


Data, Engineering, and Analytics Practice
Snowflake Premier Partner

Claude, Snowflake, and the Governed Intelligence Layer Your Data Already Needs

Why Claude on Snowflake, built and governed by Zensar, protects what an enterprise already has while multiplying what it can do with it: a practical guide for data and technology leaders.

 Whitepaper



Before we start

Every enterprise with a Snowflake investment is asking a version of the same question: how do we put a capable model to work on this data without exposing it to the two things that keep a CDO up at night — a governed asset being exposed and a confident answer that turns out to be wrong? Claude, deployed within Snowflake's AI Data Cloud and governed by Zensar's workflow layer, is our answer to both halves of that question.

Most Snowflake programs do not stall because the technology fails. In our experience running these programs across banking, insurance, healthcare and life sciences, technology, retail, manufacturing, and the public sector, the platform almost always works as intended. What stalls is the translation between what the platform can do and what the business actually needs from it and, increasingly, whether a model can be trusted to touch that data at all.

That gap shows up as the same handful of questions, over and over, on nearly every account we work:

- What should move to Snowflake, and what should stay where it is?
- How do we turn a data foundation into data products that Claude or any model can actually use?

- How do we move gen AI from demo mode to production?
- How do we govern agents that make decisions on their own?
- How do we prove the Snowflake investment is creating business value, not just consuming budget?
- And underneath it all: at what risk, and to whom, if we move fast and get it wrong?

These are not primarily technology questions. They are execution questions that explain why a well-funded Snowflake program can still take twice as long as planned to show up on a P&L.

This paper is our answer to those questions, written the way we'd walk a client through it, not a product pitch or a market survey. It covers where organizations get stuck, a simple way to tell where you are today, the framework we use to move a client through the journey, and specifically how Claude, Snowflake, and Zensar's own execution layer combine to create something none of the three provides alone.

The Thirty Second Version

Snowflake protects the data: access, lineage, and audit down to the table and column levels, across every data type your enterprise uses.

Claude protects the reasoning: a model built with safety as a design principle, for decisions that touch money, patients, or regulatory exposure.

Zensar protects the outcome: the workflow and governance layer that turns a correct answer into an audited action owned by a named person.

Start with a 1-hour workshop, no obligation: see section 8 for details.

That combination is the thesis of this paper. Snowflake protects data at the table and column level, with access, lineage, and audit controls across structured, semi-structured, and unstructured data. Claude protects the reasoning: it is a model built with safety as a design principle, which matters because the questions above concern decisions that affect money, patients, or regulatory exposure, not just convenience.

Zensar protects the outcome: the workflow, governance, and change-management layer that turns a correct answer into an audited action that a named person owns, and keeps it that way as the organization scales. Remove any one of the three and something goes ungoverned: the data, the reasoning, or the business process built on top of it.

Who Should Read This

Chief Data Officers, Chief AI Officers, data platform architects, and VPs of engineering or analytics who have already committed to Snowflake (or are deciding whether to) and want a clear answer on how Claude fits safely into that investment and how to get from platform to measurable business outcomes. If you are looking for a general comparison of LLM vendors, this is not the document for you. If you want to know how Claude, Snowflake, and Zensar's execution layer fit together, and where programs actually go wrong, keep reading.



1. Where are you today?

Before any conversation about architecture or use cases, we ask a client one question: where are you on the journey? Not because it is a nice framing device, but because the advice that is right for a team running its first proof of concept is nearly useless for a team that already has 40

production pipelines and is now trying to govern agents.

We use six stages. Most organizations we meet are somewhere between modernizing and activating: the data cloud is real, but the business has not yet felt its impact.

The Agentic Trust Model

Stage	The question it's actually asking
1. Migrating	Which workloads and data assets do we move to Snowflake, and in what order, without breaking what works today?
2. Modernizing	How do we redesign those workloads for governance, documentation, and reuse, not just move them?
3. Activating	Which business use cases actually create value, and how do we get them into production and into use?
4. Intelligent	How do we responsibly bring Claude and other AIs into pipelines, products, and agent workflows, including Snowflake Intelligence and agentic orchestration?
5. Scaling	How do we operationalize what worked, scale across business units, and optimize costs as usage grows?
6. Proving value	What business metrics demonstrate ROI, in numbers a CFO will accept, not adoption counts?

