

How to Choose the Right Snowflake Partner for Data Modernization and AI Transformation

A practical guide to building AI-ready enterprise data foundations, governed data products, and accelerated EDW and ETL modernization on the Snowflake AI Data Cloud.r

 Whitepaper



The premise. I have seen Snowflake programs work when the partner treats legacy ETL and EDW modernization as a business repair job, not a platform move. Knowing that Snowflake is table stakes. The harder test is whether the partner can take old code, old warehouses, fragile batch logic, and years of operating habits and turn them into something the business can understand, trust, and run.

A business-first guide for CIOs, CDOs, CTOs, enterprise architects, data leaders, and business executives evaluating Snowflake modernization partners.

How to read this paper. The primary decision is partner selection: choosing a Snowflake services partner for data

modernization and AI transformation. The secondary goal is the target state: an AI-ready Snowflake foundation and a portfolio of governed data products. EDW and ETL modernization are the proof path, demonstrating whether the partner can reduce risk, accelerate delivery, and leave the business with something better than a converted legacy estate.

Executive answer. Choose the partner who can make hard decisions before the migration factory starts. The executive question is not only "Can we migrate?" but also "What becomes possible after modernization that is too slow, expensive, risky, or impractical today?" Without that lens, the work may still produce a successful cutover, but it will not necessarily produce a better business.

Introduction

Everyone wants the airplane. I worry about air traffic control.

This paper could start with generative AI, agents, copilots, and automation. That would make it sound more current. It would also skip the point.

AI is the visible aircraft. Executives can see it, fund it, and ask for demonstrations. The less visible work is air traffic control: trusted, governed, traceable, secure, costed, reusable data operated at cloud scale.

The foundation has to come first. When evidence is buried in brittle warehouses, outdated ETL jobs, and reporting estates nobody fully understands, every AI use case starts by reopening the same investigation.

The defensible point is not that every byte of enterprise data must instantly leave every on-premises system. Regulated industries, residency rules, latency requirements, and transition realities can require hybrid patterns. The defensible point is that Snowflake's AI Data Cloud is a cloud-native platform; Snowflake documentation states that it uses public cloud infrastructure and cannot be installed locally or run on private cloud infrastructure. For Snowflake-centered modernization and AI transformation, the control point must

therefore shift to a cloud-native Snowflake foundation, even if some source systems remain on-premises for a period.

The practical argument is simple: do not start with AI theater. Start with the operating base that lets AI survive contact with the business.

The analogy. Applications are the airplanes. Business processes are the passengers. AI is the new navigation layer. Data is air traffic control. For Snowflake-led modernization, that control tower belongs in a cloud-native foundation, not trapped in legacy on-premises silos.

I see many organizations investing aggressively in AI while still running environments where nobody can explain what is underneath. That is not an AI problem. It is a visibility, accountability, and modernization problem.

For teams moving from Teradata, Hadoop, Oracle, Netezza, SQL Server, Informatica, SSIS, Ab Initio, BODS, SAS, or similar estates, partner choice is not a procurement detail. A migration team can land workloads in Snowflake and still preserve yesterday's design mistakes. The better partner knows when to move, when to redesign, and when to retire the old system entirely.

The enterprise is not a stack. It is an airport.

Architecture diagrams are useful, but they also lie a little. They make the enterprise look cleaner than it really is.

In reality, most companies run across ERP, CRM, HR, finance, supply chain, manufacturing, customer engagement, security, cloud services, analytics platforms, custom applications, legacy reporting estates, and systems inherited through acquisition.

Every system produces data. Every system consumes data. Every system depends on something else. Information is constantly

arriving, departing, waiting, changing, and being reinterpreted by people who may not agree on its meaning.

Snowflake can be a critical control point, but it is not the whole airport. A partner who sees only the data platform sees one runway. A partner who understands the business also sees the passengers, flight paths, maintenance constraints, weather, gates, and operating rules.

The wrong partner makes hidden cost permanent.

Leaders can see invoices for infrastructure, cloud services, software licenses, support contracts, and Snowflake consumption. Those costs matter. Many of the costs slowing the enterprise down are harder to see.

