

Data Strategy, Analytics, Warehousing, and Governance Consulting

Transforming Enterprise Data Capabilities for the Cloud-Native Era

Presented by: Shiv Iyer

Founder and CEO – MinervaDB Inc.

shiv@minervadb.com

linkedin.com/in/thewebscaledba

Executive Summary

Strategic Data Consulting That Delivers Measurable Results

What We Do:

MinervaDB provides vendor-neutral consulting for data strategy, cloud warehousing, advanced analytics, and governance. We architect scalable, compliant data platforms that transform raw information into competitive advantage.

Who We Serve:

Enterprise organizations migrating to cloud data platforms, scaling analytics capabilities, or establishing governance frameworks. Our clients span financial services, e-commerce, healthcare, and technology sectors.

Expected Outcomes:

- 40-60% reduction in query response times
- 25-45% decrease in infrastructure costs
- 70% faster time-to-insight for business users
- Regulatory compliance with automated controls
- Clear ROI within 6-12 months

This presentation outlines our methodology, platform expertise, and proven track record helping organizations build data infrastructure that scales with business growth.

The Business Challenge

Enterprise leaders face concrete obstacles that prevent them from capitalizing on data investments:

- Legacy systems cannot handle the volume and velocity of modern data sources
- Siloed data across departments creates inconsistent reporting and conflicting insights
- Cloud migration projects stall due to unclear platform selection criteria
- Analytics teams spend 80% of their time preparing data instead of generating insights
- Compliance requirements grow more complex while governance remains manual and reactive
- Infrastructure costs spiral as workloads scale without proper optimization

These challenges compound over time. Organizations that delay modernization fall further behind competitors who have built agile, scalable data capabilities.

