

How to Choose the Right Databricks Partner

For data modernization and AI transformation.

 Whitepaper




The executive premise

Databricks expands what is possible. The partner helps determine whether that possibility becomes measurable business value.

An experience-led guide to modernizing EDW and ETL estates, building an AI-ready foundation, protecting the Databricks investment, and delivering measurable value across revenue, cost, and risk.

A question we hear often

“We are about to invest millions in Databricks. How do we know the Databricks partner will help us change the business, not simply change the platform?”



Executive perspective: the platform is the foundation. Business value is the test

The enterprise is not simply choosing a migration partner. It is choosing who will help decide what should move, what should change, what should retire, which numbers the business will trust, and how the investment will show up in revenue, cost, or risk. We repeatedly see the platform decision get executive attention, while the harder value decisions are deferred until delivery. Databricks expands what is possible. The partner helps determine whether that possibility becomes measurable business performance.

Databricks is a powerful data analytics and AI platform. It can bring engineering, analytics, governance, and AI into a more coherent foundation. But it cannot decide which workloads deserve investment, which numbers the business will certify, where human judgment must remain, or when to switch off an old platform. Those choices sit with the enterprise and the Databricks implementation partner helping it modernize.

A modernization program should answer three direct questions: What revenue opportunity becomes possible? What cost can stop? What risk becomes easier to manage? It should also show that the improvement holds when the business is under pressure—not only when the project team is presenting results. If those answers remain unclear, the program may have delivered activity, but it has not yet earned the investment.

What we hear in engagements

“The platform is live. Why are we still reconciling the same reports, paying for the old estate, and debating which customer number is correct?”

This standard should govern the rest of the decision: judge modernization by what changes in the business, and judge the partner by whether it makes the difficult choices visible, measurable, and defensible.

The most important modernization decisions are rarely technical.

FROM POTENTIAL TO OUTCOMES THROUGH BETTER DECISIONS.

PEOPLE
DECISIONS
DATA
A STRONGER TOMORROW

Platform possibility

MORE DATA
HIGHER PERFORMANCE
GREATER FLEXIBILITY
NEW CAPABILITIES
BROADER ECOSYSTEMS

A BROADER HORIZON
CREATES NEW OPTIONS.



Business outcomes

Revenue
GROW FASTER

Cost
DO MORE WITH LESS

Risk
BE MORE RESILIENT

Productivity
UNLOCK HIGHER VALUE

Trusted AI
TURN POSSIBILITY INTO REAL IMPACT

BETTER DECISIONS
DELIVER A STRONGER TOMORROW.

POSSIBILITY LIVES IN DECISIONS.

SAME QUESTIONS. A BRIGHTER TOMORROW.

The most important modernization decisions are rarely technical. Platform possibility becomes business value through decisions about scope, meaning, economics, adoption, and accountability.

1. Modernization must earn its investment

The investment question

“What must improve in the business for this program to deserve the capital, disruption, and executive attention it will consume?”

A legacy estate is hard to leave because the business still depends on it every day. It closes the books, produces regulatory reports, prices products, supports claims, and keeps critical operations moving. We repeatedly see that the most important logic isn't documented in one place; it lives across code, reports, and the experience of people who know which numbers to trust. The problem is not simply that the technology is old. The problem is that changing it has become expensive, explaining it has become difficult, and using it as a trusted foundation for AI is increasingly risky.

Across banking, insurance, healthcare, retail, manufacturing, and the public sector, we repeatedly hear variations of the same concerns:

- Renewals and capacity decisions arrive before leaders can see whether the last investment created value.
- Business logic is scattered across platforms, pipelines, reports, and personal knowledge, so the lineage behind a critical number is difficult to defend.

- New requests enter the same queue, encouraging analysts to create private copies that drift apart.
- AI pilots create excitement until different teams discover they do not agree on the meaning of revenue, active policy, member, product, or risk.

Waiting is not neutral. Every quarter of delay extends duplicate platform cost, preserves manual controls, and pushes the shutdown date-and the savings tied to it-further out. It also delays the revenue, productivity, and decision improvements the new foundation was meant to enable. The business case therefore needs more than a Databricks migration estimate. Every material benefit needs a baseline, a target, an owner, a way to observe it, and a date when leadership will decide whether it was achieved.

