

SNOWFLAKE DATA OBSERVABILITY BUYER'S GUIDE

The Definitive Data Platform Leader's Guide
to Selecting the Right Fit for Your Needs



Introduction

The observability market is at an inflection point. According to Gartner's 2024 Market Guide for Data Observability Tools:

- By 2026, 50% of enterprises implementing distributed data architectures will have adopted data observability tools to improve visibility over the state of the data landscape, up from less than 20% in 2024
- Traditional, static, event-based monitoring is insufficient for managing complex data systems effectively
- Enterprises now demand comprehensive solutions that provide real-time visibility, proactive alerting, and actionable recommendations, not just isolated metrics. Yet, the market is currently fragmented, with no standard definition for data observability.
- Most vendors focus on specific areas of observation rather than providing comprehensive solutions

This guide helps you cut through the noise and select a solution that delivers measurable business impact across all critical observability domains.

Understanding Data Observability

Data observability goes beyond traditional monitoring - it's the foundation for proactive data operations.

It enables organizations to:

- Understand the state and health of their data ecosystems
- Continuously monitor data pipelines, quality, and infrastructure health
- Alert and troubleshoot issues proactively
- Analyze and optimize performance and costs
- Prevent data errors and system downtime before they impact business decisions

Gartner identifies five key observation areas



Data Content

Data quality, anomaly detection, schema changes



Data Flow and Pipeline

Pipeline monitoring, execution tracking, broken pipeline detection



Infrastructure and Compute

Resource monitoring, workload analysis, performance tracking



User, Usage, and Utilization

User behavior, access patterns, query analysis



Financial Allocation

Cost monitoring, spend analysis, resource optimization

Your observability solution options

While the need for effective observability solutions for Snowflake is clear, the exact nature and form of observability required can be decidedly opaque. The term "observability" can refer to a wide-ranging landscape of home-grown and off-the-shelf point solutions, each with a different area of focus. This analysis focuses on six of the most common types of tooling for Snowflake data platform observability:

1. Do it Yourself (DIY) Homegrown Solutions
2. Generic FinOps Tools
3. Generic DevOps Observability Tools
4. Snowflake Native Tools
5. Point Solutions for Data Quality and Observability
6. AI-Native Data Observability and FinOps



1

Do-It-Yourself (DIY) Homegrown Solutions

Custom-built monitoring using open-source technologies (like Prometheus, Grafana, ELK stack) combined with Snowflake's native monitoring capabilities and custom SQL queries against ACCOUNT_USAGE schema.

Best Fit

Large engineering teams with specialized resources prioritizing control over cost, pace of innovation, and convenience.

Strengths

Complete customization

- Create custom dashboards specific to Snowflake metrics and your organization's needs

No additional licensing costs

- Use Grafana for visualization and custom scripts using Snowflake's REST APIs

Complete control and flexibility over implementation

- Modify monitoring parameters as business needs change and integrate with internal systems

Tailored to niche business requirements

- Industry-specific data quality monitoring and custom charge-back models

Weaknesses

Requires significant internal engineering resources

- A Dedicated team is needed for development and maintenance, often costing more in FTEs than commercial solutions

Limited scalability

- Custom code and dashboards often break when Snowflake usage scales dramatically

Lack of advanced analytical capabilities

- No built-in machine learning for anomaly detection or sophisticated predictive recommendations

High maintenance drag

- Constant upgrades needed as Snowflake releases new features

Fragmented visibility:

- Siloed tools for cost, performance, and data quality with no unified view

Real-world scenario

A financial services company built its monitoring solution using Grafana dashboards and custom Python scripts. While initially effective, maintenance demands skyrocketed as their Snowflake usage expanded from 10 to 100 clusters, ultimately spending more on engineering resources than a commercial solution would have cost. This mirrors industry trends: most enterprises outgrow DIY as Snowflake's complexity increases, leading to higher costs and slower innovation.

2

Generic FinOps Tools

Cloud cost management platforms are designed for financial governance across multiple cloud providers. Examples include CloudHealth, IBM Apptio Cloudability, Flexera, and CloudZero.

Best Fit

Organizations requiring only high-level cloud cost visibility without granular Snowflake-specific optimization needs.