- Unused reports nobody reads
- Pipelines nobody owns
- Transformations nobody can explain
- ETL jobs nobody remembers creating
- Data copied repeatedly because teams do not trust the source
- Obsolete platforms retained after acquisitions
- Business rules buried in code written by people who have left
- Environments still running because nobody knows what will break if they are turned off
- Incomplete projects, duplicate assets, shadow workflows, idle storage, and unnecessary compute

This is the part of technical debt that rarely has a clean invoice. It consumes storage, compute, licenses, engineering capacity, operational attention, and time. Worse, it slows every future change because teams must first determine what exists, who uses it, what depends on it, and what can be safely removed.

A weak implementation partner moves the debt and calls it progress. A serious one puts names, costs, usage, dependencies, and retirement decisions on the table before the move begins.

The first modernization benefit. Sometimes the first value is not building something new. It is finally learning what is running, what it costs, who uses it, and whether it deserves to survive.

Why companies really modernize

Companies rarely modernize because an older platform is unfashionable. They modernize because the current estate has started to get in the way: decisions take too long, costs are hard to justify, risk is accumulating, and AI plans keep running into data that nobody quite trusts.

Cannot scale	Demand grows faster than the platform can support.
Performance issues	Business users wait longer for data, reports, and decisions.
Excessive technical debt	Unknown dependencies and obsolete assets raise cost and risk.
Skills scarcity	Critical platforms depend on difficult-to-find specialists.
Platform risk	Aging or unsupported technology increases operational exposure.
Slow analytics delivery	New questions take weeks instead of days or hours.

The visible symptom may appear technical. The real issue is that information no longer moves through the business at the speed, cost, and confidence leaders need.

Many organizations attempt to solve this through lift-and-shift migrations, tool conversions, dashboard replacements, cloud relocations, or lengthy assessments. Those activities can be necessary, but they rarely change how the business uses data. A new platform running old logic, ownership models, and processes often delivers the same outcomes at a different address.

Migration, modernization, transformation

Migration, conversion, modernization, and transformation are not the same job.

Many failed programs start with loose language. People use migration, conversion, modernization, optimization, and transformation as if they are interchangeable. They are not. Each one creates different obligations for the partner and different expectations for the buyer.

Motion	What it means	Result
Migration	Moves data, code, and workloads to a new environment.	Location changes.
Conversion	Translates legacy objects and logic into target-platform equivalents.	The old design may survive.
Modernization	Changes architecture, engineering, governance, ownership, semantics, and reuse.	Capability changes.
Optimization	Continuously improves performance, reliability, and economics.	Value is protected.
Transformation	Uses the new capability to change a business process, decision, product, or outcome.	Business performance changes.

Don't judge a modernization program by how much technology moved. Judge it by what became possible after the move and what legacy capability could finally be retired. Migration is the activity. Modernization is the capability. Transformation is the business result .

A program can convert thousands of ETL jobs, tables, views, stored procedures, and reports and still preserve the old bottlenecks, duplicate logic, ownership confusion, manual controls, and poor

adoption. That may be a completed conversion, but it is not a successful modernization.

The valuable work is uncomfortable. Some ETL logic should be converted, some redesigned, and some should be retired. Some EDW workloads deserve proper Snowflake patterns; others only exist because nobody has dared to turn them off.

Changing the engine mid-flight is exactly the point.

Legacy ETL and EDW modernization usually happens while the business is still running. Revenue transactions, financial close, claims processing, supply chains, manufacturing plans, patient workflows, customer reporting, trading, settlement, and regulatory submissions cannot wait for the new platform to be perfect.

The cost risk has also changed. Gartner forecasts AI-optimized infrastructure-as-a-service spending to reach about \$42 billion in 2026, up 96.4%, as enterprises move AI from experiments into production. The warning is straightforward: weak modernization and weak FinOps can consume the business case before AI ever reaches scale. [1]

Additional cost-risk benchmarks are included in the source notes from Harness, Flexera, and McKinsey for readers seeking a broader view of cloud waste and FinOps risk.

Migration velocity is easy to celebrate and dangerous to overvalue. The partner still has to rationalize the estate, design for consumption, reduce unnecessary movement of sensitive data, and build FinOps into the operating model from the start.

Programs lose value through a familiar set of decisions and omissions:

- Planning before the estate is understood
- Selecting workloads because they are easy rather than valuable



- Operating duplicate environments longer than planned
- Rework caused by undiscovered dependencies or buried business logic
- Technical testing without business reconciliation
- Poor adoption because users do not trust the new data
- Keeping legacy platforms alive after go-live

Technical testing proves the platform works. Business validation proves the organization trusts the result. Successful modernization programs plan for both.