The cost of waiting

Delay does not preserve the status quo. It extends legacy cost, duplicated controls, and fragmented definitions while competitors, customers, and operating expectations continue to move.

Cloud waste is not only a FinOps problem. It is evidence that ownership, workload design, adoption, and operating discipline did not keep pace with consumption. The buyer's question is not whether a partner can deploy the platform. It is whether the partner can help the enterprise explain what it is spending, why it is spending it, and which business outcome the spend supports.

How much money is at stake

More than \$1 trillion¹

in potential run-rate EBITDA across Fortune 500 companies by 2030.

~27%²

surveyed respondents estimated that cloud spend was wasted through idle resources, overprovisioning, duplicate environments, and inefficient operations

Migration alone rarely causes the waste. Poor planning, the wrong workload selection, duplicate environments, rework, insufficient testing, weak adoption, and keeping legacy platforms alive too long amplify it.

Execution discipline-not tooling alone-determines whether the value is captured.

2. Databricks expands the opportunity. It does not define the outcome.

The value-gap question

“Which decisions remain ours even after we choose a powerful platform?”


Databricks offers more than a modern technical foundation. It can reduce the effort spent moving, reconciling, and explaining data so teams can focus on improving products, operations, and decisions. The opportunity is significant, but the platform cannot determine which definitions are authoritative, what should retire, where AI may

support a decision, or who owns consumption. Those choices remain with the enterprise and the partner.

The Databricks partner matters because these decisions are made before the first workload reaches production-and their consequences last long after the implementation team leaves. A capable implementer can demonstrate features and move workloads. A trusted Databricks partner will challenge unnecessary scope, expose economic assumptions, protect business meaning, identify where growth can be enabled, define where human judgment remains, and leave evidence that executives can test.

What frustrates adoption leaders

A technically successful launch followed by low adoption, uncontrolled consumption, unresolved ownership, and an old platform that never switches off. The technology changed. The operating reality did not.




Business need	Databricks opportunity	Decision the enterprise and partner must resolve
Trusted information	Governed data, metrics, lineage, and discovery	Which definitions are authoritative, who certifies them, and where exceptions are allowed
Faster modernization	Discovery, conversion, orchestration and scalable processing	What should retire, move, be redesigned, remain, or be federated
Governed access	Unified governance for data and AI assets	How identities, domains, policies, controls, and evidence work in practice
AI and intelligent applications	A shared foundation for analytics, models and agents	Which decisions AI may support, how answers are evaluated, and where humans retain authority
Economic control	Usage data, workload identifiers and optimization capabilities	Who owns the bill, the unit metric, the budget, and the value case

From fragmented data to trusted AI outcomes

Modernize. Unify. Govern. Activate.

A trusted data and AI foundation
For a brighter tomorrow



Business Outcomes

Lower cost

Faster insight

Trusted data and AI

Resilient operations

Secure and compliant

Reusable data products

The retained modernization foundation visual connects fragmented estates to trusted data, governed AI, and business outcomes.

3. The decisions that determine value

The scope question

“What deserves to move, what deserves to change, and what should not survive the program?”

Scope is one of the largest economic levers in modernization. A common mistake is assuming that every workload deserves to move simply because it exists. Migrating everything carries yesterday’s duplication and technical debt into tomorrow’s consumption model. Redesigning everything creates a different problem: cost, delay, and transformation risk that the business may not need.

Partner judgment shapes the economics before migration begins. Every workload carried forward adds cost, complexity, and an obligation to support it.

Start with the evidence: which workloads are used, what they cost, who depends on them, which controls they support, and what happens if they fail. Then make the decision explicit. Retire what no longer earns its place. Retain or federate what does not need to move. Replatform where speed matters. Refactor or rewrite only where the business value justifies the added cost and disruption.

Automation can accelerate inventory, complexity analysis, code conversion, test generation, reconciliation, and documentation. It cannot decide what deserves to survive, certify a financial definition, approve a control, or determine how much change an operation can absorb. We frequently see automation preserve logic that nobody intended to keep because the business decision was never made. Those are judgment calls, and they need accountable owners before you move code.

A view for executives

The most expensive workloads in many programs are the ones that should never have been migrated. Ask to see the disposition logic before you see the price.

