WHITE PAPER

VEEVA CRM → SALESFORCE LIFE SCIENCES CLOUD

From Veeva to AI-Ready: Five Decisions for an Enterprise Move to Life Sciences Cloud

The data, validation and AI governance choices that decide whether your agents work on day one — for large pharma, biotech and MedTech teams.

Dec 2029
Veeva CRM end of support

Mid-2027
Lock your migration path

Data · Consent · Validation · AI governance · Agent budget
November 2026

accellor

EXECUTIVE SUMMARY

Don’t miss the opportunity! Build your foundation for Agentic Customer Engagement while you migrate off Veeva.

Teams that rebuild Veeva's data as-is on Salesforce Life Sciences Cloud (LSC) pay twice: once for the migration, and again to fix the foundation before any agent can use it.

This paper is for large pharma, biotech and MedTech enterprises that plan to move from Veeva CRM to LSC ahead of Veeva's December 2029 end-of-support date. It requires five decisions that enterprises need to make before the build starts.

01 **Treat the data model as a design decision**
Design the model agents will read, don’t just copy the old one.

02 **Make consent machine-readable**
Store consent by channel, purpose, and jurisdiction.

03 **Validate by risk tier**
Size assurance effort to process risk, in line with FDA's 2025 CSA guidance.

04 **Govern AI before the first agent ships**
Keep an AI use-case register before shipping the first agent.

05 **Budget for agents & tokens, not just seats**
Forecast Flex Credit consumption, not just seats.

> THE DAY-ONE TEST
> Can you turn on a pre-call briefing agent for one region without starting a data cleanup project the day Veeva is switched off?

THE CHALLENGE

Replacing a CRM versus building for agents

Most Veeva exit plans follow the same pattern: inventory the _vod objects, map them to LSC, migrate the data, retrain the field, and switch off the old org before December 2029. That is a sound plan for replacing a CRM. It is not a plan for what enterprises need to get ready for an agentic customer engagement platform.

The leading LSC deals were made for AI agents, not screens:

TAKEDA · MAY 2025

Agents across 3 functions

Selected LSC to run AI agents across medical, commercial and patient support.

HALEON · OCT 2025

4,500-person sales force

Chose LSC with Data Cloud and Agentforce to support its field organization.

The bet behind both is that agents will brief reps, wrap up visits, flag account risk and, in time, take on regulated work. **Every one of those agents reads from the data model designed during the migration.**

	Like-for-like migration	AI-ready migration
Data model	× Map each _vod object to its nearest LSC object	✓ HCP/HCO identity designed in Data 360: one person, one profile
Consent	× Consent kept in SOPs and a marketing tool	✓ Consent by channel, purpose and jurisdiction, readable by agents
Validation	× Everything validated to one standard	✓ Risk tiers agreed with QA; CSA-style assurance
AI governance	× Policy written after the first pilot	✓ AI use-case register before the first agent ships
Budget	× Forecast seats	✓ Forecast seats plus Flex Credit consumption

Figure 1. Two ways to leave Veeva: same deadline, very different day one.

DECISION 01

Treat the data model as a design decision

The data model is the single biggest determinant of whether agents work on day one. Veeva CRM runs on one flat family of custom objects. LSC ships purpose-built industry data models, with Data 360 (formerly Data Cloud) underneath to resolve a person or organization into one profile across CRM, ERP and third-party sources.

So the design questions change. "Where does Call2_vod go?" matters less than these:

- **Master source.** Once Veeva Network and OpenData are gone, which source serves as the master for HCP and HCO identity: OneKey, IQVIA, H1, or an in-house MDM?
- **Matching.** How is one physician matched across CRM, medical inquiries, events and patient services, so an agent sees one person instead of four?
- **History.** Which years of activity actually feed a briefing or a model, and which can be archived under retention rules?

Three weeks of identity resolution design costs far less than three months of rebuilding the field's trust after a bad briefing.

DECISION 02

Make consent machine-readable

Agents don't read SOPs; they act on whatever the data allows. Many Veeva customers handle consent through a mix of CRM fields, a marketing tool and local SOPs. If an AI-drafted follow-up goes to a German HCP who withdrew email consent, the failure is in the data, not the agent.

Model consent by channel, purpose and jurisdiction from the first release. GDPR in Europe, the growing list of U.S. state privacy laws, and HIPAA wherever hub or patient-services data touches the org all have to show up as fields an agent checks before it acts.

