zensar

Orchestrating the Best of All Worlds with Databricks: Turning CRM Migration into a Future Ready Intelligence Strategy for Life Sciences

White paper

An RPG Company

Executive summary

Veeva's published roadmap—moving Veeva CRM off the Salesforce platform with legacy support ending December 31, 2029—forces a decision now. The risk is treating this as a cutover: without a clear data and governance strategy, organizations inherit fragmented ownership, higher reconciliation costs, slower decisions, and avoidable audit risk. In the first 90 days, Zensar aligns stakeholders on the target operating model and success metrics, stands up the first set of governed data products on Databricks, and delivers 1–2 priority use cases with baseline ROI/TCO measurements.

Introduction

This paper is about a decision point many life sciences organizations face: CRM change coinciding with rising expectations for governed data, advanced analytics, and generative AI. The case is simple—keep the system of engagement strong, and add a system of intelligence that turns engagement data into reusable, trusted enterprise insight.

It is written for commercial, medical, clinical, quality/compliance, supply chain, and IT leaders—especially those responsible for CRM strategy, data governance, and outcomes tied to On-Time In-Full (OTIF), launch readiness, and regulatory confidence.

You will not find platform "winner" rhetoric here. The goal is to frame decisions through a business lens: where to invest to lower total cost of ownership, accelerate time-to-insight, and reduce compliance and operational risk—without compromising field execution.

The story starts where these decisions usually start: with pressure you can feel in the business.

Here's what clients who've done this well have achieved:

Metric	Typical target
Reduction in manual reconciliation effort	20%–40%
Time-to-insight for priority questions	Days - hours
Run cost for reporting and integrations (TCO)	10%–25% reduction
Duplicate pipelines/reports retired through standardization	15%–30%
Data assets with policy-based access + lineage	70%–90% coverage for scoped domains

Note: Targets vary by baseline, scope, and regulatory context; use these as planning ranges and validate during discovery.

The context

Most major technology investments in life sciences don't start with technology questions. They start with business pressure.

A commercial leader studies a pre-launch plan and realizes the field model will need to flex faster than traditional planning cycles allow. A CIO looks across multiple phase II and phase III trials and feels the pressure of rising data volume, validation demands, and timelines that refuse to slow down.
A manufacturing leader balances clinical supply across borders, languages, and regulations, knowing that "on time" is no longer good enough. OTIF means delivering the right product, in the right quantity, at the promised time. It helps prevent shortages, delays, and waste.

When OTIF performance slips, the cost shows up quickly: site disruptions, expediting and waste, missed revenue, and—most importantly—risk to patients and trial timelines. In many organizations, the hidden driver is not effort but visibility: insights arrive late because data is fragmented, and teams spend time reconciling instead of acting. Just as important, fragmented data and manual work make it harder to produce consistent, audit-ready evidence when quality and regulatory scrutiny increase.

In life sciences, OTIF is not just a supply chain metric. Patients need therapy on time. Clinical trials need the right supply at each site. Regulators expect reliable, traceable delivery. Partners and providers depend on it, too.

Zensar has seen this play out directly. For a global clinical diagnostics leader operating in 127 countries, manual ordering of consumables created delays and billing issues. Zensar automated the integration between ordering, inventory tracking, and ERP to remove manual handoffs and reconciliation. Outcomes included faster, more reliable order booking, improved availability, and higher satisfaction—because visibility and data flow were fixed at the source.

That is the same problem—and the same opportunity—inside CRM migration decisions today.

These are day-to-day realities in life sciences—and why CRM migrations (and migration decisions) between 2026 and 2029 matter. On the surface, this looks like a platform transition. In practice, it is a moment when an organization is willing to look under the hood. That creates an opportunity not just to replace a system, but to rethink how insight, intelligence, and value move across the enterprise for the next decade.

Veeva and Salesforce: Proven foundations that should stay strong

Let's start with clarity and respect.

Veeva Vault CRM and Salesforce Life Sciences Cloud are two of the strongest, most credible, forward-looking paths for life sciences engagement. Both support compliant, structured interaction with healthcare professionals and stakeholders at scale, with the auditability and operational rigor the industry requires.