Strengths

Cost tracking and financial management

- High-level tracking of overall Snowflake spend by account

Multi-cloud support

- Unified view of costs across AWS, Azure, and GCP Snowflake deployments

Mature financial reporting

- Automated monthly billing reports for chargeback and showback

Governance and compliance

- Role-based access control for financial data and audit trails

Integration with financial systems

- Export capabilities for ERP systems and automated invoice reconciliation

Weaknesses

Limited Snowflake-specific insights

- Unable to correlate warehouse failures with increased costs or understand Snowflake-specific concepts like Query Acceleration Service

Lack of deep technical performance analysis

- Can't identify inefficient queries driving up costs or recommend performance tuning opportunities

Minimal data quality monitoring

- No visibility into data corruption issues or pipeline failures that result in costly downstream delays

Superficial observability

- Only tracks account-level costs, not warehouse or query-level efficiency

Reactive rather than proactive

- Identifies cost spikes after they occur, with limited predictive capabilities

Real-world scenario

A retail company implemented CloudHealth to track Snowflake costs but found they still needed separate tools for performance monitoring and troubleshooting. When faced with escalating costs, they could see the increase but couldn't identify which specific workloads or users were causing the problem without additional tools. Ultimately, they had to deploy additional tools for performance monitoring and troubleshooting, highlighting the gap between high-level financial visibility and actionable, Snowflake-specific insights.

3

Generic DevOps Observability Tools

Infrastructure and application monitoring platforms designed for DevOps teams. Examples include DataDog, Dynatrace, New Relic, AppDynamics, Splunk, Elastic Observability.

Best Fit

Teams prioritizing infrastructure optimization and reliability over Snowflake-specific platform optimization and management.

Strengths

Comprehensive infrastructure monitoring

- Detailed metrics on CPU, memory, disk I/O for Snowflake virtual warehouses

Real-time performance tracking

- Live dashboards giving immediate visibility into warehouse performance

Robust alerting capabilities

- Customizable alert thresholds for key metrics and incident management workflows

APM capabilities

- Distributed tracing across microservices and code-level performance insights

Broad technology support

- Monitoring capabilities for databases, containers, and cloud services

Weaknesses

Not specialized for data platforms

- Lack of deep understanding of Snowflake query execution plans and optimization opportunities

Limited data-specific insights

- Cannot correlate data quality issues with performance problems or provide visibility into data lineage

Requires domain expertise

- Users need to know what they are looking for to create the monitoring and alerts required

Expensive for comprehensive coverage

- Agent-based pricing models become costly at scale for data platforms

Complexity in the Snowflake environment configuration

- Difficult to deploy monitoring on serverless and auto-scaling Snowflake resources

Gaps in data platform context

- Limited understanding of Snowflake-specific features like clustering keys, time travel, or zero-copy cloning

Real-world scenario

A healthcare company deployed Dynatrace to monitor their Snowflake environment, but struggled to correlate infrastructure metrics with data pipeline failures. While they could see when a cluster was under stress, they couldn't easily determine which specific data transformations were causing the problems or how to optimize them, ultimately requiring additional tools to close these gaps.

4

Snowflake Native Tools

Built-in tools and features provided by Snowflake for monitoring their platform, including Account Usage views, Query History, Resource Monitors, and the new cost management interface.

Best Fit

Organizations using Snowflake primarily for analytics workloads needing basic monitoring rather than comprehensive optimization and reliability management.

Strengths

Tight platform integration

- Native understanding of Snowflake components with no additional deployment needed

Basic performance and cost metrics

- Warehouse utilization dashboards and query execution histories

No additional tool procurement

- Included with Snowflake licensing and simplified vendor management

Growing capabilities

- Recent enhancements include cost budgets, Query Acceleration Service insights, and improved observability features

Native understanding of Snowflake concepts

- Built-in visibility into features like clustering, time travel, and search optimization

Weaknesses

Limited depth of observability

- Basic metrics without advanced correlation or long-term historical trend analysis

Limited cost and performance optimization

- Few options to identify problems with code, data pipelines, infrastructure over-allocation, and other problems

Lacks advanced predictive capabilities

- No AI-driven recommendations or sophisticated forecasting

Minimal cross-platform insights

- No integrated view of upstream/downstream systems or external integrations

Basic alerting mechanisms

- Limited customization of alerts and minimal integration with external notification systems

Fragmented tooling

- Different interfaces for various monitoring aspects with no unified view across all observation areas

Real-world scenario

A manufacturing company relied on Snowflake's native monitoring tools but found they couldn't effectively analyze long-term trends or get proactive recommendations. When troubleshooting performance issues, they had to piece together information from multiple interfaces manually, and they could not set up sophisticated alerts to identify potential problems.