A few observations from sitting inside enough of these programs to notice the pattern:

Migrating and modernizing get confused. Moving workloads onto Snowflake is migration. Redesigning them so they are governed, documented, and reusable is modernization. Plenty of programs report the first as though it were the second.

Activating is where many programs stall. Teams can show a working Cortex Analyst demo. Getting business users to trust it enough to use it instead of the old report is a change-management problem, not an engineering one.

Intelligence is arriving faster than most governance functions anticipated. Agentic AI, including Snowflake's own agentic and multi-agent capabilities, does not wait for an organization to finish modernizing before appearing as a pilot somewhere in the business. Governance has to be ready earlier than the roadmap assumes. This is also the stage where the choice of model stops being a technical footnote and becomes a decision someone in the room needs to own.

2. How we move a client through it

We use a seven-stage framework for every Snowflake engagement, whether the starting point is a single proof of concept or a multi-year transformation. It is not a menu of services. Each

stage produces something the next stage depends on, and the last stage feeds back into the first, proving value is not an endpoint; it is what justifies the next round of investment.

The Agentic Trust Model

Stage	What it answers
1. Assess	Where are we today — architecture, workloads, cost, data quality, and AI readiness, baselined honestly rather than optimistically.
2. Architect	What should the target data and AI architecture look like, and which workloads genuinely belong on Snowflake versus somewhere else?
3. Modernize	How we migrate and rationalize the existing estate, not just lift-and-shift, but redesign for governance and reuse.
4. Govern	How we make data, models, and agents trusted and controlled through Snowflake Horizon and AI governance tooling, before scale makes it harder to retrofit.
5. Activate	Which business use cases actually create value (chosen by impact, not technical novelty) and get them into production.
6. Scale	How we operationalize what worked, scale it across business units, and continue optimizing cost and performance as usage grows.
7. Prove value	What business metrics demonstrate ROI in numbers a CFO will accept, not in adoption counts alone.

A Few Things We've Learned the Hard Way

- Snowflake Horizon gets implemented as a technical control long before it becomes an operating discipline. The second part is where the value actually lives.
- A successful gen AI demo can hide an unsuccessful gen AI program. The demo proves the model works. It says nothing about adoption, cost at scale, or whether the answer was trustworthy enough for someone to act on.
- Not every workload belongs on Snowflake, and saying so out loud earns more trust than pretending otherwise.

2.1 The stage most programs skip: The people side

Each stage above has a technological and an organizational answer, and most programs budget only for the first. Snowflake can move from lift-and-shift to a governed data cloud in a quarter. The underwriter, the adjuster, or the analyst whose job changes because of it does not adjust on the same timeline. If the operating model around them doesn't change too, the technology can succeed while the transformation fails.

We have seen this play out in a specific, avoidable way. A company automates a meaningful share of a process. The technology

delivers exactly the productivity gain it promised. But the roles, approval chains, and incentives around that process never get redesigned — people keep doing the old workflow alongside the new tool instead of letting it replace part of their day. Retained labor costs absorb the productivity gain, the AI's operating costs, and the rework required to reconcile the two. The net benefit: close to zero. The AI worked. The transformation didn't.

The fix is not a training deck. It is answering six questions for every workflow the platform touches, before go-live rather than after:

Question	What changes
Process	Which steps disappear, change, or get added, not just which system they run in.
People	Which roles change, and what happens to the capacity the automation frees up.
Skills	What people need to learn to work alongside the system, not just use it once.
Decision rights	Who owns the outcome when the agent recommends and a person still approves — named, not assumed.
Governance	Who can approve, modify, or stop the agent, and how an exception actually gets handled.
Incentives	Whether people are actually rewarded for using the new process, or quietly penalized for the short-term disruption it causes.

This is not a separate workstream we bolt on. It is why Govern and Activate, in the framework

above, are where we spend the most time with a client's own people, not just their data.

2.2 Pricing the risk, not just the technology

We evaluate every recommendation against three things at once: (i) what it does for revenue, (ii) what it does for cost, and (iii) what it does to the organization's risk, not sequentially, as if risk were a footnote to check after the business case is approved. Speed makes this sharper, not simpler. Moving faster on a technology decision has real value: faster time-to-insight, faster time-to-market, and a first-mover advantage in a use case a competitor hasn't yet automated. But speed that outruns the people and process work behind it is exactly how the zero-net-benefit

scenario above happens, quietly, eighteen months in.