The risks that must be priced before movement

Risk	What happens	Business consequence	Measure
Data loss	Missing records or failed reconciliation	Rework, restatement, and delayed cutover	Rows and balances outside tolerance
Business disruption	Reports, interfaces, or processes fail	Revenue impact and missed service levels	Downtime, failed interfaces, missed SLAs
Data quality	Duplicate or inconsistent rules move forward	Low adoption and unreliable AI	Reconciliation and quality pass rates
Cost overruns	Hidden dependencies expand scope	Higher TCO and delayed value	Budget, schedule, and rework variance
Security and compliance	Access, masking, retention, or lineage gaps	Audit findings and regulatory exposure	Open exceptions and remediation time

Bad partners treat conversion risk as a staffing issue. Good partners attack it with automation, reconciliation, delivery evidence, and enough judgment to stop the wrong work from moving forward.

The AI-ready foundation

AI does not fix trust. It exposes whether trust exists.

The AI problem I see most often is not a shortage of ideas. It is the inability to repeat the work without restarting the data hunt every time.

The pattern is familiar. A promising use case emerges. Then the questions start: Where is the data? Can it be used? Who owns it? What does the field mean? Which rule is correct? Is it current? Can the answer be traced? Are we allowed to expose it to this user or model?

If every project pays that tax again, the company lacks AI capability. It has a series of pilots with good intentions and fragile foundations. Snowflake positions the AI Data Cloud as a unified platform for data engineering, analytics, AI, applications, collaboration, governance, security, and observability. Its cloud-native architecture matters because AI-ready foundations need elastic capacity, controlled access, traceability, and the ability to work across multiple workload patterns. [2][3]


Snowflake also changes the modernization opportunity by letting the enterprise simplify the estate rather than recreate it. Its separation of storage and compute, integrated governance, support for structured and unstructured data, secure collaboration, application capabilities, and AI services allow organizations to bring more workloads to governed data, rather than repeatedly moving data between

disconnected platforms. The partner's responsibility is to use those capabilities deliberately, not reproduce the legacy architecture inside a newer environment.

That distinction matters. The issue is not whether a company can keep some systems on premises during transition. Many will. The issue is whether the enterprise data foundation for AI remains constrained by on-premises EDW and ETL patterns that were never designed for elastic scale, integrated governance, cross-cloud collaboration, AI workloads, and reusable data products. This is where the services partner becomes strategic. The work is not selling an AI story. It is helping the enterprise build the foundation first, then proving the capability by practically modernizing legacy EDW and ETL.

Data modernization comes first because AI transformation doesn't scale on confusion. Snowflake matters. Still, the platform alone will not assign ownership, resolve definitions, retire waste, or make people trust the answer.

The measurable outcome is not Snowflake adoption. The destination is a portfolio of governed data products that improve decisions, simplify operations, accelerate analytics, and create a foundation the business can repeatedly use for AI.



The executive question is rarely whether a workload can be migrated. The better question is what modernization makes possible that is difficult, expensive, slow, or impossible today. Faster decisions, real-time insight, trusted AI, enterprise-wide access to information, and reusable data products matter more than the technology move itself.

That is the executive test for modernization. If the program only changes where data sits, the ambition is too small. The work should make something new practical: faster underwriting decisions, cleaner margin analysis, better inventory actions, safer AI assistance, or simply fewer teams arguing over whose number is correct.

What an AI-ready enterprise foundation requires

 Trusted and explainable The business can reconcile results, trace answers to sources, and understand how rules were applied.	 Governed and secure Access, policy, retention, lineage, accountability, and sensitive-data controls are explicit.
 Reusable and discoverable People, systems, analytics, and AI can find and reuse the same governed assets instead of rebuilding them.	 Scalable and observable The foundation supports growth in volume, concurrency, cost visibility, and performance expectations.
 Semantically meaningful Business terms, metrics, and definitions are consistent enough for decisions and AI use.	 Owned Each important asset has accountable business and technical owners, not anonymous platform ownership.

AI does not fix trust. It exposes whether trust exists.

The destination is not Snowflake adoption. It is a portfolio of governed data products that support decisions, analytics, applications, and AI.