80%

Wasted Analyst Time

Spent on data preparation vs. analysis

3-5x

Cost Overruns

Common in unoptimized cloud migrations

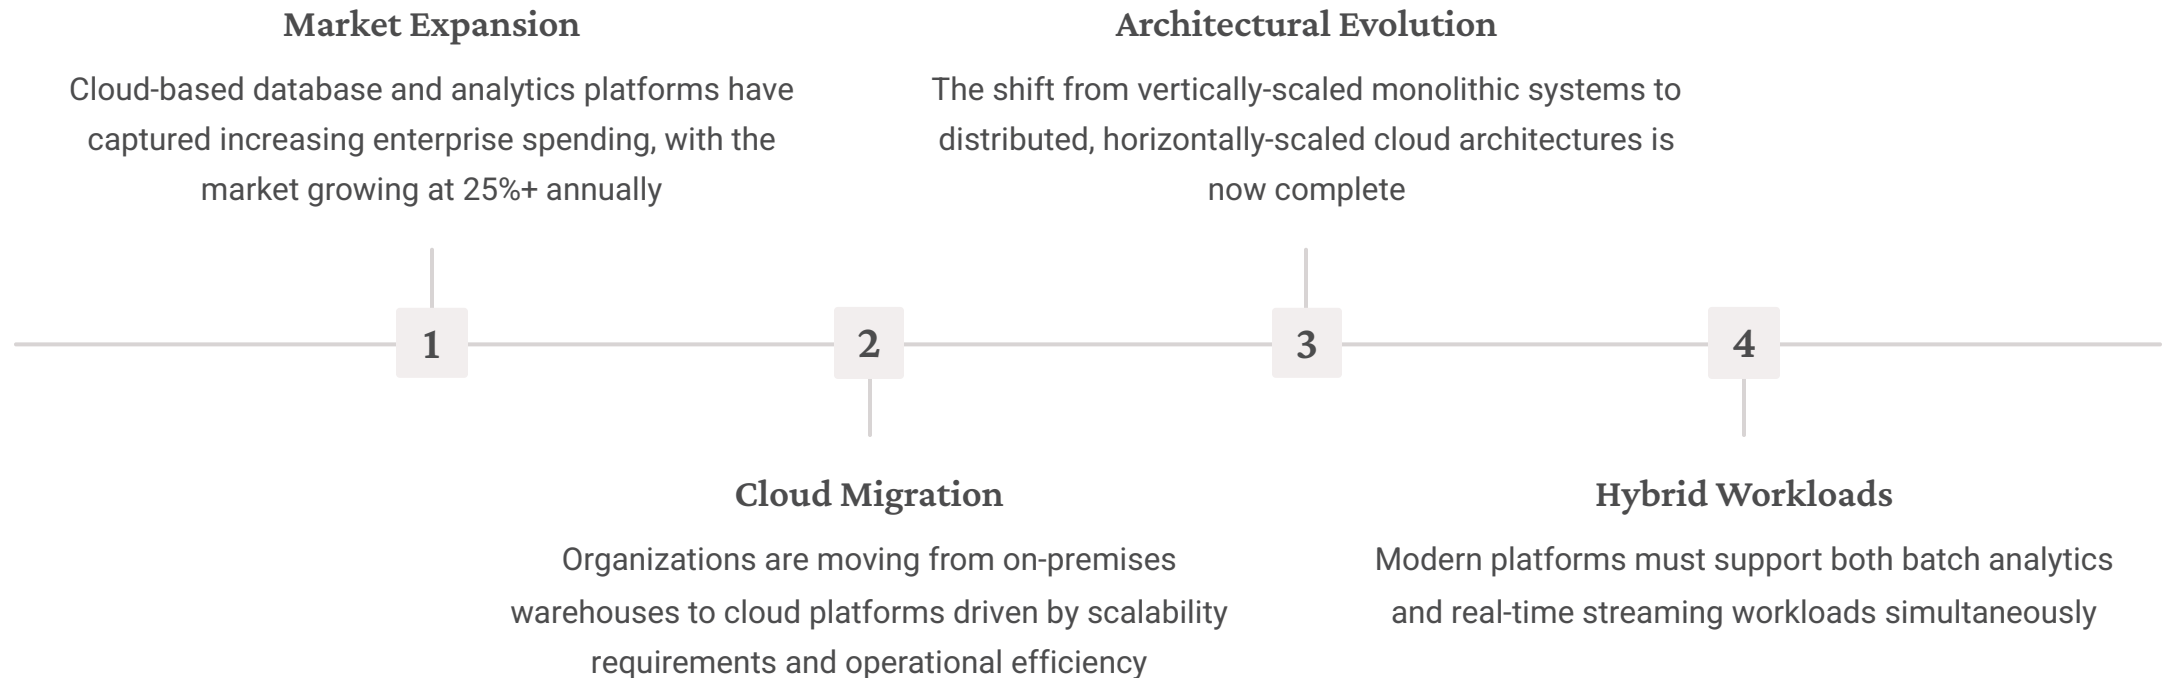
62%

Failed Initiatives

Data projects that don't meet objectives

The Evolving Data Landscape

Enterprise data architecture has shifted fundamentally over the past decade. Cloud-native platforms now deliver elasticity, performance, and cost efficiency that legacy systems cannot match.



This evolution creates opportunity but also complexity. Organizations must navigate competing vendor claims while ensuring technology choices align with business requirements and technical constraints.

Key Industry Challenges

Organizations encounter recurring obstacles when attempting to modernize data infrastructure:

Technology Fragmentation

Multiple specialized databases and analytics tools create architectural complexity. Organizations struggle to integrate point solutions, leading to operational inefficiency and increased maintenance costs.

Persistent Data Silos

Departments deploy independent systems that prevent enterprise-wide analysis. Silos duplicate storage costs, create conflicting reports, and undermine data-driven decision-making.

Governance Gaps

Regulatory requirements like GDPR and CCPA demand strict data controls. Organizations without mature governance processes face compliance violations, quality issues, and financial penalties.

Skills Shortages

Demand for data engineers and architects with cloud platform expertise exceeds supply. Talent gaps create bottlenecks that delay initiatives and limit return on technology investments.

MinervaDB's Role and Positioning

Specialized Data Consulting

MinervaDB specializes in enterprise data modernization. We architect cloud data warehouses, implement real-time analytics platforms, and establish governance frameworks that scale with organizational growth.

Our consultants average 15+ years of experience with production data systems. We have implemented solutions across Snowflake, BigQuery, Databricks, Redshift, ClickHouse, and Apache Druid—enabling us to match platforms to workload requirements rather than applying standard templates.

This depth of experience translates to faster implementations, fewer costly mistakes, and infrastructure that performs as designed.

Vendor-Neutral Consulting

MinervaDB maintains no financial relationships with platform vendors. We receive no referral fees, partnership incentives, or implementation bonuses. Our recommendations are based solely on your business requirements and technical constraints.

This independence matters when platform selection can represent millions in annual spending. We evaluate solutions objectively, negotiate contracts on your behalf, and ensure you avoid vendor lock-in.