That is why our ROI models for these programs are not single-year technology paybacks. We evaluate value over multiple years, accounting for when value actually materializes rather than when the project ends, the compounding effect of capabilities reused across use cases, and a risk-adjustment factor that prices in regulatory, brand, and compliance exposure alongside the direct cost of the technology and the transformation around it.



3. What actually makes the difference

The market is converging on a shared architectural vocabulary: governed data, agents, orchestration, and secure access. These capabilities are becoming table stakes. Technology alone does not create business value.

Business context and execution do, and they only show up after you have done the work in a specific industry, with a specific client's data, more than once.

We think about this in three layers.

3.1 Implementation excellence

- **Quality:** Validated, trusted answers, not just fast ones. A wrong answer delivered quickly is still wrong.
- **Interoperability:** Working inside the systems a team already uses, not asking them to adopt a new one to get value from this one.
- **Real-time relevance:** Insight arriving while the decision is still being made, not in next week's report.
- **Industry pattern, not generic template:** An approach proven in claims or fraud or trade operations, not a reference architecture that could apply to anyone.

3.2 Four layers of context

Data alone answers "what happened." Getting to a decision someone can actually act on needs three more layers on top of it:

- **Data:** what happened.
- **Business:** why it's happening.
- **Operational:** what's in the way of acting on it.
- **Industry:** how it compares to what everyone else in the sector is seeing.

3.3 From insight to action

An answer that only describes the past is of little value to an executive under time pressure. We push every engagement toward prescriptive (what to do next, ranked by impact) and

predictive outputs, so a leader can simulate an alternative before committing to it. Done well, every decision made this way becomes a signal that sharpens the next one.

Table Stakes Vs. Differentiation

Governed data, business context, an agent runtime, and secure access are now table stakes. Every serious platform vendor offers these capabilities, Snowflake included. They are no longer differentiators; they are the price of entry. The model itself is not the differentiator either. Section 4 explains why Claude is well suited to

this architecture, but two competitors can deploy the same model on the same platform and deliver very different outcomes. The difference between a services brochure and a program that changes how a business runs is whether the partner is outcome-led or service-led. We build for the former.

4. The foundation, the layer, and the model

This is where most vendor material collapses three distinct elements into a single story: the platform, the orchestration layer that makes it usable for a real business process, and the model that performs the reasoning. They are not the same

thing, and confusing them is exactly how a program ends up with a great model demo and no governed way to put it to work. We separate the three deliberately.

4.1 The foundation: Snowflake AI Data Cloud

Snowflake and Anthropic have entered a strategic partnership that makes the Claude model family accessible within the AI Data Cloud on AWS, Azure, and Google Cloud. That gives you a governed data platform: Snowflake Horizon for cataloging, access control, and lineage; Cortex AI for LLM-powered functions; Cortex Analyst for natural language to SQL; Cortex Search for semantic retrieval; and Snowflake Intelligence, Snowflake's conversational analytics experience, for multi-step reasoning across governed data sources. Claude is one of the models available within this environment, callable through Cortex functions rather than requiring a separately provisioned AI system.

When deployed in a properly configured Snowflake account, model inference for supported Cortex functions runs on Snowflake-managed infrastructure in the customer's designated cloud region. Data handling and residency

characteristics depend on the Snowflake account configuration, the specific Cortex feature in use, and the customer's region settings. Customers should confirm the data processing behavior for their specific account tier, region, and use case with Snowflake before making any external data residency representations. Zensar's delivery approach includes this validation as a standard step in the architecture phase.

As of mid-2026, Claude models available through Snowflake Cortex include Claude 3.5 Sonnet and Claude 3.5 Haiku, with availability varying by Snowflake account region and tier. Anthropic and Snowflake update the model roster on a rolling basis; current availability should be confirmed directly with Snowflake. No separate Anthropic API contract or model infrastructure is required to invoke Claude through Cortex functions.

4.2 The workflow layer: ZenseAI.AgentMesh for Snowflake

The platform gives you governed data and a model that can reason over it. Neither one, by itself, is a business process. The gap between "the model gave a correct answer" and "the business acted on it, with an audit trail, inside an approval chain someone actually owns" is where most AI programs stall. Closing that gap is the job of the enterprise workflow layer we add on top of Snowflake-native AI and governance: ZenseAI.AgentMesh for Snowflake.