A data product is a governed business asset with a defined purpose, named ownership, known consumers, agreed-upon meaning, measurable quality, access policies, lineage, service expectations, and evidence of use. It may include curated domain data, packaged

metrics, semantic models, feature sources, retrieval-ready content, or governed APIs and shares.

A customer data product, for example, is not simply a consolidated customer table. It has an agreed-upon business definition, named owners, traceable sources, quality thresholds, access policies, service expectations, known consumers, and measures that show whether it improves decisions or operations.

The shift is practical:

Current state	Modernized state
Monolithic warehouse	Modular domain products
Duplicated metrics	Governed definitions
Bespoke pipelines	Reusable patterns
Technical ownership	Business and technical co-ownership
Data delivery	Measured consumption and outcome
Repeated AI preparation	Reusable AI context

Choosing the right partner

Automation tells you more than certification counts.

In EDW and ETL modernization, a polished certification slide does not tell me enough. I want to see how the partner discovers the estate, decides what not to move, uses automation without hiding behind it, applies industry judgment, and leaves the customer with data products and operating discipline that endure after the project team leaves.

People and certifications matter, but they vary by assignment. Packaged IP matters because it captures lessons the firm should not

have to relearn at the customer's expense. The asset still has to prove itself in delivery: does it reduce risk, shorten time-to-value, help establish governance, and leave behind a foundation the business can operate on?

The better selection question. Show me the accelerators. Show me what they produce. Show me where they stop and where human judgment takes over. If the answer is only "we have smart people," the risk has not gone away.

Why acceleration matters in Snowflake consulting services

Legacy EDW and ETL modernization is where partner capability becomes visible. The work is repetitive enough to benefit from automation and risky enough to punish shallow judgment. A strong SI can scan the estate, classify complexity, convert what is safe, flag what should not move, generate target patterns, automate validation, and show evidence before cutover.

EDW and ETL modernization shows whether the partner can accelerate delivery, reduce error, preserve business continuity, and

still make the judgment calls automation can't.

The SI differentiator. The real asset is not a logo, a methodology slide, or a bench of résumés. It is delivery knowledge turned into tools, checks, patterns, and evidence the next program can reuse.

Ten tests for an EDW and ETL modernization partner

1. **Business outcome clarity:** Can the partner connect the program to a measurable decision, process, KPI, or business outcome?
2. **Estate intelligence:** Can the partner discover ETL mappings, scheduler dependencies, SQL logic, lineage, usage, technical debt, and business rules before defining scope?
3. **Automation and packaged IP:** Can the partner demonstrate working accelerators for code extraction, conversion, profiling, pipeline generation, validation, governance automation, and documentation?
4. **Modernization judgment:** Can the partner distinguish what to retire, consolidate, convert, replatform, refactor, redesign, automate, or retain temporarily?
5. **Industry context:** Does the partner understand the data, controls, and business processes specific to the industry?
6. **Business continuity:** Can the partner preserve operations through wave planning, parallel runs, automated reconciliation, rollback, controlled cutover, and decommissioning discipline?
7. **Data product capability:** Can the partner establish ownership, semantics, quality, policy, lineage, and lifecycle management?
8. **AI production capability:** Can the partner move from a demonstration to governed, monitored, integrated, and economically sustainable production use?
9. **TCO and FinOps discipline:** Can the partner right-size workloads, optimize consumption, attribute cost, and remove legacy spend?
10. **Operations and knowledge transfer:** Can the customer operate and continuously improve the platform after implementation?

How to test the partner's claims

These tests matter only if the partner can provide evidence. Ask for artifacts, not descriptions: estate-discovery outputs, accelerator results, conversion samples, reconciliation evidence, data-product operating models, AI production patterns, FinOps controls, and knowledge-transfer plans. A credible partner will specify what is automated, what still needs expert judgment, and which business decisions must be made before the work moves forward.

- Show an example of estate discovery and dependency mapping.
- Show what the accelerators actually produce.
- Show how conversion is validated against legacy results.
- Show how sensitive data stays within governed boundaries where possible.
- Show how cost, egress, and consumption are controlled during and after migration.
- Show how data products are owned, measured, improved, and retired.
- Show how internal teams are enabled after go-live.

Credentials are evidence, not the decision.