5

Point Solutions for Data Quality and Observability

Specialized platforms focused primarily on data quality, content monitoring, and pipeline observability. Examples include Monte Carlo, Acceldata, Bigeye, Anomalo, Metaplane, Soda, and Revefi.

Best Fit

Teams primarily needing cross-platform data quality enforcement, often used in combination with more comprehensive solutions.

Strengths

Strong focus on data quality

- Automated detection of schema changes and data anomalies with statistical analysis

Detailed data content monitoring

- Continuous tracking of null values, data freshness, and quality metrics

Rule-based anomaly detection

- Custom business rules for data validation and automated testing

Metadata management

- Tracking of data lineage and impact analysis for downstream dependencies

Specialized data profiling

- Column-level statistics and pattern recognition with trend analysis

Weaknesses

Narrow observability coverage

- Limited or no infrastructure monitoring and minimal cost optimization capabilities

Primarily focused on data content

- Most tools cover only 1-2 of Gartner's five observation areas

Limited infrastructure and cost optimization

- Cannot recommend resource sizing changes or detect idle warehouse costs

Fragmented capabilities

- Need to combine multiple tools for comprehensive coverage with different interfaces for quality versus performance

Varying levels of Snowflake integration

- Some tools have limited support for Snowflake-specific features like Streams, Tasks, or Query Acceleration Service

Real-world scenario

A media company implemented Monte Carlo for data quality monitoring but still struggled with performance and cost issues in their Snowflake environment. While they could detect data quality problems, they lacked visibility into which users were running inefficient queries or how to optimize their cluster configurations, ultimately requiring additional tools to address these gaps.

6

AI-Native Data Observability and FinOps

Purpose-built, AI-driven data observability platforms designed specifically for modern data platforms like Snowflake, delivering unified, actionable insights across all Gartner-defined observability domains. Example: Unravel Data.

Best Fit

Large users of Snowflake with mission-critical reliance requiring comprehensive, full-stack optimization, performance, and reliability, along with AI-native automation capabilities to drive actionability at scale.

Strengths

Comprehensive cross-domain visibility

- Unified coverage across all five Gartner observability areas with correlation of issues across multiple domains

AI-powered insights and automation

- Automatic root cause analysis for failed queries and warehouse issues, ML-based anomaly detection, and proactive notifications

Deep Databricks integration

- Native Snowflake App framework support with applications available on Snowflake Marketplace, enabling installation directly within customer accounts without data movement

Granular cost optimization

- Query-level and warehouse-level cost attribution with specific recommendations and estimated savings

Actionable recommendations

- Concrete, implementable guidance with expected impact, including automated SQL optimization and self-service optimization tools

Advanced Capabilities

SQL code optimization

- Advanced SQL query rewriting and optimization recommendations with before/after performance comparisons

Automated remediation

- Self-healing capabilities that can automatically implement optimization recommendations

Native app security

- Data never leaves the customer's Snowflake environment, addressing security and compliance concerns

Free health check

- Available free health check applications in Snowflake Marketplace that analyze up to 30 days of usage and provide optimization recommendations

Weaknesses

Higher initial investment

- Higher pricing compared to native tools

Organizational adoption challenges

- May necessitate process changes to leverage capabilities fully

Feature richness can be overwhelming

- Requires prioritization of which features to adopt first

Evolving technology area

- Market definition and standards continue to mature

Real-world scenario

A global financial services company implemented Unravel Data in their Snowflake environment and achieved a 40% reduction in compute costs within three months by optimizing queries, cluster configurations, and reducing resource waste. The platform also helped identify data quality issues and link them to specific pipeline failures, avoiding downstream business impacts and significantly reducing mean time to resolution for critical incidents.

Key differentiators to a holistic approach to data observability

1. AI-Powered Insights
2. Comprehensive Coverage
3. Full-Spectrum Observability



1. AI-Powered Insights

1. Automatic root cause analysis for failed queries and warehouses
2. Predictive analysis for resource bottlenecks
3. ML-based anomaly detection for performance and costs
4. Automated recommendations for query optimization
5. Proactive notifications of potential issues before they impact users

2. Comprehensive Coverage

All Five Critical Observation Areas

- Unifies all of Gartner's observability dimensions—data content, data flow, infrastructure, usage, and financial monitoring

End-to-End Visibility

- Tracks data from ingestion through processing to consumption with correlation of user activities and infrastructure performance

