. The forecast becomes a discussion, not a negotiation.



Churn Signal Detection (Customer Success and RevOps)

Churn is rarely a surprise to the data. Tableau Next monitors account health across Salesforce objects — case volume, product usage from Data 360, renewal timelines, and CS engagement frequency — building a composite risk picture continuously.

The output is not a churn score someone checks monthly. It is a proactive signal: four tickets in thirty days, active users down 40%, renewal in sixty days, CSM has not logged a touchpoint in three weeks. Act now. That specificity, delivered at the right moment inside the workflow, is what separates Tableau Next from a reporting tool.

Tableau Pulse vs. Tableau Next: Understanding the Difference

Tableau Pulse and Tableau Next are not two versions of the same thing. Understanding the distinction matters, because organizations that confuse them will underinvest in one or over-expect from the other.

Tableau Pulse is a metric-monitoring feature within Tableau Cloud. It tracks KPIs, surfaces natural language insights, and delivers updates through Slack and email. It makes data accessible to people who do not build dashboards. It is reactive, answering questions about what is happening with a metric. Pulse is something you turn on to improve visibility.

Tableau Next is the platform. It encompasses Pulse but extends into Agentic AI through Agentforce, with AI agents — Concierge for conversational analytics and Data Pro for deep analytical reasoning — that do not wait to be asked. They monitor, analyze, surface what matters, and take actions within Salesforce workflows. They are powered by Data 360 and operate on the semantic layer the organization has built. Tableau Next is proactive.

Pulse is what you consume, Next is the environment in which intelligence acts. One is a feature you enable. The other is a platform you architect your operating model around.

That distinction separates incremental improvement from meaningful organizational change. And it explains why the full data readiness foundation — the semantic models, the governed data, the CoE, the AI governance policy — is what Tableau Next requires. Pulse can operate with less. Next demands more because it does more.



Pulse is what you consume, Next is the environment in which intelligence acts. One is a feature you enable. The other is a platform you architect your operating model around.

Structure Your Data to Optimize Tableau Next Implementation

Most organizations approach AI analytics the same way: choose a tool, secure budget, launch. Then the numbers don't match. Teams argue over definitions. Trust erodes. The AI was confident, but it was reasoning over contradictory, ungoverned data.

Inside Salesforce organizations, this follows the same arc. AI systems are pattern-seeking engines. They don't fix bad data, they scale it. Metric definitions that drift across teams, Salesforce reports with overlapping filters and hidden joins, missing business context like territory models and fiscal calendars — these don't disappear when you add AI. They become automated, amplified, and harder to unwind.

The risk of building on a weak data foundation

AI is not a tooling decision; it is a data maturity outcome. Jump to AI without a data readiness foundation and you will automate inconsistency, amplify ambiguity, and scale risk.

In Salesforce environments, this shows up in specific, familiar ways:

- “Bookings” changes based on which Opportunity filters are applied
- “Closed Won” depends on whether overrides, partial deals, or late-stage churn are included
- Picklist evolve over years, undocumented and inconsistently enforced
- Territory logic lives in someone's head
- Fiscal calendars differ between reports

For organizations also working across Data 360, the complexity compounds further. Harmonizing CRM data with external sources introduces new inconsistencies if the definitions were never settled inside Salesforce to begin with. When AI is introduced on top of this, the result is not clarity, it is the erosion of trust. Once trust is lost, adoption collapses.

The fragmentation challenge

A standard Salesforce org is rarely a single system, but an ecosystem.

Sales works in Opportunities and Forecasts. Service lives in Cases and SLAs. Marketing operates in Campaigns and attribution models. RevOps stitches performance together. Finance reconciles everything after the fact.

Each group has built reports optimized for its own decisions, often under deadline pressure, often without clear governance. Ask three people in this environment what “Closed Won” means and you will get three different answers. All three answers are defensible in isolation. None of them coexist safely inside an AI system.

Before any AI initiative can succeed, an organization must answer a simpler and more fundamental question: Do we have a single shared, trusted version of our data?

The Role of the Center of Excellence

A Center of Excellence (CoE) is the organizational structure that keeps data AI-consumable over time. Its purpose is to define the truth, maintain the truth, and protect the organization from the cost of losing it. Without this structure, your data quietly degrades as teams evolve, leadership changes, and Salesforce configurations drift.



Metrics and KPI stewardship

The CoE defines and ratifies KPIs with the business, publishing metric contracts that document the full logic: filters, time windows, attribution rules, and deprecation policies. When a metric changes, the change is deliberate, documented, and communicated. Silent drift is eliminated.



Data models and semantics

The CoE owns the canonical data model for Sales, Service, and Marketing, managing hierarchies, lifecycle states, and the semantic layer that BI tools and AI systems consume. This is the organization's shared language for its own business.



Data quality and governance


The CoE establishes data contracts with source systems and monitors quality dimensions (completeness, consistency, timeliness, validity) with defined thresholds and accountable owners. Quality is not assumed, it is measured.



AI governance and risk

The CoE approves AI use cases before activation, requires full traceability for every AI output, and defines human-in-the-loop checkpoints and rollback procedures for any automated action. AI without governance is a liability.

Run the CoE with a product mindset. Treat data assets as living products that are owned, versioned, and continuously improved. Measure success not by the documentation it produces, but by the business outcomes it enables: faster pipeline movement, improved win rates, shorter case resolution times.