Veeva Systems has set the standard for life sciences engagement by embedding compliance, content governance, and operational discipline directly into the workflows that commercial and medical teams rely on every day. Salesforce brings unmatched flexibility and ecosystem strength, enabling life sciences organizations to seamlessly connect CRM with enterprise-wide data, analytics, and customer engagement capabilities at scale.

Our recommendation is additive: keep the CRM strong and let it do what it does best. Bring Databricks into the equation as a governed intelligence layer that extends what Veeva and Salesforce already enable—without disrupting field execution or established compliance workflows.

These platforms do exactly what CRMs should do—and they do it exceptionally well. They power engagement, support execution, and keep the field productive, compliant, and focused. That strength is not in question.
What has changed is how important the data generated by these systems has become beyond CRM itself.

Two simple examples:

- A demand signal model that blends interaction activity with order history, inventory, and shipment status—plus emails and weekly reports—to flag risk early and keep OTIF stable.
- A governed "medical truth" dataset that brings together structured records with slide decks, PDFs, and image-based evidence from SaaS, on-prem, and cloud sources—so insights can be reused with consistent lineage and controls.

Engagement data is no longer "commercial data" — it's enterprise intelligence

For years, CRM engagement data lived comfortably within commercial operations. It answered questions about activity, reach, frequency, and adoption.
That era is over—today, engagement data has become enterprise intelligence.

Clinical operations teams increasingly look at engagement patterns to understand investigator responsiveness, site readiness, and operational bottlenecks. Medical communications teams rely on engagement and content signals to ensure scientific accuracy, consistency, and responsiveness across channels. Pharmacovigilance organizations face growing pressure to detect early safety signals, and engagement data often provides valuable context when correlated with narratives and inquiries.

Supply chain leaders use engagement as an early signal of demand to improve OTIF performance and reduce waste. Quality and compliance teams need clear, traceable data flows to meet GxP requirements without slowing innovation. R&D organizations need secure, interoperable data foundations that support collaboration while protecting intellectual property.

None of these use cases was the original reason CRM data was captured—and yet many now depend on it.

At the end of the day, regardless of industry, organizations must stay profitable and resilient. One of the most practical ways to do that is to use the data they already produce to react and adapt faster. Engagement data is one of the richest sources of value—if it can be used beyond a single system.

A simple metaphor: You don't run the kitchen in the dining room

Think about the best restaurant you've ever visited. The dining room is where the experience happens: polished, fast, and consistent. That's where the guest interacts with the brand.

But the dining room is not where ingredients are prepped, recipes refined, or quality controlled. That work happens in the kitchen—designed for scale, repetition, and precision.

Veeva Vault CRM or Salesforce Life Sciences Cloud is the dining room. It should remain focused on engagement, execution, and user experience.

Enterprise intelligence is the kitchen. It's where data from many sources is refined, governed, and transformed into insight that feeds the entire organization.

When organizations try to do everything in the dining room, service slows down. When they build the kitchen the right way to support the dining room, the whole experience improves.

Why the 2026–2029 window matters

The industry often talks about 2030 as a deadline. Leaders know the real forcing function is earlier.

Veeva has announced its roadmap to move Veeva CRM off the Salesforce platform to Vault CRM, with legacy support ending December 31, 2029. That timeline makes proactive planning a necessity.

The hard part is not copying data or re-creating screens. It is maintaining trust while integration, analytics, validation, governance, localization, and operating models change.

Teams that define governance and data ownership before go-live spend far less time cleaning up after.

This is also the moment to challenge the operating model behind your CRM. In "stay vs. move" talks, don't just compare platforms—test whether your model fits today's market and can scale. Many teams are running a model built years ago, and the gaps show up slowly: missed handoffs, manual work, and late decisions. Treat this as "just another migration," and you may still win the cutover. Then the business asks harder questions.

We've seen this movie before. Not every organization gets it right the first time.

Why did uptake differ across regions? What signals suggest risk before launch? Where are we over- or under-supplying?

If those answers require manual reconciliation, brittle logic, or weeks of debate, the foundation was not strong enough.

Weak foundations rarely collapse on day one. They fail when the business outpaces the architecture.

We've also seen what changes when the foundation is built early. In one HIPAA-regulated environment, a unified Databricks platform with audit logging, lineage, and role-based controls cut processing time from hours to under 20 minutes. It reduced the infrastructure footprint—creating a foundation the organization could extend without rework.