Our compensation comes from consulting fees, not vendor commissions—aligning our incentives with your long-term success.

Proprietary Methodologies and Technical Frameworks

Intellectual Property That Accelerates Implementation

Assessment Frameworks

MinervaDB has developed proprietary assessment methodologies refined across hundreds of engagements:

- **Workload Profiling Matrix:** Systematic classification of query patterns, data volumes, and latency requirements to match optimal platforms
- **TCO Modeling Engine:** Comprehensive cost projection accounting for compute, storage, data transfer, licensing, and operational overhead across platforms
- **Migration Risk Scorecard:** Quantitative assessment of technical complexity, organizational readiness, and execution risk factors
- **Performance Benchmarking Suite:** Standardized test harnesses for comparing platform performance with client-specific workloads

Optimization Techniques

Our consultants apply advanced optimization techniques developed through production experience:

- **Query Optimization Patterns:** Systematic approaches to rewriting queries, restructuring joins, and leveraging platform-specific features for 10-100x performance gains
- **Data Modeling Strategies:** Dimensional modeling, denormalization patterns, and materialized view strategies optimized for cloud warehouses
- **Cost Optimization Playbook:** Resource right-sizing, workload scheduling, storage tiering, and query result caching strategies reducing costs 30-50%
- **Concurrency Management:** Connection pooling, query queuing, and workload isolation techniques enabling high-concurrency user-facing analytics

These frameworks represent years of production experience distilled into repeatable processes. Clients benefit from proven approaches rather than experimental implementations, reducing risk and accelerating time-to-value.

Why Organizations Choose MinervaDB

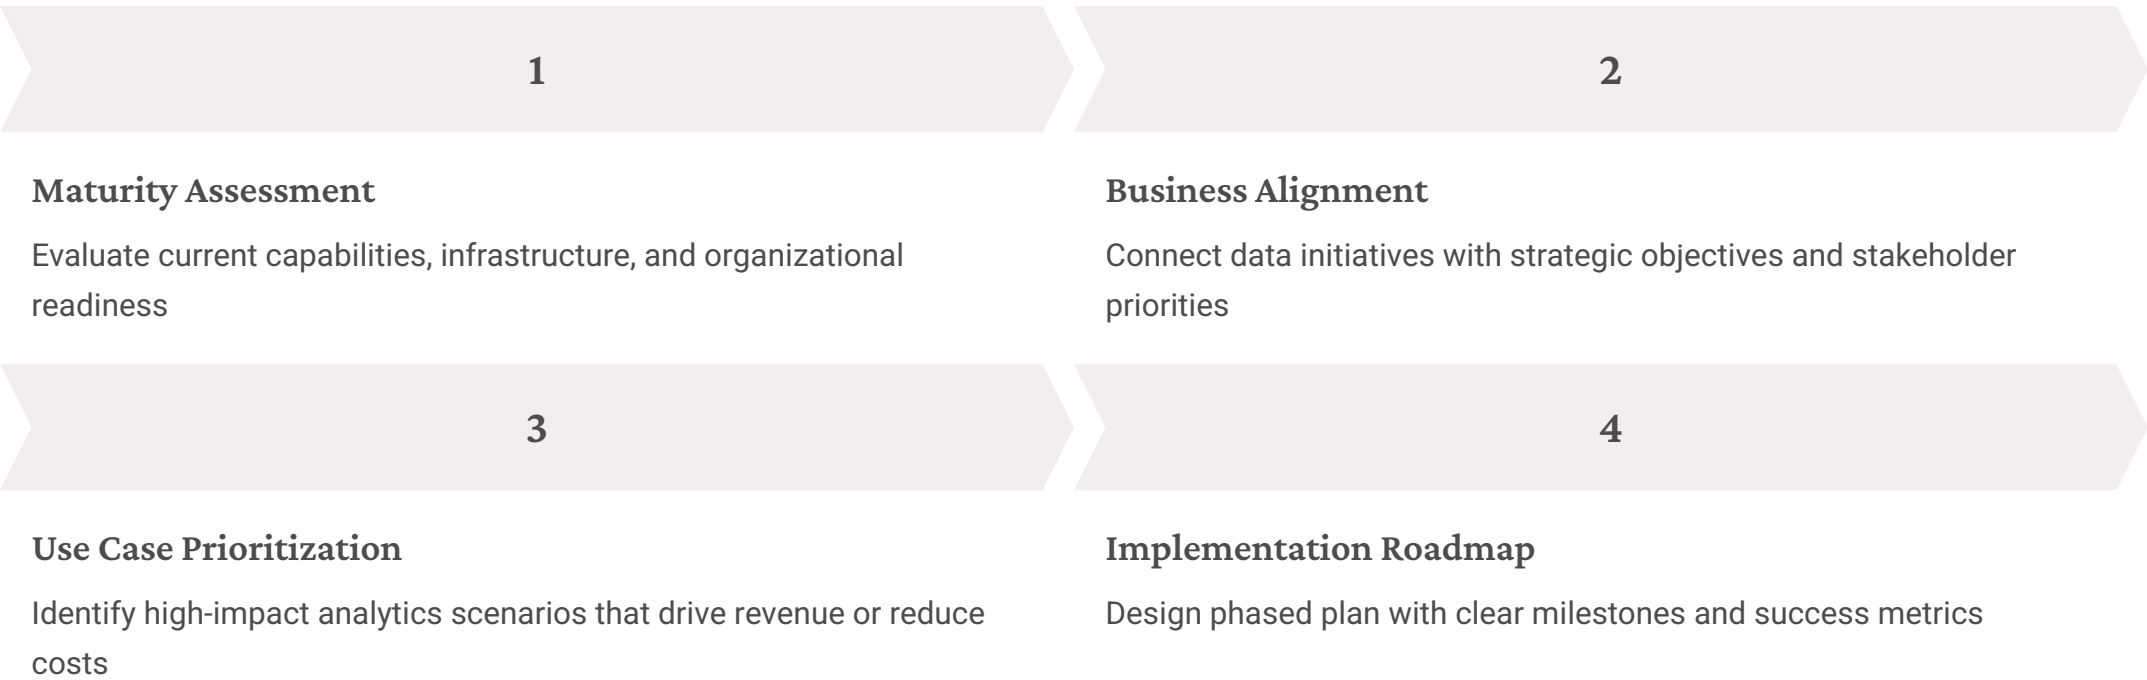
Clear Differentiation in a Crowded Consulting Market