It adds reusable agent patterns, context engineering, and orchestration that turn a Snowflake-governed answer into a decision, an exception, an approval, or an audited action delivered through Streamlit in Snowflake or a client's own enterprise applications, without standing up a second data platform. This is the part of the stack that is distinctly ours, not a restatement of what Snowflake or Anthropic already provide.

4.3 Where Claude fits and what's changed in 2026

Within the Snowflake AI Data Cloud, Claude can be invoked through several mechanisms. The primary programmatic path is the `SNOWFLAKE.CORTEX.COMPLETE()` function, which accepts a model name and a prompt and returns a completion that can be used in SQL worksheets, Snowpark Python, and dbt models. Beyond direct completions, Claude also participates in the following Snowflake AI surfaces:

- **Cortex analyst:** Translates natural-language business questions into governed SQL, grounded in a semantic model that Zensar builds and maintains to reflect the client's actual business definitions, not raw table schema.
- **Cortex search:** Semantic retrieval from unstructured documents (PDFs, emails, notes, claim narratives) stored in Snowflake stages or tables, enabling hybrid retrieval-augmented generation (RAG) patterns.
- **Snowflake intelligence:** Snowflake's multi-agent conversational analytics experience, announced and expanded in 2025–2026, enables Claude to reason across multiple governed data sources and agent tools in a single conversation thread, with every action governed by Horizon.
- **Cortex agents and multi-agent orchestration:** Snowflake's agent framework enables Claude to be wired into multi-step workflows, calling tools, executing SQL, and passing context between agents, all within the Snowflake execution environment and under the AI Trust Layer's governance.
- **Snowflake MCP server:** The official integration between Claude Desktop (and other MCP-compatible clients) and Snowflake, enabling Claude to access a Snowflake workspace to query data, inspect schemas, and call Cortex functions within a conversation. Every tool call is subject to the calling user's existing Snowflake RBAC permissions.
- Each of these paths bills against standard Snowflake consumption credits. There is no separate AI API metering for customers invoking Claude through Cortex, and every path is subject to the access control policies already defined in Snowflake Horizon. Claude does not receive elevated

4.4 Two layers of governance, not one

It is worth being precise about what governs what, because the terms are often used interchangeably and are not the same thing. Snowflake Horizon is the governance layer for data and AI assets, including tables, files, models, and functions, as securable objects with access control, lineage, data classification, and privacy policies. Snowflake's AI Trust Layer, built on top of Horizon, serves as the governance surface for AI runtime behavior: model access controls, prompt safety filters, output validation, and auditability of AI-generated actions. Horizon governs what a person or agent can see. The AI Trust Layer

governs what AI is allowed to do and provides an audit trail for every AI-mediated action.

Because this governance architecture applies uniformly whether a Claude call originates from a scheduled Snowflake Task, a Streamlit-in-Snowflake application, a Cortex Analyst query, or a Snowflake Intelligence conversation, an enterprise gets a consistent audit trail rather than a patchwork of separate logging systems. That is a meaningfully different answer to give a CISO than "we'll build a second logging system to reconcile with the first."

4.5 A reference architecture


From data sources to governed agentic workflows, a production Claude-on-Snowflake deployment spans five layers. The exact configuration, particularly data residency, model

availability, and AI Trust Layer settings, should be validated for each customer's Snowflake account tier and region.

Layer	What it contains
Data Sources	Structured (ERP, CRM, transactional), semi-structured (JSON, Parquet, event streams), and unstructured (documents, emails, claims notes), ingested via Snowpipe Streaming, Kafka connector, or native connectors.
Snowflake AI Data Cloud	Tables, Dynamic Tables, Snowpark, Apache Iceberg, governed by Horizon with column-level security, Dynamic Data Masking, and row-level access policies applied uniformly across all data assets.
Semantic and context layer	Governed business metrics, shared entity definitions, Cortex Search indexes, and Snowflake Intelligence agent definitions, the layer that grounds Claude's answers in what the business actually means, not what the raw table schema says.
Claude via Cortex and intelligence	CORTEX.COMPLETE() for direct LLM calls; Cortex Analyst for NL-to-SQL; Cortex Search for semantic retrieval; Snowflake Intelligence for multi-agent conversational analytics; Cortex Agents for orchestrated multi-step workflows, all governed by Horizon and the AI Trust Layer.
ZenseAI Workflow Layer	Reusable agent patterns, approval chains, exception handling, audit trails, and Streamlit in Snowflake applications — the layer that turns a governed answer into an audited business action with a named owner.