It's avoidable.

What decisions would you make differently if your leadership team could trust a single set of answers—globally, in near real time?

Why CRM-only strategies start to strain

CRMs are optimized for engagement—responsiveness, workflow control, and compliance at the point of interaction—and Veeva and Salesforce do that extremely well.

What's changing is the demand placed on data: leaders expect governed insight across structured and unstructured sources and functions, at enterprise speed. In that environment, a CRM-only strategy starts to strain—not because the CRM is weak, but because the job has changed.

As organizations push CRM platforms to absorb these responsibilities, complexity grows. Data models become rigid. Reporting logic becomes tightly coupled to schemas. Manual processes creep in to bridge the gaps.

That's where friction shows up. Progress stalls.

The issue is rarely the CRM. It is fragmented data ownership and delivery that treat enterprise intelligence as a series of disconnected tasks rather than a business imperative with clear accountability. In life sciences, that fragmentation has a second cost: it complicates GxP-aligned governance and can weaken audit readiness when expectations include traceable lineage, controlled access, and documented change.

Executives need numbers: agree up front on a small set of outcomes—e.g., reconciliation reduction (%), time-to-insight (days to hours), and lower run cost for reporting and integrations (TCO).

If this isn't treated as a business imperative with clear ownership, what breaks is predictable: metrics diverge by region, reconciliation becomes the operating model, and AI efforts get stuck in pilots because teams can't reuse governed data and approved logic.

This is not a failure of Veeva or Salesforce. It reflects the natural limits of any CRM when asked to carry enterprise-scale analytics, unstructured data, and cross-domain governance.

The future requires CRM to stay excellent at engagement. It also needs a trusted companion—an intelligence layer that helps it scale and handles analytics, AI, and cross-enterprise insights.

Enter Databricks: The system of intelligence

Databricks fits this moment not because it stores data, but because it changes how organizations create, govern, and reuse intelligence at scale.
Rather than competing with systems of engagement such as Veeva Vault CRM or Salesforce Life Sciences Cloud, Databricks serves as an independent system of intelligence. In this place, data from many domains is standardized, enriched, governed, and turned into durable data products that serve the enterprise.

Most importantly, it standardizes meaning. Definitions, metrics, and features are codified and governed so insights are repeatable—not dependent on individual interpretation—and improve as data products mature.

At enterprise scale, Databricks supports:

- The refinement of engagement data into reusable intelligence
- Secure interoperability across clinical, commercial, safety, manufacturing, and research systems
- Analytics and AI workloads that demand both performance and explainability

At the center of this model sits Unity Catalog, providing a consistent control plane for governance, lineage, access control, and policy enforcement across analytics and AI. For regulated work, that consistency helps teams produce audit-ready evidence—who accessed what, what logic ran, what changed, and when—supporting GxP-aligned processes and 21 CFR Part 11 expectations without relying on manual, after-the-fact documentation.

What GenAI actually looks like in this context

GenAI's promise in life sciences is clear; delivery is harder because outputs are only as trustworthy as the data underneath. When definitions are not standardized, lineage is not tracked, and access controls are inconsistent, outputs cannot be explained or reused—and they stay in pilots.

Zensar has already built past this. For a global pharmaceutical company operating in more than 50 countries, we deployed a multi-agent RAG platform to automate literature search and extraction across patents, clinical dossiers, regulatory filings, and safety documents. The system delivered citation-backed answers — every response traceable back to its source — so researchers and regulatory teams could trust the output and use it in submissions. The result: up to 50 percent time savings in data extraction and summarization for analysts, 30 percent faster discovery timelines, and up to $1.5 million in annual savings potential. Regulatory confidence improved because the system couldn't answer without showing exactly where it came from.

That is what ZenseAI.Data brings to the CRM intelligence context. Not a generic AI layer — a governed retrieval and reasoning capability built on standardized data products, with lineage back to source, role-based access, and outputs that can be explained and audited.

In practice, this means three things for life sciences organizations going through CRM migration:

For commercial operations: a governed Q&A experience over standardized engagement and territory performance data — so a regional director asking "what changed this week and why?" gets an answer grounded in approved definitions, not someone's spreadsheet interpretation.