Capability	Large Consulting Firms	MinervaDB
Platform Expertise	Generalist consultants with broad IT knowledge	Specialists with 15+ years in production data systems
Vendor Relationships	Partnership agreements and referral incentives	Zero vendor relationships—completely independent
Implementation Approach	Standard frameworks applied across clients	Custom architecture based on workload analysis
Team Composition	Junior consultants supervised by senior staff	Senior practitioners on every engagement
Engagement Model	Large teams with high overhead costs	Lean teams focused on knowledge transfer
Post-Implementation	Handoff to support teams	Ongoing advisory relationships and optimization

Our clients choose MinervaDB when they need specialized expertise, objective recommendations, and consultants who remain accountable for long-term results.

Comprehensive Data Strategy Consulting

Data transformation requires strategy before technology selection. MinervaDB's methodology ensures infrastructure investments support business priorities and deliver measurable returns.



Strategic Focus: We prioritize measurable business value over technical complexity. Every recommendation includes ROI projections, risk assessment, and implementation feasibility.

Strategy engagements typically span 6-12 weeks and produce actionable blueprints for implementation. We remain engaged throughout execution to ensure alignment with strategic objectives.

Engagement Model and Delivery Framework

Structured Approach with Predictable Outcomes

01	02	03
Discovery & Assessment (2-3 weeks) Current state analysis, stakeholder interviews, workload profiling, and requirements documentation. Deliverable: Assessment report with findings and recommendations.	Architecture & Planning (3-4 weeks) Platform selection, solution design, migration strategy, and implementation roadmap. Deliverable: Technical architecture document and project plan.	Implementation (8-16 weeks) Infrastructure deployment, data pipeline development, integration, and testing. Deliverable: Production-ready data platform with documentation.
04	05	
Optimization & Tuning (2-4 weeks) Performance optimization, cost analysis, query tuning, and workload balancing. Deliverable: Optimized system meeting performance and cost targets.	Knowledge Transfer (Ongoing) Team training, documentation, runbooks, and operational handoff. Deliverable: Internal team capable of independent operations.	