The semantic and context layers are worth calling out specifically. Horizon governs access to raw tables and files; the semantic layer is where governed business metrics and shared definitions live. Snowflake's research on Cortex Analyst accuracy shows that grounding NL-to-SQL in a well-defined semantic model yields substantially

better results than querying raw table schemas directly. Building and maintaining that layer for a specific client's business, not the generic version, the one that reflects how this insurer defines a claim or this bank defines a risk-weighted exposure is core delivery work, not a one-time configuration step.





5. Proof, not promises

Frameworks are easy to write. What we'd rather show you is what happened when we ran this playbook on real programs. The examples below are anonymized at the client's request. Where specific metrics are shown, they reflect directional outcomes from Zensar-led engagements; precise figures are available to qualified prospects through a governed discovery process, pending client approval for sharing.

It is worth being precise about what these examples prove: they are Snowflake-platform engagements. They demonstrate the underlying governed pattern, including Horizon, a shared entity model, and an audited path from question to action that a Claude deployment inherits when it runs within this same architecture. The pattern is what carries over, irrespective of which specific Cortex capability or Claude model version powered it.

US general insurance carrier: Intelligent data source onboarding

Baseline: New data source onboarding (mapping external schemas to internal data models) was taking an average of several weeks per source, blocking pipeline delivery and slowing regulatory reporting cycles.

Intervention: Claude, called via Snowflake Cortex COMPLETE() over Snowflake's Information Schema, analyzed incoming schema metadata and generated candidate column mappings with confidence scores, thereby reducing the manual mapping step in the onboarding workflow.

Faster onboarding	Measurable reduction in new data source onboarding time; specific figures available through discovery. The primary driver was the removal of manual schema analysis, not the automation of the full process.
Within-account inference	Inference executed through Cortex functions within the Snowflake account; data residency characteristics confirmed through Zensar architecture review for this client's account configuration.
Audit trail	Cortex function calls logged within Snowflake Query History, providing lineage from source schema to generated mapping for every onboarding run.

Stack: Snowflake AI Data Cloud, CORTEX.COMPLETE(), Information Schema, Snowpark Python

North American bank: Transaction narrative risk annotation

Baseline: The existing rule-based fraud system flagged a narrow range of transaction patterns, missing social engineering and mule account activity that did not match pre-defined rules. Fraud operations lacked a scalable way to analyze transaction narratives at scale.

Intervention: Claude deployed through Snowflake Cortex to analyze payment transaction narratives at ingestion time, generating structured risk annotations alongside each transaction. Analysts reviewed flagged transactions rather than raw data.

New signal class	Narrative-based mule and social engineering indicators surfaced that were not capturable by the prior deterministic rule set.
Analyst prioritization	The fraud operations team shifted from raw transaction review to structured annotation review, improving throughput in investigation queues.
Governed at source	Dynamic Data Masking applied to PII fields before the Claude prompt is constructed; full Horizon access trail maintained throughout.

Stack: Snowflake AI Data Cloud, CORTEX.COMPLETE(), Dynamic Data Masking, Snowpipe Streaming, Streamlit in Snowflake

Specialty insurance market: Intelligent underwriting assistance

Baseline: Underwriters spent the majority of their time assembling structured briefings from broker-submission data, prior claims history, and external risk signals, all of which were stored in separate systems, none of which were unified or consistently formatted.

Intervention: Claude synthesized submission data, claims history, and external signals in Snowflake into a structured underwriting briefing document, automatically generated for each new submission via a Streamlit application in Snowflake.

Faster triage	Underwriters reported spending less time on information assembly and more time on risk judgment. Precise time savings are confidential and subject to client approval before external citation.
Consistent format	Every submission presented in the same structured format regardless of broker, improving cross-submission comparability and reducing data-gathering errors.
No new platform	The application was built and deployed within the client's existing Snowflake environment; no additional vendor procurement was required for this specific engagement.