Three Steps to Prepare Your Data for AI Activation

Becoming data-ready is a journey of maturity. The organizations that get it right advance through the three deliberate Crawl/Walk/Run phases outlined below, each building the trust and infrastructure the next one requires.

Skipping a phase does not accelerate the timeline. It guarantees a return to the beginning. The most common sequencing mistake is attempting to build governance (the Walk step) before the data itself has been audited and agreed upon. Organizations that do this find themselves governing definitions nobody trusts, which is indistinguishable from having no governance at all.

CRAWL: Collect. Refine. Clean.

Before any platform enters the conversation, the data itself must be trustworthy. This phase is unglamorous and non-negotiable. Data stewards and subject matter experts across Sales, Service, Marketing, RevOps, and IT work together to reach a shared, signed-off understanding of what the data means and who is accountable for its quality.

A Center of Excellence can drive this, but it does not need to be a large or formal body. The right people already exist in your organization. The work is to give them ownership of this problem and the mandate to resolve it.

That means auditing Salesforce fields, picklists, and reporting logic for inconsistencies; identifying data quality gaps with named, accountable owners; and defining the core KPIs that the business will stake its decisions on, signed off by the stakeholders who use them. This work will surface disagreements that have existed for years. That is the point. In practice, this phase exposes political friction as often as technical debt.

Data readiness is not about forcing agreement on day one, but about establishing ownership, escalation paths, and a shared commitment to resolve ambiguity deliberately rather than letting it persist invisibly. Better to surface them in a data audit than to discover them embedded in an AI output that leadership has already acted on.



WALK: Document. Unify. Govern.

Before AI can deliver reliable insights, there needs to be a stable analytics layer that the business already trusts and uses. This is the proving ground where metrics are validated, dashboards tell consistent stories, and teams develop confidence in what the numbers are showing them.

This phase is about institutionalizing definitions. KPIs must be documented with their full logic: the filters applied, the time windows used, the attribution rules followed, and the edge cases handled. Master data principles need to be applied so that Account hierarchies, territory structures, and product taxonomies are consistent across every report and every team.

The smarter move during this phase is to build on what already exists. If teams are using Tableau, deepen it. If CRM Analytics or Tableau is in place, build on it. The goal is not to introduce new tools, but to lock in the core BI metrics, stabilize the analytical stories the business relies on, and enforce access controls and row-level security that match organizational reality. If dashboards are multiplying faster than decisions improve, the organization is stuck in Walk.

An AI layer is only as credible as the dashboards and metrics beneath it. That stability is what AI will build on.



RUN: Activate. Govern. Act.

The Walk phase stabilizes analytics and reveals where AI will have the most immediate impact. By the time an organization reaches the Run stage, the opportunities are already visible. The question shifts from whether to introduce AI to how to govern it well.

A semantic layer must be built: the layer that tells AI not just what the data is, but what it means. Entities must be defined. Hierarchies must be established. Time logic must be encoded. Fiscal calendars, territory models, and product hierarchies must be explicit so that AI can reason over them correctly rather than approximate them incorrectly.

And governance must travel with AI. Every AI output needs traceability. Every automated action needs a human-in-the-loop checkpoint. Every use case needs an approval process and a rollback plan. The organizations that skip these steps find themselves explaining to leadership why the AI recommended something the business would never have sanctioned.



The AI Activation Checklist

Use this as your phase-by-phase readiness guide. Complete each stage before moving to the next. Each unchecked box is not a gap to paper over, but your next investment priority.



CRAWL

Clean the data

- ☐ Core KPIs defined and signed off by business stakeholders
- ☐ Data quality gaps identified with accountable owners assigned
- ☐ Salesforce fields, picklists, and reporting logic documented and agreed upon



WALK

Structure and govern

- ☐ Governed BI datasets built in CRMA or Tableau with documented lineage
- ☐ Business metrics stabilized and consistently telling the right story across dashboards
- ☐ Access controls and row-level security mapped to your organizational structure



RUN

Bring in AI

- ☐ Semantic layer in place exposing entities, hierarchies, time logic, and business context
- ☐ AI governance policy defined: approvals, audit trail, explainability standards, and fallback procedure
- ☐ Human-in-the-loop decision points established for every automated action

Conclusion

Organizations do not fail at AI because the technology is wrong. They fail because they expect it to operate without changing the organization around it. Tableau Next makes that reality visible, and when the foundation is in place, something shifts.

Teams stop waiting for reports and start acting on signals.

Analysts stop building dashboards and start validating AI-generated narratives.

Leaders stop asking what happened and start acting on what should happen next.

The workforce does not simply have more information, it has better information, at the moment of decision, embedded in the workflow where that decision is made.



About Mphasis Silverline

As part of Mphasis, Silverline tailors digital transformation solutions to meet your specific needs by leveraging insights acquired through 10+ years in the business. We've completed thousands of Salesforce engagements informed by real-world expertise gained across Media and Entertainment, Financial Services, and Healthcare industries. From strategic planning and implementation to managed services, we guide clients through every phase of their journey, enabling continuous value with the Salesforce platform. We also offer CalendarAnything, a popular scheduling application on the AppExchange, as well as industry-proven accelerators.

To learn more, visit www.silverlinecrm.com