Typical Engagement Duration: 4-6 months for full implementation

Team Size: 2-4 senior consultants

Client Involvement: Weekly stakeholder reviews, technical team collaboration

Flexible Engagement Models:

- Full implementation partnership
- Advisory-only consulting
- Proof-of-concept projects
- Ongoing optimization retainers

Return on Investment Framework

Quantifying the Business Case for Data Modernization

Cost Reduction Opportunities

- Infrastructure Optimization: 25-45% reduction in cloud compute and storage costs through workload optimization and right-sizing
- Operational Efficiency: 40-60% decrease in data pipeline maintenance time through automation and modern tooling
- Analyst Productivity: 70% reduction in time spent on data preparation, freeing analysts for higher-value work
- Vendor Consolidation: 30-50% savings by replacing multiple point solutions with integrated platforms
- Compliance Risk: Avoid regulatory penalties ranging from \$50K to \$20M+ through automated governance controls

Typical ROI Calculation:

Annual benefit (cost savings + revenue impact): \$2M - \$8M

Implementation investment: \$400K - \$1.2M

Payback period: 6-12 months

3-year ROI: 300-600%

Revenue Enhancement

- Faster Time-to-Market: Launch data products 3-5x faster with modern infrastructure
- Customer Insights: Improve conversion rates 15-25% through real-time behavioral analytics
- Operational Intelligence: Reduce churn 10-20% by identifying at-risk customers earlier
- New Revenue Streams: Monetize data assets through analytics products and services

These figures represent median outcomes across our client base. Actual results depend on organization size, current infrastructure maturity, and implementation scope.

From Vision to Roadmap

Stakeholder Collaboration

Data strategies require alignment across executive leadership, business units, and technical teams. We facilitate workshops and interviews to surface priorities and build consensus.

Key activities:

- Executive interviews to define strategic goals
- Business unit sessions to identify pain points and use cases
- Technical assessment of existing infrastructure
- Data quality and governance reviews

Goal Definition Framework

Customer Experience

Personalization, recommendations, churn prediction, sentiment analysis

Operational Efficiency

Supply chain optimization, predictive maintenance, workforce analytics

Innovation & Growth

New product development, market expansion, competitive intelligence

The output is a prioritized roadmap that sequences initiatives based on business impact, technical dependencies, and resource constraints. This roadmap guides multi-year platform evolution.

Advanced Analytics: From Descriptive to Prescriptive

Organizations progress through increasingly sophisticated analytics capabilities, each delivering greater business value:



Descriptive Analytics

What happened? Historical reporting, dashboards, and KPI tracking.
Foundation for data-driven culture.

Diagnostic Analytics

Why did it happen? Root cause analysis, drill-down capabilities, and anomaly detection.

Predictive Analytics

What will happen? Forecasting, machine learning models, and propensity scoring to anticipate outcomes.

Prescriptive Analytics

What should we do? Optimization algorithms, recommendation systems, and automated decisioning.

Our analytics practice specializes in user-facing analytics requiring sub-second response times and high concurrency. These demanding workloads require specialized platform architectures and optimization techniques.

Real-Time and User-Facing Analytics

Powering Interactive Experiences at Scale

Digital products embed analytics directly into user workflows, requiring real-time processing and instantaneous query responses. We architect solutions that maintain performance under extreme concurrency and data volumes.



Product Analytics

Embedded dashboards showing user engagement, feature adoption, and behavior patterns updated in real-time



Ad Tech Optimization

Bid decisioning and campaign performance monitoring requiring millisecond-latency queries across billions of records



Customer Behavior Tracking

Session replay, funnel analysis, and cohort comparison supporting rapid product iteration



Operational Intelligence

Live monitoring dashboards for infrastructure, application performance, and business metrics with alerting

Sub-second query performance at high concurrency requires careful platform selection, data modeling, indexing strategies, and query optimization. We have implemented these patterns across diverse industries and workload profiles.

Scalable Data Warehousing – Core Platforms

We maintain deep expertise across leading cloud data warehouse platforms, enabling optimal technology matching:

Platform	Key Strengths	Typical MinervaDB Use Cases
Snowflake	Compute-storage separation, multi-cloud portability, zero-copy cloning, data sharing	Enterprise-wide analytics hubs, cross-organizational data sharing, workload isolation with separate virtual warehouses
Google BigQuery	Serverless architecture, ML integration, GCP ecosystem, flat-rate pricing	Rapid prototyping, ML-driven analytics, organizations standardized on Google Cloud Platform
Amazon Redshift	AWS integration, Spectrum for data lake queries, predictable pricing	AWS-centric architectures, migration from traditional data warehouses, mixed batch and interactive workloads
Azure Synapse	Unified analytics, Power BI and Azure ML integration, dedicated SQL pools	Microsoft-aligned enterprises, integrated BI and analytics workflows, Azure data estate consolidation
Databricks Lakehouse	Delta Lake ACID transactions, unified batch and streaming, ML workflows, open formats	Advanced analytics on data lakes, ML pipeline integration, organizations requiring both BI and data science capabilities

Platform selection depends on query patterns, data volumes, latency requirements, ecosystem alignment, and total cost of ownership.