Stack: Snowflake AI Data Cloud, CORTEX.COMPLETE(), Cortex Search, Streamlit in Snowflake, Snowflake Horizon

Global retailer: Natural language analytics democratization

Baseline: Business teams needed SQL literacy or analyst support to query commercial performance data, creating a bottleneck that slowed reporting cycles and left front-line managers waiting for answers to straightforward questions.

Intervention: Cortex Analyst, grounded in a Zensar-built semantic layer reflecting the client's actual business definitions, deployed as a Streamlit application in Snowflake for non-technical business users.

Self-service insight	Business users received governed, query-grounded answers without requiring analyst intermediation for standard questions.
Role-aware access	Snowflake RBAC is applied uniformly. A regional manager sees only their region's data, regardless of how the natural-language question is phrased.
Full audit trail	Every natural language question and the resulting SQL are logged in Snowflake Query History, preserving traceability between business intent and data retrieval.

Stack: Snowflake Cortex Analyst, Claude, Snowflake Horizon, Streamlit in Snowflake, semantic layer on Snowflake data model


A note on evidence

The outcomes above reflect directional results from Zensar-led engagements. Specific productivity figures, time reductions, and cost savings cited in Zensar sales conversations are available through a structured discovery process and are subject to client approval before external publication. If we cannot substantiate them, we do not publish them. Prospects seeking benchmarked metrics for business case construction should engage Zensar's ZenseAI.Data practice directly.

6. The accelerators behind the framework

Snowflake builds the platform. Anthropic builds the model. Neither one ships an audit trail, an approval chain, or a named accountable owner for the decision an agent makes; that is the layer we build, and it is what the accelerators below actually do.

The Value Framework in Section 2 is the methodology. ZenseAI is the intellectual property that makes each stage faster than starting from a blank page. It is not a product catalog to browse; each accelerator earns its place because it does specific work at a specific stage of the journey.



Framework stage/Zensar IP	What it does at that stage
Modernize ZenseAI.Data	Automated metadata harvesting, semantic layer construction, and validation for legacy data platforms migrating to Snowflake. The accelerator that makes Cortex Analyst and Snowflake Intelligence accurate from the start, not after months of prompt tuning. Recognized as a Snowflake Premier Partner solution.
Modernize ZenseAI.Engineering	Application transformation and code understanding, for the legacy application code a Snowflake program inevitably has to interact with, not just legacy data.
Govern ZenseAI.ADLA	The agentic AI development lifecycle encompasses the delivery and governance methodology for every stage, from design through continuous operation, including Claude-specific prompt governance, AI Trust Layer configuration, and output validation frameworks.
Govern ZenseAI.AssureAI	Quality engineering built for AI: continuous validation of data integrity, model behavior, and agent decisions, aligned to NIST AI RMF, ISO/IEC 42001, and the EU AI Act. Covers both Cortex-native AI functions and ZenseAI.AgentMesh agent actions.
Activate ZenseAI.AgentMesh	Building and governing multi-agent systems on Snowflake, with pre-built agent patterns across industry and cross-functional use cases. Integrates with Snowflake Intelligence, Cortex Agents, SAP, Salesforce, ServiceNow, and Databricks.
Activate ZenseAI.Guidewire	Underwriting, policy, and billing modernization specifically for P&C insurers on Guidewire, drawing on a practice with 40-plus purpose-built accelerators, grounded in Snowflake data where applicable.
Activate ZenseAI.Content	Regulated content creation and review for Life Sciences, built for the audit and approval requirements pharma, payer, and provider content teams operate under.
Scale ZenseAI.FinOps	Cost visibility and chargeback for AI and LLM spend, the piece most programs add too late, after Cortex credit consumption has already outpaced the value it is producing. Covers token-level cost attribution across Cortex functions and Intelligence workloads.

Where Most Programs Start

Most enterprise Snowflake modernizations start with the same three: ZenseAI.Data to build the semantic layer and modernize the estate, ZenseAI.AgentMesh to build and govern agents, and ZenseAI.AssureAI to keep both trustworthy as they scale. Together they cover the hardest parts of enterprise AI: trusted data and safe, governed deployment, which is why we recommend them as the default starting point rather than the full list.