Disposition	Use it when	The value decision
Retire	The asset is unused, duplicated, or no longer justified	Remove cost before paying to migrate it
Retain	The value of change does not justify the effort or risk	Keep it deliberately and set a review date
Federate	Data should remain in place for now	Avoid movement while preserving governed access
Replatform	The logic is sound and speed matters	Move efficiently without pretending the design has been modernized
Refactor	Known debt blocks reliability, scale, or economics	Invest selectively to improve the operating model
Rewrite	The current logic or design cannot support the future need	Rebuild only where the business value justifies the disruption

4. Protect the numbers the business trusts

The question finance eventually asks

“Can I still trust the number?”

The first month-end close after cutover is often where confidence is won or lost. The platform may be running, the tables may be loaded, and the row counts may match. What we hear frequently from finance leaders is simple: if revenue, reserves, or risk changes and nobody can explain why, the technical success no longer matters. Finance remembers that the number moved.

Technical reconciliation asks whether data arrived correctly. Business acceptance asks whether the answer still supports the decision, obligation, or control that depends on it. A row count can match while revenue, reserves, eligibility, or risk changes. That is why you must prove technical parity and business trust separately.

Validation discipline protects business trust. Technical equivalence is not enough when a financial, regulatory, or operating meaning changes.

Sometimes the right choice is to reproduce the legacy answer exactly because a contract, regulator, or financial process relies on it. Sometimes the legacy answer is wrong, and modernization creates the moment to correct it. Either choice needs a named business approver and a record of the decision. This isn't merely an engineering fix.

Don't wait until cutover to discover that engineering, finance, operations, and risk were proving different things. Agree on the evidence and the decision owners before migration begins. Once a disputed number reaches production, the argument slows, costs more, and becomes harder to contain.

Level of proof	Question	Evidence	Accountable owner
1. Structure	Did everything arrive?	Counts, schemas, types, and null checks	Data engineering
2. Content	Do the values match?	Checksums, record differences, and tolerance rules	Data engineering
3. Meaning	Does each KPI still mean the same thing?	Signed definitions, joins, date logic, and lineage	Business data owner and finance
4. Business outcome	Do users receive a trusted answer when they need it?	Operational-cycle comparison and user acceptance	Business sponsor
5. Control	Did the controls survive?	Role mapping, masking, lineage, audit, and retention tests	Security, compliance, and risk

What we hear

“The migration team says the data reconciled. Finance says the result changed. Who decides whether the new answer is right?”

5. Make the economics hold

The CFO question

“When do the savings reach the P&L, and when does the investment begin enabling growth?”

Comparing platform bills is not a business case. The full economics include implementation, redesign, parallel operations, cloud and platform consumption, support, adoption, legacy shutdown, and benefit timing. They also include the growth the new foundation is expected to enable: faster product launches, better pricing, stronger retention, more effective cross-sell, or improved service capacity. Separate cash savings, avoided costs, released capacity, revenue effects, and risk reduction so you don't count the same benefit twice. The model should show when each benefit hits the P&L, what must happen before it appears, and who is accountable when it doesn't.

Realized value requires ownership. Consumption must be tied to a workload, an accountable owner, and an outcome the business can defend.


Consumption is not value. A rising bill may signal healthy adoption or uncontrolled growth. A lower bill may signal optimization or a platform that nobody is using. We repeatedly see organizations explain the invoice long before they can explain the outcome. Executives need to see which workload, product, decision, or business transaction consumed the spend-and whether the result justified it.

Ask two direct questions of every partner: who owns the legacy shutdown date, and who owns the cost model after go-live? One lesson we have learned is that the old platform rarely survives because of technology. It survives because nobody is willing to own the decision to turn it off. If ownership is vague, the savings are probably presentation assumptions rather than operating commitments.

The economic truth

Go-live creates capability. Legacy shutdown releases savings. Adoption creates operating value. Measurement proves whether any of it happened.

Economic loop	Executive question	Evidence
Retirement	What costs, duplication, or dependencies can we stop?	Contracts removed, infrastructure retired, duplicate runs ended, named shutdown date
Production	Can trusted products move from idea to daily use faster?	Lead time, reliability, recovery, reuse, adoption, and quality
Business outcome	Did a decision, process, product, or control improve?	Revenue enabled, cycle time, productivity, loss avoided, service, or compliance outcome



6. From trusted data to business-ready AI

The AI question beneath every AI question

“What makes us believe the answer?”