Real-Time Analytics Engines

Workloads demanding sub-second latency, real-time ingestion, or extreme concurrency often require specialized analytics engines. We help organizations integrate these platforms strategically:



ClickHouse

Sweet spot: Time-series analytics, log aggregation, event tracking
Columnar storage with exceptional compression. Aggregates billions of rows in milliseconds. Open-source with commercial support options.



Apache Druid

Sweet spot: High-concurrency dashboards, exploratory analytics, streaming ingestion
Purpose-built for user-facing analytics with sub-second queries. Native support for approximate algorithms and real-time streams.



Apache Pinot

Sweet spot: User-facing analytics at massive scale, real-time dashboards
Developed at LinkedIn for trillion-row datasets. Optimized for OLAP queries with high concurrency and low latency.



SingleStore

Sweet spot: HTAP workloads, operational analytics, mixed transactional and analytical
Distributed SQL supporting row and columnar storage. Enables real-time analytics on operational data without ETL.



StarRocks

Sweet spot: Modern lakehouse architectures, high-performance queries on data lakes
MPP architecture with vectorized execution. Strong object storage integration for cost-effective analytics on open formats.

Platform selection requires rigorous analysis across multiple dimensions. We employ a structured assessment methodology:

Technical Requirements

- **Query patterns:** Interactive vs. batch, complexity, concurrency expectations
- **Data volumes:** Current scale, growth projections, retention policies
- **Latency needs:** Real-time, near-real-time, or batch acceptable
- **Workload mix:** BI reporting, data science, operational analytics, streaming
- **Integration points:** Existing tools, data sources, downstream consumers

Business Constraints


- **Compliance mandates:** Data residency, audit trails, encryption requirements
- **Budget parameters:** Capital vs. operational expense, predictability preferences
- **Skills availability:** In-house expertise, training capacity, managed service appetite
- **Cloud strategy:** Multi-cloud, cloud-specific, hybrid considerations
- **Time to value:** Urgency of delivery, phased vs. big-bang approach

The outcome is a platform recommendation matrix documenting fit scores, implementation complexity, TCO projections, and risk factors. This evidence-based approach builds stakeholder confidence and establishes clear success metrics.

Robust Data Governance Frameworks

Balancing Agility with Control

Data governance must enable innovation while ensuring compliance, quality, and security. We help enterprises establish governance operating models that scale with organizational growth:



Strong governance reduces risk, improves data quality, accelerates compliance, and builds stakeholder trust. Organizations with mature governance experience fewer data incidents and faster time-to-insight.

Operationalizing Governance with Modern Tooling

Manual governance cannot keep pace with modern data ecosystems. We implement automated governance platforms that embed controls throughout the data lifecycle:

01	02	03
Data Discovery and Classification	Policy Definition and Enforcement	Access Request Workflows
Automated scanning identifies sensitive data (PII, financial, health) and applies appropriate tags and policies	Centralized policy authoring with automatic propagation to databases, data lakes, and analytics platforms	Self-service portals for data access requests with automated approval routing and audit logging
04	05	
Continuous Monitoring	Compliance Reporting	
Real-time detection of policy violations, unusual access patterns, and data quality anomalies	Automated audit reports demonstrating regulatory compliance and control effectiveness	

Leading governance platforms—including Collibra, Alation, and open-source alternatives—provide these capabilities. We help select, implement, and optimize these tools for seamless integration with your data infrastructure.

Technology Stack Alignment & Vendor Neutrality

Our Philosophy: Technology Serves Strategy

MinervaDB maintains no financial relationships with platform vendors. We receive no referral fees or partnership incentives. This independence ensures recommendations remain unbiased and aligned with client interests.

We guide organizations through complex migrations from legacy warehouses—including Teradata, Oracle Exadata, SAP HANA, and Netezza—to cloud-native platforms. These transitions require workload analysis, TCO modeling, risk mitigation, and phased execution.

We architect for flexibility, enabling you to adopt emerging technologies as requirements evolve without vendor lock-in or proprietary format constraints.