Apply the same discipline to transfers of value. When meals, speaker fees, and travel are captured cleanly at the point of interaction, CMS Open Payments reporting comes out of the system instead of a quarter-end spreadsheet.

DECISION 03

Validate by risk tier

Validating the whole CRM to one standard is expensive and slows every release after go-live, including every agent added later. Many large companies still do it because that is how it was done in 2012.

FDA's final Computer Software Assurance (CSA) guidance, issued in September 2025, points the other way: assurance effort proportional to process risk, and critical thinking ahead of scripted test volume. It explicitly covers cloud and SaaS solutions and says the framework can apply to AI tools used in production or the quality system.

CSA was written for device production and quality systems, so it does not govern a commercial CRM directly. Its risk-based logic, though, is now how most QA leaders approach regulated SaaS. In practice, agree three tiers with QA before the first sprint.



Figure 2. Validation effort sized to process risk, following FDA CSA logic (final, Sept 2025).

WHY SAMPLE PROCESSES SIT AT THE TOP

21 CFR Part 203 requires a practitioner-signed request and a signed receipt for drug samples, and electronic signatures bring 21 CFR Part 11 controls into scope.

DECISION 04

Govern AI before the first agent ships

Regulators are writing AI rules faster than most companies are writing AI policies.

- **JAN 2025 · FDA (DRAFT)**
  A risk-based credibility assessment for AI used to support regulatory decisions on drugs and biologics.
- **2025 · EUROPEAN COMMISSION**
  Consultation on a revised GMP Annex 11 for computerized systems and a new Annex 22 for AI in medicines manufacturing.
- **DEC 2027 · EU AI ACT**
  High-risk obligations set to apply to stand-alone systems, after the May 2026 omnibus agreement.
- **AUG 2028 · EU AI ACT**
  High-risk obligations set to apply to AI built into regulated products.

Most commercial CRM agents, such as pre-call briefings and visit summaries, typically fall into the lowest AI risk categories. However, effective AI governance requires these use cases to be formally documented rather than assumed to be low risk.

Establish a lightweight AI use-case register that records for each agent:

- Risk classification and business impact
- Data sources accessed and processed
- Required human oversight or review checkpoints
- Underlying foundation model and provider
- Applicable compliance, privacy, and security controls

This governance approach enables compliance, privacy, and risk teams to rapidly approve low-risk productivity agents while focusing deeper scrutiny on agents that handle medical content, pharmacovigilance signals, regulated communications, or patient data.

Life Sciences Cloud (LSC) further strengthens governance by allowing organizations to select from multiple approved foundation models, including Claude, Gemini, OpenAI, and Amazon Bedrock-hosted models. This provides privacy, compliance, and security stakeholders with greater control over model selection, data handling requirements, regulatory alignment, and enterprise AI risk management.

DECISION 05

Budget for agent usage, not just seats

Veeva costs scaled per user; LSC adds a consumption layer that agents draw on. Salesforce's current LSC list prices:

EDITION	PER USER/MONTH	FLEX CREDITS/ORG/YEAR
Core	$350	500,000
Advanced	$525	1,000,000
Customer Engagement add-on (field users)	+$150	—

Teams that forecast seats but not agent volume usually find out around month nine. Put an agent usage forecast in the business case, next to the overlap between the Veeva contract end and the LSC license start.

MEASURING READINESS

The day-one test

The answer to this one question shows whether a migration worked.

On the day you switch off Veeva, can you turn on a pre-call briefing agent for one region without starting a data cleanup project?

IF YES

You moved to Life Sciences Cloud.

IF NOT

You only moved Veeva to a new address.

Accellor scores six readiness dimensions before any agent reaches the field. Each maps back to one of the five decisions, plus the integrations and adoption work that make agents usable in the field.