7. Why Zensar

None of the above matters if the team behind it cannot deliver at the scale a real enterprise program needs. A few numbers, as of this writing

(figures subject to verification through Zensar Alliances before use in formal sales materials):

10,700
employees

1,500+
data consultants

100+
Snowflake
certifications

130+
global
customers

Zensar is a Snowflake Premier Partner, and ISG, Avasant, and Everest Group rate our work. Our delivery teams have built Claude-powered applications on Snowflake. This paper draws on experience, rather than a formal Anthropic

partnership credential. Any credential claims beyond Snowflake Premier Partner status should be verified with Zensar Alliances and Marketing before use in any customer-facing context.

7.1 What customers say they chose us for

- Deep industry expertise across every vertical we serve, delivered by a practice focused on data and analytics work, not a generalist team handling whatever comes in.
- Proven modernization programs on cloud data platforms, not our first.
- Strong data engineering and AI delivery capability, not just advisory.
- Accelerators that measurably speed Snowflake adoption, covered in Section 6.
- End-to-end delivery, from strategy through managed services, including build, operate, transfer, or a combination, depending on what the client actually needs.


7.2 The platform-neutrality argument

This is probably our least expected differentiator, and possibly our strongest. We don't start an engagement by asking whether Snowflake or Databricks should win. We start by asking what architecture actually serves the business in front of us. Our job is to make already-made work harder, not to force a client into a new platform strategy just because it's the one we're certified in this quarter.

That matters because the modern enterprise stack is not simplifying. It is completely normal for one organization to run Snowflake alongside a data lakehouse, use multiple hyperscalers by business unit, and use multiple foundation model providers. An integration strategy built on the

assumption that heterogeneity is temporary creates rework the moment the stack evolves, which, in our experience, is always sooner than the roadmap assumed.

We can say this with confidence because it is not a slide claim; it is how the practice is actually built. Beyond Snowflake and Databricks, Zensar holds active technology partnerships with Microsoft Azure, AWS, and Google Cloud; automation platforms such as Automation Anywhere; and application ecosystems such as Salesforce and Guidewire. We are not trying to move every client onto one stack. We are trying to make the stack they already use work harder.



8. How we work together

We deliberately keep our engagement low-friction because the biggest risk in a first conversation with a new partner isn't a bad

recommendation; it's a slow, expensive way to find out whether the fit is right.

8.1 The first conversation: A workshop, not a pitch

Our door-opener is a half-day executive workshop built around one question: is your Snowflake investment ready for AI? We walk through platform maturity, governance, data products, AI readiness, and the operating model and economics behind them, with the CIO, CDO, or data leader in the room, not just their team. Model choice is part of that conversation, Claude

included, but it comes after we understand the use case, not before. You leave with five specific things: a maturity score against the stages in Section 1; a prioritized list of candidate use cases; the governance gaps standing between here and Activate; an ROI hypothesis built on the approach in Section 2.2; and an outline of a 90-day proof-of-value plan.

8.2 A low-risk path from first call to scale

- Discovery workshop: Align on priority domains, pain points, and target business outcomes.
- Current-state assessment: Baseline maturity, tooling, and architecture across the Snowflake landscape.
- Vision and roadmap: Define the Snowflake target state and a phased, value-driven plan to get there.
- Proof of value: Prove measurable impact fast, on a single use case, before either side

commits further. Zensar co-invests in PoC efforts for qualifying engagements; terms are confirmed through Zensar sales.

From there, the onboarding is intentionally light: rapid, low-friction discovery to get teams moving; a complimentary initial engagement that proves value with no cost commitment; dedicated advisory hours committed up front once the fit is clear; and a defined path to a larger transformation program. Land small, prove value fast, then expand.

9. Where to start, by industry

Banking and Financial Services

- Transaction narrative analysis and risk annotation for fraud operations
- Automated regulatory reporting narrative generation (BCBS 239, AML SAR drafting)
- Natural language querying of risk and performance data via Cortex Analyst and Snowflake Intelligence

Insurance


- Underwriting briefing generation from submission data
- Claims triage and severity classification from unstructured loss descriptions
- Policy document entity extraction and obligation mapping, integrated with Guidewire where relevant

Manufacturing and Supply Chain

- Maintenance log analysis and predictive failure narrative generation
- Supplier risk assessment from unstructured procurement and news data
- Production quality annotation and defect pattern narration

Healthcare and Life Sciences

- Clinical, safety, and market-access analytics unified across EDC, CTMS, CRM, and EMR/EHR data through Snowflake Horizon
- Clinical decision support grounded in governed Snowflake data products
- Payer and provider data quality management and reconciliation across claims and clinical systems
- Regulatory-compliant content review and adverse-event pattern detection from unstructured safety data