AI is the accelerator, not the destination. The business opportunity begins with a decision or workflow that needs to improve: helping an underwriter price risk with greater confidence, giving a claims handler better context, identifying a manufacturing exception earlier, reducing service effort, strengthening clinical or research insight, or helping a frontline employee act with reliable information. Databricks can help operationalize that opportunity, but value appears only when trusted data, business context, human authority, and measurable outcomes come together.

AI programs lose credibility when the enterprise cannot explain the data, definition, authority, or acceptable use behind an answer. We repeatedly see the model become the focus while the business decision remains undefined. A stronger model cannot resolve two departments that define

“active customer” differently, and automation cannot replace the owner who must decide which definition governs the action.


AI credibility depends on context, evaluation, and authority. The partner should make each explicit before a model influences a business decision.


The path to business-ready AI is practical: govern access, agree on meaning, measure quality and lineage, benchmark the answers, define where people retain authority, observe how the system behaves, and connect its use to a measurable return. One lesson we have learned is that trust is easier to build before deployment than to recover after a confident answer proves wrong in front of a customer, regulator, clinician, or executive team.

Industry experience matters when it helps translate business processes, evidence, and controls into definitions, evaluation sets, and operating workflows. It should not substitute for customer validation. The customer still owns its meaning; the partner should make that ownership actionable.

What industry context changes

A table name does not explain how a policy moves through its lifecycle, how a patient identity should be resolved, how a product hierarchy changes margin, or which manufacturing exception requires human judgment.





Same question, two starting points

What changes is the context the agent starts from

Starting from raw tables



Question

"How many active policies do we have?"



The agent sees table and column names

e.g. policy, policy_status, eff_date ...



It guesses which table, status code and date to use

Based on patterns, not business meaning.



Three people get three answers—and nobody can trace why.

Starting from governed context



Question

"How many active policies do we have?"



Business definitions

e.g. what "active policy" means



Entities and relationships

e.g. policy, customer, product, with defined joins



Industry process knowledge

e.g. policy lifecycle, status transitions, effective dates



Rules and access

e.g. calculation rules, data security, user permissions



Evaluation set

e.g. test questions and expected answers



One answer, traced to a certified definition. Returned under the asker's permissions. Re-tested on every change.

Ownership of semantics

Databricks provides

Rules and access, governance, security, platform

Zensar delivers

Industry process knowledge (domain expertise, accelerators)

Joint with customer owners

Business definitions, entities and relationships, Evaluation set (co-created and owned by the business)

7. Make value last beyond go-live

The accountability question

“Who protects the value when the implementation team leaves?”

78%³

of 315 surveyed US cloud professionals estimated that 21%-50% of their cloud spend was waste that could be optimized.

Adoption becomes durable when handover creates operating independence. Internal teams must be able to run, improve, and defend the platform without routine partner intervention.


Value lasts only when ownership lasts. Successful organizations assign named owners to data products, business metrics, controls, consumption, and benefits after go-live. They apply the same discipline to release quality, cost reviews, platform


optimization, retirement, and knowledge transfer. The organizations that struggle allow those responsibilities to disappear with the project team-and eventually recreate the cost and complexity they paid to remove.

- Make parity and control tests part of every material release, not a one-time migration exercise.
- Treat the first months after cutover as a tuning period with explicit reliability, adoption, and economic reviews.
- Review cost and performance by workload and business unit, not only at the total account level.
- Re-run retirement and duplication tests so the new estate does not recreate the sprawl it replaced.
- Measure handover by what internal teams can operate and improve independently, not by training attendance.

The long-term test

The right Databricks partner should make the enterprise less dependent on the partner over time, while increasing the organization's ability to govern, improve and scale the platform.





8. Choose the Databricks partner on evidence

The decision question

“What would we need to see before trusting this partner with the economics, business truth and operating future of the platform?”

Certifications, partner tier, and technical capacity establish eligibility. They do not prove that a provider can protect the economics, business meaning, and operating future of your estate. In executive conversations, we hear that the difficult questions begin after reviewing credentials: What will the partner challenge? What will it refuse to migrate? How will it protect the numbers the business trusts? Ask every provider—including Zensar—to show how it makes scope decisions, validates critical numbers, models all-in economics, governs AI, assigns accountable delivery leaders, and leaves the customer able to operate independently.