Partner tier, competencies, certifications, accelerators, and scale are important qualification signals. They do not replace evidence from the named delivery team. Buyers should request comparable programs, relevant source-platform experience, industry knowledge, sample accelerator outputs, a test-and-cutover approach, an optimization plan, and clear post-go-live accountability. Select the partner whose people, methods, assets, and accountability align with the business problem. Everything else is supporting evidence.

already been made. It should show whether the partner can discover a real slice of the estate, handle representative logic, reconcile trusted outputs, expose operating risks, estimate consumption cost, and connect the work to a business measure that matters.

A pilot should not be a polished demo with friendly data. It should answer whether value appears soon enough, whether the pattern survives real quality, access, lineage, cost, and adoption constraints, and whether the same approach can scale without becoming a custom one-off.

Before the work starts, define the owner, users, budget guardrails, success measures, quality gates, and scale path. If the team cannot identify who owns the outcome, where value will be measured, how quality will be checked, how cost will be tracked, and what would make the pilot production-ready, the exercise is not proof. It is activity.

Run a proof of value, not a canned demo.

A useful proof of value should generate evidence early enough to shape the program, not after major scope and budget decisions have

The Zensar approach

Zensar's advantage is not a product list.

The hard work starts before platform design. Someone has to find what exists, decide what matters, identify what to retire, and turn the remaining estate into data the business can trust.

That is the outcome executives care about. ZenseAI.Data, ZenseAI.AgentMesh, and related capabilities are proof points only if they help deliver that outcome. They should never become the headline.

A cleaner frame is this: strong consulting method, experienced people, industry context, automation where it actually helps, and platforms that prove the approach can scale.

Zensar's differentiation is connecting decisions that are normally separated. We connect estate discovery to retirement decisions,

conversion automation to business reconciliation, platform engineering to data-product ownership, AI ambition to governance and economics, and implementation to the enterprise systems and operating processes around Snowflake. That connection matters because most modernization failures occur between workstreams, not inside a single technology.

What makes this useful for a Snowflake buyer is not that Zensar has data engineers, accelerators, platforms, or industry knowledge. Many firms can claim those. The difference is how those pieces connect in delivery: modernization choices tie to retirement decisions, data products tie to ownership and adoption, AI use cases tie to governance and cost, and implementation ties to the surrounding enterprise systems that determine whether the work actually sticks.

A business-first Snowflake modernization service with packaged IP

Frame the outcome	Identify the business problem, decision, process, risk, or opportunity that legacy EDW and ETL modernization must improve.
Understand the process	Map inputs, outputs, handoffs, exceptions, reports, ETL jobs, batch windows, approvals, manual work, systems, and controls.
Assess the operating model	Evaluate roles, ownership, skills, governance, data responsibilities, testing responsibilities, release constraints, and operational handoffs.
Examine the ecosystem	Understand source platforms, ETL tools, orchestration, files, databases, warehouses, reporting layers, security, cloud services, analytics, AI, and operational tooling.
Define the target state	Specify business, process, data, technology, risk, and AI measures that describe success.
Choose the path	Use accelerators to decide what to eliminate, consolidate, convert, refactor, modernize, automate, govern, productize, optimize, and decommission.
Deliver in measurable waves	Tie each wave to a working business capability, acceptance criteria, owner, baseline, and KPI.
Operate and improve	Maintain named ownership, proactive monitoring, managed operations, FinOps, and a backlog of continuous improvements.

The service principle. A successful modernization has accelerators for repeatable work, human judgment for business choices, KPIs that measure success, owners who can accept the result, and an operating model that keeps improving after deployment.

The enterprise is the system.

A typical enterprise has dozens or hundreds of systems in its stack. The data platform must work with that ecosystem, not sit beside it. Source applications create the data. Business processes give it meaning. Security determines how it can be used. Cloud infrastructure supports it. Digital channels and applications activate it. Quality engineering validates it. Operations keep it reliable. That is why Zensar's Snowflake work must connect beyond the data platform itself: enterprise applications, cloud infrastructure, security, quality engineering, digital channels, AI, and managed operations all

affect whether modernization can run in the real business. That breadth matters because modernization often reveals issues outside the data platform: weak source quality, brittle interfaces, identity gaps, release constraints, unclear handoffs, and business processes that were never redesigned.

Our point of view. Snowflake provides the control tower. Zensar helps align the business, operating model, ecosystem, and delivery accountability around it.

Zensar combines three layers of capability.