Migration Expertise

1

Assessment & Planning

Workload profiling, compatibility analysis, success metrics definition

2

Proof of Concept

Validate performance, test critical queries, quantify improvements

3

Phased Execution

Iterative migration minimizing business disruption and risk


4

Optimization & Tuning

Post-migration refinement achieving target KPIs


Client Success: Performance and Cost Outcomes

Our consulting engagements consistently deliver measurable results. Clients experience substantial improvements across performance and financial metrics:




Query Performance Improvement

Average reduction in query execution time following platform optimization and migration




Total Cost of Ownership Reduction

Cost savings achieved through cloud-native platform adoption and resource optimization



Concurrency Scalability

Increase in simultaneous user capacity for user-facing analytics applications



Faster Time-to-Insight

Reduction in time from data ingestion to actionable analytics delivery

"MinervaDB enabled us to achieve sub-second query performance on datasets exceeding 50 billion rows while reducing our infrastructure costs by half. Their expertise in real-time analytics platforms transformed our product capabilities."

— VP of Engineering, SaaS Analytics Provider

These outcomes result from deep technical expertise, proven methodologies, and focus on aligning technology investments with business value.

Case Study – Global E-commerce Modernization

Databricks Lakehouse Migration Driving Operational Excellence

Client Profile

A multinational e-commerce platform processing millions of daily transactions across multiple regions, struggling with siloed data systems and delayed inventory visibility.

Business Challenge

- Fragmented analytics across regional warehouses
- Hours-long ETL processes delaying decision-making
- Inability to track inventory in real-time
- Poor visibility into supply chain performance

MinervaDB Solution

Architected and implemented a unified Databricks Lakehouse consolidating data from disparate sources, enabling batch and streaming analytics on a single infrastructure.

Key Technical Components

- Delta Lake for ACID transactions on data lake
- Streaming ingestion from transaction systems
- Medallion architecture (bronze/silver/gold layers)
- ML pipelines for demand forecasting
- Unified governance and access controls

70%

ETL Time Reduction

From hours to minutes for critical reporting pipelines

15%

Fulfillment Efficiency Gain

Improvement in order-to-delivery cycle time

Real-Time

Inventory Visibility

Live tracking across warehouses and distribution centers

Case Study – Financial Services Governance

Enterprise-Wide Data Governance Reducing Risk and Accelerating Compliance

Client Profile

A regional financial services institution managing sensitive customer data across lending, investment, and insurance business units, facing regulatory scrutiny and audit complexity.

Business Challenge

The organization lacked centralized governance, resulting in inconsistent data definitions, duplicated customer records, manual compliance reporting, and audit preparation cycles consuming weeks of staff time.

MinervaDB Solution

Designed and implemented a comprehensive data governance framework including:

- Data catalog with business glossary and lineage tracking
- Automated data classification and tagging
- Role-based access controls aligned to business units
- Policy-driven masking for sensitive data elements
- Continuous monitoring and compliance dashboards