Support for Complex Architectures

- Multi-cloud and hybrid deployment support with integration to upstream and downstream systems

3. Full-Spectrum Observability

1. Advanced FinOps Optimization

- Granular cost attribution to users, departments, and applications with idle resource detection and automated scaling recommendations

2. Application Performance Management

- End-to-end visibility into query execution with query-level performance analysis and data skew detection

3. Data Infrastructure Management

- Warehouse configuration optimization and resource utilization analysis with automated scaling recommendations

4. User Management and Optimization

- User behavior analytics and query pattern analysis with self-service optimization recommendations

Unique Value Proposition for Snowflake Environments

Native Snowflake Integration

Modern Snowflake observability solutions leverage the Snowflake Native App Framework to provide:

Security by Design

- Data never leaves the customer's Snowflake environment, eliminating security risks and compliance concerns while leveraging Snowflake's built-in encryption and governance features

Frictionless Deployment

- Install applications directly from Snowflake Marketplace with one-click deployment, similar to downloading an app on a smartphone

Zero Data Movement

- Applications run entirely within Snowflake, eliminating the need to move data for processing and reducing integration complexity

Snowflake Credit Usage

- Pay using existing Snowflake credits with transparent, usage-based pricing models

Free Health Check Offerings

A key differentiator in the Snowflake market is the availability of free optimization health checks through the Snowflake Marketplace:

Unravel Health Check

- Available free on Snowflake Marketplace, analyzes up to 30 days of Snowflake usage across warehouse usage, SQL workloads, and storage optimization

Immediate Value

- Results available within minutes of installation with no security concerns

Risk-Free Assessment

- Organizations can evaluate optimization opportunities before committing to complete solutions

SQL Code Optimization and Actionability

Advanced solutions now include automated SQL optimization capabilities:

Query Rewriting

- Automatic optimization of SQL queries with performance improvements

Before/After Analysis

- Clear performance comparisons showing expected impact

Automated Implementation

- Ability to automatically apply optimization recommendations

Self-Service Tools





- Enable users to optimize their queries with guided recommendations

Selecting the right solution for you

To choose the right solution for your needs, evaluate each solution category against your specific requirements. For example, smaller organizations may be able to get by using existing generic DevOps tools like DataDog, while organizations with very large engineering teams and unique requirements may choose a DIY homegrown option. Most alternatives offer free trials or health check reporting to better understand ROI for your environments. The comparison matrix and evaluation framework can help you determine your best option.



Comparison Matrix

		Data Content	Data Flow & Pipeline	Infrastructure & Compute	User, Usage & Utilization	Financial Allocation
 Limited	DIY Solutions	Basic	Basic	Moderate	Limited	Basic
 Moderate	FinOps Tools	Limited	Limited	Basic	Limited	Strong
 Basic	DevOps Tools	Limited	Moderate	Strong	Moderate	Moderate
 Strong	Snowflake Native	Moderate	Moderate	Moderate	Basic	Basic
	Data Quality Tools	Strong	Moderate	Limited	Limited	Limited
	Unravel Data	Strong	Strong	Strong	Strong	Strong

Recommendation Framework

When assessing data observability solutions for your Snowflake environment, consider these factors:

1

Depth of observability

- How many of the five observation areas does the solution cover?
- What is the granularity of the insights provided?
- Can it correlate issues across different domains?
- Does it support chargeback and showback models?

2

AI and predictive capabilities

- Does the solution offer automated root cause analysis?
- Can it provide proactive recommendations?
- Does it learn from patterns in your environment?
- How accurate are its predictions and anomaly detection?

3

Snowflake-specific features

- Does it understand Snowflake-specific concepts and components?
- How deeply does it integrate with Snowflake?
- Does it support Snowflake Native Apps?
- Can it monitor Query Acceleration Service, clustering, and other advanced features?

Cost optimization features

- How granular is the cost attribution?
- Does it provide actionable recommendations with estimated savings?
- Can it detect idle resources and recommend rightsizing?
- Does it help identify who is accountable for optimization decisions?

Ease of implementation

- What is the deployment time and complexity?
- Does it require data movement or code changes?
- How much configuration is needed?
- Is it available as a Snowflake Native App?

Scalability

- Can it handle enterprise-scale Snowflake deployments?
- How does performance scale with increased data volumes?
- Does pricing scale reasonably with usage?
- Can it support multi-region and multi-account deployments?