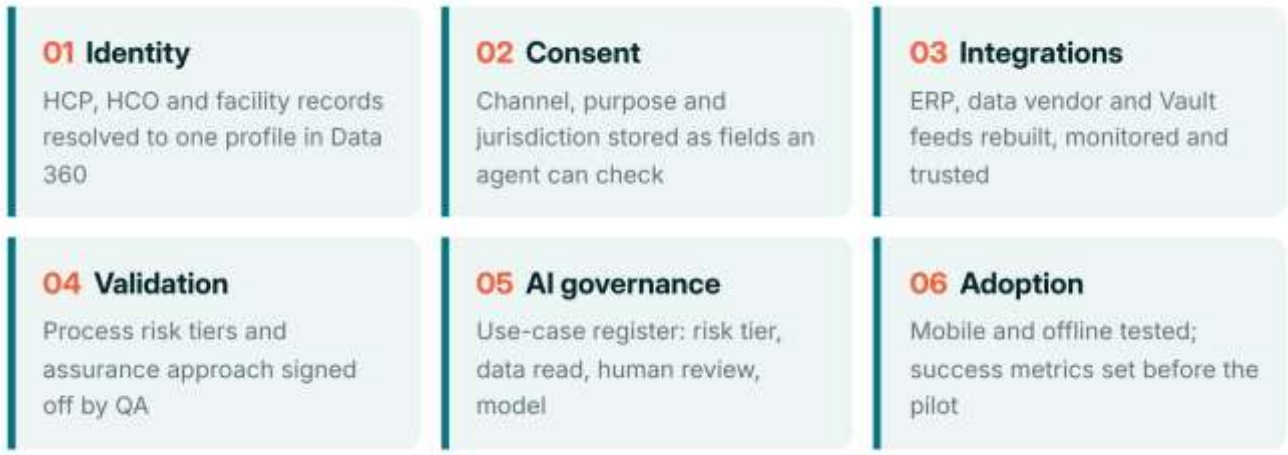


Figure 3. Six readiness dimensions scored before any agent reaches the field.

GETTING STARTED

The Accellor LSC Launchpad

A 4 to 6 week Data and AI readiness assessment that turns the five decisions into a high-confidence plan that your team and stakeholders can align to.