The evidence should be visible before you sign the contract. Ask each provider, including Zensar, to show:

- A representative estate inventory and the logic used to retire, retain, federate, replatform, refactor, or rewrite workloads.

- A comparable client example with scope, named outcomes, and a reference the buyer can speak with where permissions allow.
- A parity and business-acceptance approach that covers structure, content, meaning, outcome, and control.
- An all-in economic model with customer inputs, transparent assumptions, sensitivity, and a named shutdown path.
- The named delivery leadership and the people who will actually perform the work.
- How AI use cases will be grounded, evaluated, governed, adopted, and measured.
- The operating model, knowledge transfer, and evidence that the customer can run and improve the platform after handover.

Begin with one consequential workload family and a measurable business hypothesis. Use the first investment to inspect real dependencies, test the partner’s judgment, prove business meaning, and determine what can retire. The result must be allowed to say “proceed,” “redesign,” “not yet,” or “stop.” If the exercise can only confirm the original plan, it is a sales demonstration, not evidence.

The final test

Ask a prospective Databricks partner what it would refuse to migrate, what it would require the business to decide, and what evidence would make it stop or change direction. A polished demonstration shows what the platform can do. Judgment shows whether the partner can be trusted when the economics, business meaning, and operating trade-offs become real.

A note on Zensar

Zensar sponsored this paper and may be one of the providers under consideration. The argument should remain useful even if we choose another partner. Hold us to the same standard: ask what we would refuse to migrate, which assumptions we would challenge, how we would protect business meaning, when savings would hit the P&L, and what evidence would make us change direction.

Our role is to connect decisions that are too often managed separately: what should retire, which business definitions must survive, where Databricks can improve the operating model, how AI will be evaluated, when legacy cost can stop, and how the customer will run the platform independently. ZenseAI.Data can accelerate estate discovery and

modernization analysis, while our industry and delivery experience provides a practical starting point. The customer remains the owner of business meaning, operating decisions, and outcomes.


We work across more than one data platform, so retaining or federating a workload remains a legitimate answer when it better serves the business. A partner should be willing to recommend “do not move this” when the economics, risk, or operating need do not support migration. Ask us to show where we made that call, how we validated it, and what the customer ultimately owned.

Practical starting point

Start with one visible business problem: a close that takes too long, a claims process that requires too much manual judgment, a supply chain decision made with stale information, a growth opportunity constrained by fragmented data, or a platform cost nobody can explain. Ask the

partner to trace that problem through the data, workloads, controls, operating decisions, and economics. Then require a credible answer to four questions: What changes? What retires? Who owns the outcome? How will we know it worked?

A migration is an activity. A business outcome is evidence.





Sources and disclosures

Independent research is cited in the footnotes. Databricks capability descriptions are based on published Databricks materials; vendor-reported and Zensar-reported information should be read as originating from the named organization. Any illustrative economics are examples not forecasts and depend on the stated assumptions. Unattributed executive questions and engagement

observations are representative composites drawn from experience. They are not verbatim quotations from individual customers or executives. Zensar sponsored this paper and may be one of the providers considered; the selection principles are intended to remain useful when evaluating any partner.

¹James Kaplan, Will Forrest, Soyoko Umeno, Trevor Schroeck and Jonathan Reilly, “Cloud’s trillion-dollar prize is up for grabs,” McKinsey Quarterly, February 26, 2021, “About the research” and discussion of the value at stake. McKinsey’s review estimates more than \$1 trillion in run-rate EBITDA across Fortune 500 companies as up for grabs in 2030 and describes this as an estimate of what should be possible, not a prediction.

²<https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/clouds-trillion-dollar-prize-is-up-for-grabs>
Flexera, “2025 State of the Cloud Report,” March 19, 2025. Survey of 759 cloud decision-makers and users worldwide; respondents estimated that 27% of cloud spend was wasted.
<https://www.flexera.com/blog/finops/the-latest-cloud-computing-trends-flexera-2025-state-of-the-cloud-report/>

³Stacklet, “Cloud Cost & Usage Optimization Survey,” Fall 2024, pages 2 and 6. Stacklet and Centiment surveyed 315 U.S. cloud professionals responsible for cloud infrastructure management and FinOps; 78% estimated that 21–50% of their cloud spend was waste that could be optimized. <https://stacklet.ai/survey-usage-optimization-2024/>



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