Business Outcomes



Data Accuracy Improvement

Customer record quality metrics



Audit Prep Time Reduction

From weeks to days



Regulatory Compliance

Full adherence to mandates

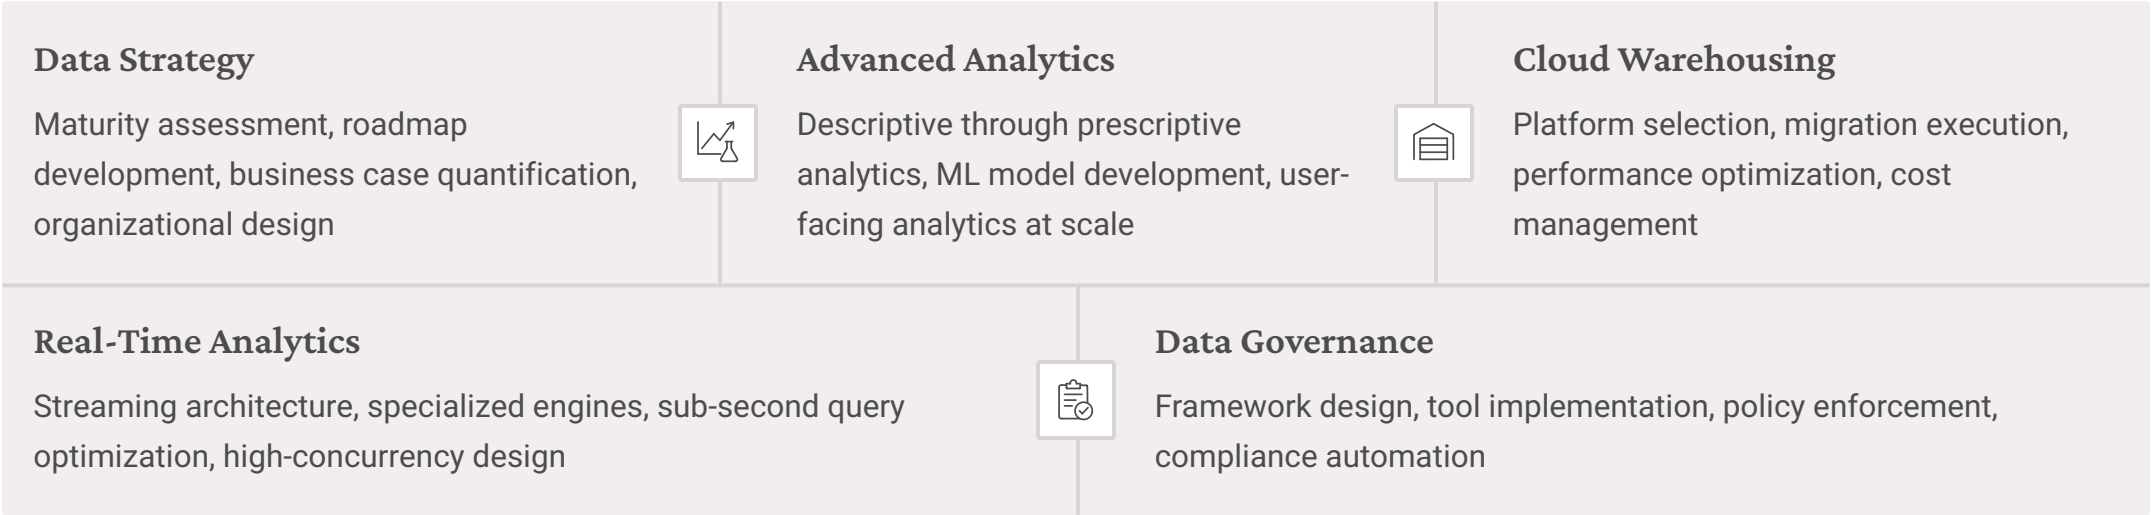
Strategic Impact

Enhanced customer trust through data stewardship, reduced operational risk, and enabled new data-driven products previously blocked by governance gaps.

Partnering for Data Excellence

End-to-End Capabilities Delivering Competitive Advantage

MinervaDB provides comprehensive consulting services across the full spectrum of enterprise data capabilities, from strategy through ongoing operations:



Our platform expertise—spanning Snowflake, BigQuery, Redshift, Databricks, ClickHouse, Druid, and emerging technologies—combined with our vendor-neutral approach, positions MinervaDB as the partner for organizations seeking data excellence as a competitive advantage.

We transform how organizations leverage data to drive growth, efficiency, and innovation.

Transform Your Data Infrastructure

Next Steps: Multiple Ways to Engage

Ready to modernize your data infrastructure? We offer flexible engagement models to match your needs and timeline:

Engagement Options:

- Strategic Assessment Workshop (2-3 days)
Rapid evaluation of current capabilities with prioritized recommendations and ROI projections
- Proof-of-Concept Project (4-6 weeks)
Validate platform selection and performance with real workloads before full commitment
- Full Implementation Partnership (4-6 months)
End-to-end migration and optimization with knowledge transfer to your team
- Advisory Retainer
Ongoing strategic guidance for organizations with internal implementation teams
- Optimization Review
Performance and cost analysis of existing cloud data platforms

All engagements begin with a complimentary consultation to understand your requirements and determine the best path forward.

Contact MinervaDB

Shiv Iyer

Founder and CEO

Email: shiv@minervadb.com

LinkedIn: [linkedin.com/in/thewebscaledba](https://www.linkedin.com/in/thewebscaledba)

Web: <https://minervadb.xyz/data-strategy-and-analytics/>

Schedule a Consultation:

Contact us to discuss your data modernization objectives. We'll provide an honest assessment of whether our expertise aligns with your needs.

What to Expect:

- 30-minute discovery call to understand your situation
- Preliminary recommendations and engagement options
- Transparent pricing and timeline estimates
- No obligation—we only proceed if there's clear mutual fit