The better story is not the asset list. It is what those assets help the customer do: find waste, make better choices earlier, reduce risk, establish governance, create data products people use, and move AI into production without pretending the platform will do the work by itself.

1. **Enterprise modernization accelerators.** These reduce delivery effort and implementation risk by helping teams discover the estate, analyze legacy ETL and EDW complexity, identify what to retire, accelerate conversion where it makes sense, automate validation, support governance, and make cutover evidence visible.
2. **Packaged platforms.** Platforms such as ZenseAI.Data and ZenseAI.AgentMesh are proof points when they help customers modernize governed data foundations and connect trusted data to operational decisions. They should support the outcome, not compete with it for attention.
3. **Industry-specific solutions.** Zensar brings business context through solutions and patterns for domains such as insurance, healthcare, manufacturing, financial services, retail, and life sciences. That context matters because a claims data product, a manufacturing quality signal, a banking risk measure, and a retail customer view do not modernize the same way.



Why acceleration is key: manual conversion scales the problem. Good acceleration breaks the work apart: what can be scanned, what can be converted, what deserves redesign, what needs business review, and what should finally be retired. That is how speed becomes safer rather than reckless.

Proof should look like outcomes, not adjectives.

Proof should be concrete. What risk went down? What was retired? What moved faster? Which data products went live? What can the business now do that it could not do before? If the evidence cannot answer those questions, adjectives will not help.

What proof can look like

Industry/client	Challenge	What Zensar did	Snowflake use	Impact
Food service Global food services leader	Millions of daily events outpaced legacy reporting; no real-time KPI visibility.	Real-time streaming KPI core and AI-ready governed data.	Snowflake + CoCo, Kafka/Flink, Cortex-ready.	85-90% lower latency; 5x faster decisions; 60% less manual effort.
Banking US regional bank	Highly complex legacy SQL exceeded conventional AI-assisted conversion.	Agentic SQL-to-DBT conversion factory on LangGraph.	CoCo extended with agentic orchestration.	4-6x faster migration; 80%+ complex SQL auto-converted.
Life sciences Project AI	Project data fragmented; status and risk answers slow to find.	Natural-language project intelligence on a medallion core.	Cortex AI / Cortex Analyst + CoCo	PoC validated; path from 1 to ~200 projects; answers in Teams and Power BI.
Retail US home furnishings retailer	Two legacy on-prem warehouses, Netezza and SQL Server, were costly to run.	Consolidation into one cloud EDW; legacy retired.	Snowflake EDW + IICS + MicroStrategy.	45% lower TCO; 20% faster insights; 2 → 0 legacy platforms.
Financial services US financial services firm	Quantum reporting coupled to the OLTP core caused contention.	Decoupled analytical store modeled as a data product.	Snowflake + Cortex Gen AI.	No resource contention; scales concurrency; AI-ready datasets.
Insurance UK insurance provider	No curated ML data or governed path from development to production.	End-to-end MLOps lifecycle with CI/CD automation.	Snowflake ML + Model Registry, Azure DevOps.	Governed, versioned, automated deployment across all environments.
Hi-tech manufacturing Hi-tech US manufacturer	On-prem Teradata and Hadoop blocked connected enterprise insight.	Hybrid-cloud modernization; data-as-a-service across 20+ domains.	Snowflake on AWS + ML.	\$1M OPEX saved; ~10x faster integration; 10+ TB/day, 1+ PB managed.
Semiconductor Global manufacturer	Siloed supply, demand, and yield data; reactive manual exception management.	Agentic alert-to-resolution with predictive risk intelligence.	Snowflake-native GenAI + ZenseAI.Data agents.	10x faster risk detection; 40% lower planner workload; RCA in minutes.

Modernization must close the economic loop.

The business case is not complete when Snowflake goes live. It closes when you remove unnecessary legacy costs, and the new platform delivers measurable operating value.

Most modernization business cases are built from five measurable value levers: lower operating cost, reduced risk, faster delivery, improved AI readiness, and retirement of legacy technical debt. The strongest programs establish a baseline for each one before work begins and measure progress throughout the journey.

Many modernization programs achieve technical go-live but delay decommissioning for months or years. Until legacy platforms, support contracts, duplicate environments, and specialist dependencies are retired, much of the expected financial benefit remains unrealized.