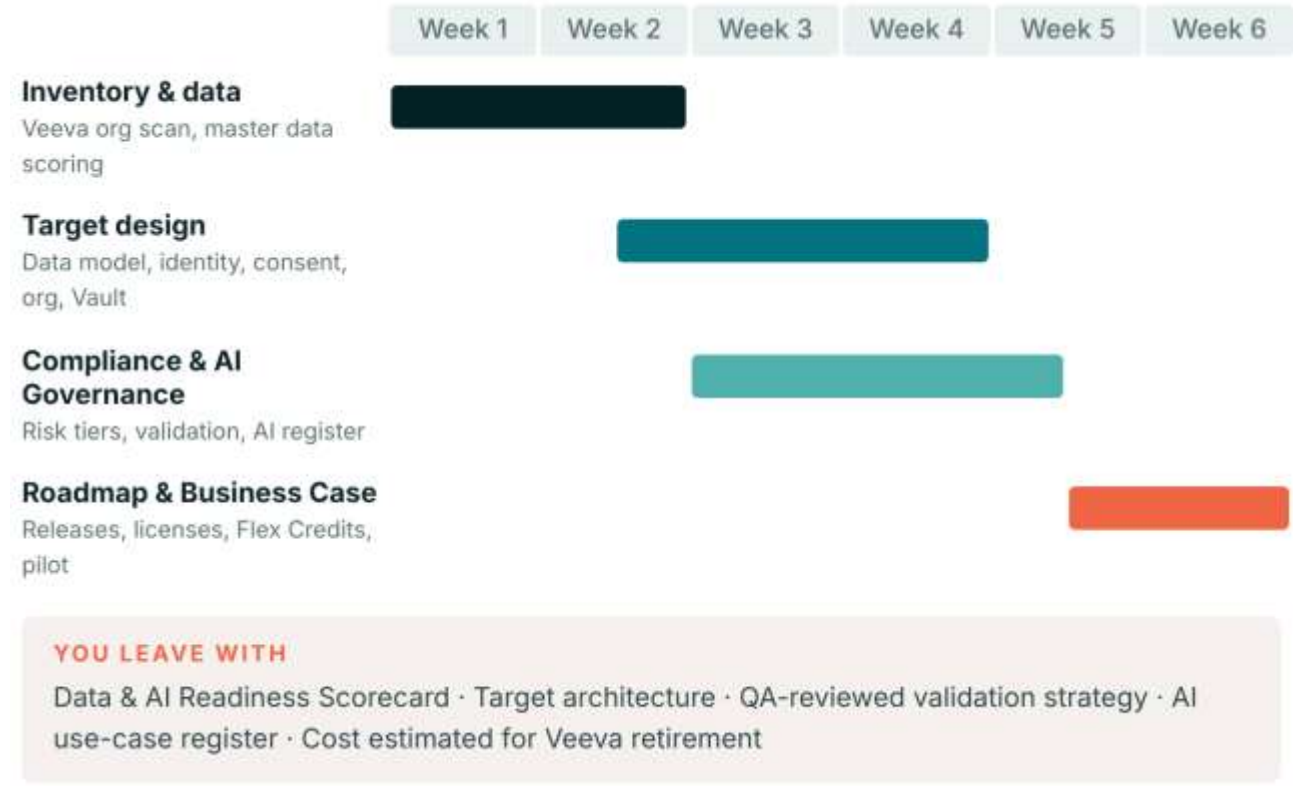


Figure 4. The Accellor LSC Launchpad, enterprise track.

Weeks 1–2 · Inventory and data assessment

Scan of the Veeva org (customized _vod objects, Apex, integrations, volumes), HCP and HCO master data quality scoring, and the Network and OpenData replacement decision.

Weeks 2–4 · Target design

LSC data model, Data 360 identity resolution, consent model, org strategy (transform in place or start clean) and the Vault decision.

Weeks 3–5 · Compliance and AI governance

Process risk tiers agreed with QA, a validation approach built on CSA-style risk thinking, and an AI use-case register with the first three agents scored.

Weeks 5–6 · Roadmap and business case

Release plan to Veeva retirement, license and Flex Credit forecast, and a pilot plan for the first agent.

> **RECOMMENDATION**
>
> Lock your migration path by mid-2027. The Launchpad gives the steering committee the inputs it needs to make that call.

CONCLUSION

The migration and the Agentic AI foundation are the same program

Splitting them doubles the cost. Enterprises that address identity, consent, validation, AI governance, and agent budgeting before the build, will switch on agents the day Veeva goes dark. Those that don't will fund a second project to get there.

ABOUT ACCELLOR

Accellor is a Summit-level Salesforce partner and has worked with Salesforce since 2011.

300+	**4.9 / 5**	**4.88**
Salesforce projects delivered	Client satisfaction score	AppExchange rating across 41 reviews

- Healthcare and Life Sciences industry competency on AppExchange
- Certified Agentforce implementation partner
- ISO 9001 and ISO 27001 certified

Accellor's life sciences work covers field inventory and trunk stock automation for MedTech, AI for regulatory and compliance workflows, and commercial intelligence and forecasting. Delivery runs on a US-led team with a large India-based bench across Salesforce, Data 360, MuleSoft and Boomi. The people who design the data model also rebuild the integrations and run the data migration, so nothing gets lost between teams.

NEXT STEP

Scope your LSC Launchpad

If you are planning or running a Veeva exit, contact Accellor to scope a 4 to 6 week Launchpad.
accellor.com

REFERENCES

Sources

1. Veeva: end-of-support date for Veeva CRM moved to December 2029 (Mar 3, 2026)
2. Business Wire: Takeda selects Salesforce Life Sciences Cloud for Customer Engagement (May 2025)
3. Nasdaq: Haleon selects Agentforce Life Sciences Cloud for Customer Engagement (Oct 2025)
4. Salesforce: Life Sciences Cloud pricing
5. Salesforce: Life Sciences Cloud October 2026 newsletter
6. FDA: Computer Software Assurance for Production and Quality System Software (final, Sept 2025)
7. FDA: Considerations for the Use of AI to Support Regulatory Decision-Making for Drug and Biological Products (draft, Jan 2025)
8. European Commission: consultation on EudraLex Vol. 4 Chapter 4, Annex 11 and new Annex 22 (2025)
9. Gibson Dunn: EU AI Act omnibus agreement and postponed high-risk deadlines (May 2026)
10. eCFR: 21 CFR Part 11, Electronic Records; Electronic Signatures
11. eCFR: 21 CFR Part 203, Prescription Drug Marketing
12. CMS: Open Payments
13. HHS: HIPAA and health information privacy
14. EUR-Lex: General Data Protection Regulation (EU) 2016/679
15. Accellor: Salesforce services
16. Accellor: Life Sciences
17. Salesforce AppExchange: Accellor consulting listing

Pricing reflects Salesforce list prices as of November 2026. EU AI Act dates reflect the May 2026 provisional omnibus agreement. Regulatory references are for context and do not constitute legal advice.