For supply chain: GenAI-assisted triage that reads late-shipment summaries, demand signals, and constraint flags from both structured planning systems and unstructured emails and reports — so teams act before OTIF slips, not after.

For medical and regulatory affairs: rapid synthesis of medical inquiries, engagement signals, and document content into decision-ready briefs — with every claim linked back to an approved source, supporting both speed and traceability.
The governance is what makes AI usable at scale in a regulated environment—and that is what Zensar delivers.

Why Zensar turns this into outcomes

Technology alone does not deliver outcomes. Execution does — and in life sciences, execution means navigating compliance requirements, field-disruption risk, legacy-system complexity, and stakeholder alignment across functions that don't always agree on definitions, let alone priorities.

Zensar has done this work in life sciences and regulated environments. In short:

- Governed analytics on Databricks with audit-ready logging, lineage, and role-based control in regulated environments.
- Operational data integration that removes manual handoffs and improves OTIF performance through reliable visibility.
- Governed GenAI (RAG) that produces citation-backed, traceable outputs that can be explained and audited.

What Zensar brings to CRM migration specifically:

Pre-built life sciences data models that standardize common domains — engagement, customer, product, territory, content, and outcomes — so teams aren't starting from scratch on definitions that should already be agreed upon.

Migration accelerators that operationalize CRM engagement data into governed intelligence on Databricks without disrupting field execution or compliance workflows during the transition.

ZenseAI.Data — governed retrieval and reasoning over approved data products, with role-based access, audit-ready lineage, and outputs that can be explained to a regulator or a CFO.

What the timeline looks like:

In the first 90 days: stakeholder alignment on a target operating model, the first governed data products stood up, and one to two priority use cases delivered with baseline metrics established — time-to-insight, reconciliation effort, and audit friction.

By month six: ingestion and governance scaled across domains, reusable pipelines and definitions industrialized, additional use cases in flight, and measurable progress against TCO and compliance targets.

Together with Databricks, the result is not more technology. It is better outcomes — measurable ones — built on the infrastructure that the 2026–2029 window makes possible to get right.

Conclusion: A business case for the 2026–2029 window

The 2026–2029 window is more than a technology transition. It is a chance to reset the operating model behind engagement and data, so future change is faster and less risky. Veeva and Salesforce remain foundational to life sciences engagement; Databricks adds a governed intelligence layer to connect structured and unstructured data, enabling analytics and GenAI to be applied responsibly; and Zensar leads the work to operationalize it, delivering measurable improvements in time-to-insight, reduced duplicated effort, TCO, and audit readiness.

To keep ROI real, quantify targets up front (time-to-insight, reconciliation reduction, and reporting/integration run cost) and track them as outcomes:

- Growth: faster, more confident decisions across launch, lifecycle, and global execution
- Efficiency (TCO): fewer one-off pipelines, reports, and manual reconciliations to maintain
- Risk and compliance: stronger control, traceability, and audit readiness as data and AI scale

Zensar's track record in life sciences makes this practical. We've delivered governed platforms in regulated environments that cut processing time from hours to minutes, improved operational reliability, and produced audit-ready, citation-backed outputs for AI use cases. The point is simple: build the governed foundation now, so you're not forced to rebuild it for the next data source, regulator question, or AI initiative.

Decision checklist (answer yes/no)

1. Do you know where your CRM engagement data lives (and who owns it) after it leaves the system of record?
2. Do you have standardized, governed definitions for key metrics so regions and functions get the same answers?
3. Do you have a quantified plan—time-to-insight, reconciliation reduction, and TCO—plus audit-ready lineage and access controls for regulated use?

Authored by

Daniel Gomez,

Head - Partner Sales, Data and Analytics

daniel.gomez@zensar.com

zensar

An RPG Company

At Zensar, we're 'experience-led everything.' We are committed to conceptualizing, designing, engineering, marketing, and managing digital solutions and experiences for over 145+ leading enterprises. Using our 3Es of experience, engineering, and engagement, we harness the power of technology, creativity, and insight to deliver impact.

Part of the $4.8 billion RPG Group, we are headquartered in Pune, India. Our 10,000+ employees work across 30+ locations worldwide, including Milpitas, Seattle, Princeton, Cape Town, London, Zurich, Singapore, and Mexico City.

For more information, please contact: **info@zensar.com | www.zensar.com**