Technology and Media

- Customer lifecycle and hyper-personalization data pipelines
- Usage and telemetry analytics for product and engineering teams
- Configure-price-quote and content generation workflows grounded in governed product data

Retail

- Market basket and demand forecasting on governed transaction data
- Inventory and route optimization across store and distribution networks
- Conversational store and contact-center assistants grounded in governed product and policy data

Public Sector and SLED

- Unifying citizen and constituent data across agencies on a single governed Snowflake platform
- Data-driven policy analytics and predictive modeling for service delivery
- Compliance-first AI adoption where procurement and audit requirements are as demanding as any regulated industry
- Three things get a program started with us:
- An active Snowflake account with Cortex AI features and Snowflake Horizon enabled. We can assist with feature enablement if any features aren't active yet.
- A defined business use case with a measurable success criterion. If you don't have one yet, that's what the discovery workshop is for.
- Access to the relevant data within Snowflake. Our accelerator framework handles schema analysis, semantic layer design, prompt and tool design, and integration architecture.

Closing thought

Snowflake delivers the platform. Claude is built to reason over it responsibly, ensuring safety as a design principle, not an afterthought, which matters most exactly where the stakes are highest: underwriting, claims, patient data, regulatory exposure. What turns that pairing into a business outcome is industry and operational context, and the judgment that only comes from having done this more than once.

We've found that the technology on either side of that pairing is rarely the limiting factor in these programs. The limiting factor is almost always the translation between what the platform and the model can do together and what the business is actually trying to accomplish.

That is the practical promise underneath everything in this paper. Claude and Snowflake


are already built to work well together. The partnership is real, and we do not take credit for it. What we do is make sure that pairing protects the investment you have already made in both platforms, and gets more out of it every quarter, not just at go-live. That is the specific thing to bring us in for: not a Snowflake migration, not a model evaluation, but the Claude-on-Snowflake integration that is architected, governed, and operated so it keeps paying off as your data, your regulations, and your use cases keep changing.

If this paper is useful, the next conversation is easy to start: a half-day workshop, no obligation, built around the five questions at the top of this paper. We'd rather spend that time finding out whether we're the right fit than convincing you we are.

"Snowflake delivers the platform. Claude delivers the reasoning. We deliver the context that turns both into a business outcome you can trust."

— Data, Engineering and Analytics Practice, Zensar Technologies

Reach the Data and AI practice at dataaipresales@zensar.com or www.zensar.com.



About Zensar Technologies

Zensar Technologies is an experience-led digital and AI engineering partner — we don't just advise, we build and run. At scale and in depth, that means 10,700 people, 24x7 operations, and a dedicated Snowflake practice that takes clients from strategy through managed services: build-operate-transfer, turnkey transformations, centers of excellence, GCCs and ODCs, knowledge academies, and continuous data and AI operations.


Zensar is a Snowflake Premier Partner and holds partner status with Microsoft, AWS, and Google Cloud. Our client portfolio spans banking and financial services, insurance, healthcare and life sciences, technology and media, retail, manufacturing, and public sector, delivered across the Americas, EMEA, and South Africa — with nearshore delivery capability from Colombia and Mexico for clients who need Americas-timezone coverage without an offshore handoff. ISG, Avasant, and Everest Group have rated our work.

Disclaimer: This whitepaper is intended for informational purposes only. The use cases, patterns, and directional outcomes described are based on Zensar's delivery experience and publicly available information about the Snowflake and Anthropic platforms. Client engagements referenced are anonymized; specific metrics are subject to client approval before external citation. Actual results vary based on client-specific data environments, Snowflake account configuration, use case characteristics, and organizational factors. Data residency and inference location characteristics depend on the specific Snowflake account tier, region, and Cortex feature in use and should be confirmed with Snowflake for each deployment. Snowflake and Claude are trademarks of Snowflake Inc. and Anthropic PBC, respectively. Zensar Technologies is a Snowflake Premier Partner. Anthropic partnership credentials referenced in any Zensar sales or marketing materials should be verified through Zensar Alliances prior to use. Nothing in this paper should be read as a claim of formal partner status with Anthropic PBC.



By Vivek Divakar

Principal Solution Architect, Data,
Engineering and Analytics, Zensar Technologies



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