Total cost of ownership

- What is the licensing model?
- What internal resources are required for maintenance?
- How does the solution impact Snowflake credit consumption?
- What is the expected ROI?

4

5

6

7

Decision Matrix Based on Primary Needs

Primary Need	Recommended Solution	Runner-Up
Comprehensive Observability	Unravel Data	Combination of Multiple Tools
Data Quality Focus	Data Quality Tools	Unravel Data
Cost Optimization	Unravel Data	FinOps Tools
Performance Tuning	Unravel Data	DevOps Tools with Customization
Basic Monitoring	Snowflake Native Tools	DevOps Tools
Limited Budget	Snowflake Native + Open Source	DIY Solution

Rolling out your solution

Rolling out a data observability solution for Snowflake involves a phased approach—starting with assessment and testing, followed by full deployment and continuous optimization. With clear metrics and real-world validation, organizations can drive measurable ROI through improved performance, reduced costs, and faster issue resolution.



Implementation Roadmap

For organizations considering a data observability solution for Snowflake, consider this phased approach.

Phase 1: 1-2 Weeks

Assessment and Planning

- Identify current pain points, observability gaps, and actionability challenges
- Define key metrics and KPIs for success
- Document current Snowflake architecture and workload patterns
- Establish baseline metrics for performance and cost
- Try free health checks available in Snowflake Marketplace

Phase 3: 4-8 Weeks

Solution Selection and Implementation

- Select preferred solution based on PoC results
- Develop implementation plan and success criteria
- Deploy solution across development, test, and production environments
- Configure alerts, dashboards, and integrations
- Train teams on optimization recommendations and automated features

Phase 2: 4-6 Weeks

Proof of Concept

- Select 2-3 candidate solutions based on evaluation criteria
- Implement each solution in a limited environment
- Leverage Snowflake Native Apps where available for faster deployment
- Test against real-world scenarios and pain points
- Evaluate effectiveness against defined metrics

Phase 4: Ongoing

Optimization and Expansion

- Refine monitoring parameters based on initial results
- Implement recommended optimizations, including SQL code improvements
- Measure and document realized benefits against Snowflake bills
- Expand coverage to additional workloads and environments
- Leverage automation features to scale optimization efforts

ROI Considerations

When calculating potential return on investment for Snowflake observability solutions, consider these factors.

Cost Factors

- Solution licensing costs
- Implementation and professional services
- Internal resources for management
- Training and enablement

Benefit Factors

- Infrastructure cost reduction (typically 20-40%)
- Improved performance leading to faster time-to-insights
- Reduced downtime and mean time to resolution
- Engineering time savings from automated troubleshooting and SQL optimization

Example ROI Calculation

For a mid-sized Snowflake environment (\$500K annual spend)

Average infrastructure cost savings	30% = \$150K/year
Engineering productivity gains	15% of 5 FTEs = \$150K/year
Downtime reduction	50% reduction in 10 annual incidents = \$100K/year
SQL optimization savings	Additional 10% through automated query optimization = \$50K/year
Total annual benefit	\$450K
Total solution cost (Example)	\$150K/year
Net annual benefit	\$300K
RETURN ON INVESTMENT	200%

Conclusion

As Snowflake environments scale, comprehensive data observability is no longer optional—it's a competitive necessity. The right solution delivers measurable cost savings, faster insights, and greater reliability through AI-driven automation and actionability.

While several solution categories offer partial capabilities, AI-native solutions with Snowflake Native App integration stand out for their ability to provide:

- Comprehensive coverage across all five Gartner observation areas
- Deep Snowflake integration without compromising security
- Free health check offerings to demonstrate value before purchase
- Advanced SQL optimization capabilities for true actionability
- Automated remediation to implement recommendations at scale

Organizations should evaluate solutions based on their specific requirements, prioritizing native integration, automation capabilities, and the ability to drive actual implementation of optimization recommendations.

Take the next step: Explore free health check applications available in the Snowflake Marketplace to assess your optimization opportunities risk-free, then implement a comprehensive solution that transforms insights into automated actions for sustained Snowflake optimization.

Additional Resources

- Gartner Market Guide for Data Observability Tools (2024)
- Snowflake Marketplace Native App listings
- Snowflake's best practices for monitoring and optimization
- Industry benchmarks for Snowflake performance and cost optimization

Disclaimer

This buyer's guide is based on Gartner's 2024 Market Guide for Data Observability Tools, industry research, and best practices. Organizations should conduct their own evaluation based on their specific requirements and environments.