TCO should account for legacy licenses, Snowflake consumption, cloud infrastructure, ETL and orchestration tools, monitoring and governance, specialist labor, duplicate environments, testing, training, compliance, delayed decommissioning, and the cost of correcting unreliable data.

The business case should also consider the cost of delay. Every additional year on a legacy estate carries continuing license expense,

specialist dependency, operational risk, duplicate environments, delayed AI initiatives, slower decision cycles, and opportunities that competitors may already be exploiting.

ROI should be tied to business change. A stronger case shows which costs fall, which risks shrink, which delivery cycles get shorter, which AI use cases can move into production, and which legacy assets can be retired. Faster reporting and higher engineering productivity matter, but only if they are connected to a decision, process, customer experience, compliance obligation, or financial outcome someone owns.

The question is not only whether modernization creates value. The question is whether the organization can afford to defer the value while modernization is postponed.

The financial discipline. Every value measure needs a baseline, a target, an owner, and a date. If ownership, budget, success metrics, and quality checkpoints are not defined, the program has activity, not accountability.

Conclusion

AI will change the business. Readiness is the open question.

Most leaders already know AI will matter. The harder question is whether the enterprise is ready to use it responsibly and repeatedly. AI will expose the enterprise's uncomfortable parts: data people do not trust, processes nobody owns, rules buried in old code, platforms teams avoid touching, technical debt nobody has measured, and opportunities the business cannot yet see.

The stronger companies will be the ones that turn trusted information into better decisions, and then into action faster than their competitors. That advantage will not come from demonstrations. It will come from ownership, quality, governance, economics, and reuse.

Modernization is bigger than migration. AI does not start with the model; it starts with the foundation that lets the business explore, question, trust, govern, and act without rebuilding the same plumbing every time.

Snowflake can provide the cloud-native control tower. Zensar's role is to help the enterprise understand the airport, protect operations during the transition, modernize old ETL and EDW patterns, create data products people use, control cost, and remain accountable after go-live.

Your data is not just another enterprise asset. It is the air traffic control system for the future you are building.

How effectively you modernize, govern, trust, and activate it will determine how confidently you navigate what comes next.

A practical starting point

Start with one high-value business domain. Identify the decision or process that must improve, the data and systems behind it, the repeated preparation effort, the technical debt and dependencies, the expected business outcome, and the owner accountable for the result.

Then ask each prospective Snowflake partner to walk through the path from current state to target outcome. What would they move? What would they refuse to move? Where would they automate? Where would they slow down and ask the business to decide? That is where the real partner difference appears.

Sources and notes

Sources


- [1] Gartner, "Gartner Forecasts Worldwide AI-Optimized IaaS Spending to Grow 96% in 2026," August 2026; forecasts AI-optimized IaaS spending of \$42.276 billion in 2026, up 96.4%, with inference spending surpassing training.
- [2] Snowflake, "The AI Data Cloud Explained," platform overview; describes Snowflake as a unified platform and connected ecosystem for data, apps, and AI.
- [3] Snowflake Documentation, "Snowflake key concepts and architecture"; states that Snowflake uses public cloud infrastructure, has no hardware or software for customers to install or manage, and cannot be installed locally or run on private cloud infrastructure.
- [4] Snowflake, "AI Data Cloud Services Partners," partner-program overview.
- [9] Harness, "FinOps in Focus 2025"; estimates that 21% of enterprise cloud infrastructure spend, or \$44.5 billion in 2025, is wasted on underutilized resources because of weak FinOps and development-team visibility.
- [10] Flexera, "2026 State of the Cloud Report" and related press release; reports wasted cloud spend rising to 29% as AI adoption accelerates cloud cost, governance, and risk demands.
- [11] McKinsey & Company, "Procurement power plays: Getting strategic about cloud," June 2025; states that recent McKinsey research found 28% of cloud spend is wasted for many companies because cloud architecture and consumption are difficult to manage in a demand-driven, continuous way.
- [12] Zensar, "Snowflake Case Studies.pptx," portfolio-at-a-glance and case-study slides; includes nine Snowflake engagements across eight industries with documented outcomes such as latency reduction, faster decisions, TCO reduction, legacy platform retirement, OPEX savings, SQL-to-DBT modernization, MLOps, and Snowflake-native AI use cases.

Authored by

Daniel Gomez,

Practice Head - Data Engineering 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