



MONTE  
CARLO

# Data Observability Evaluation Guide

Learn the must-have features for every Data Observability solution, what the analysts have to say, plus an RFP template.



# Table of Contents

## I. Introduction

## II. Key Components of Data Observability

- What's data observability again?
- End-to-end coverage
- Incident management
- Integrated data lineage
- Comprehensive root cause analysis
- Quick time-to-value
- Enterprise readiness
- AI observability



## III. Analyst Perspective

- Gartner
- GigaOm
- G2 Crowd
- Ventana

## IV. Evaluating Alternatives

- RFP Template

## V. The Future of Data Observability



# I. Introduction

Data observability has been one of the hottest emerging data engineering technologies of the last several years. And with data quality and reliability becoming a central topic in the data product and AI conversation, that trajectory isn't changing any time soon.

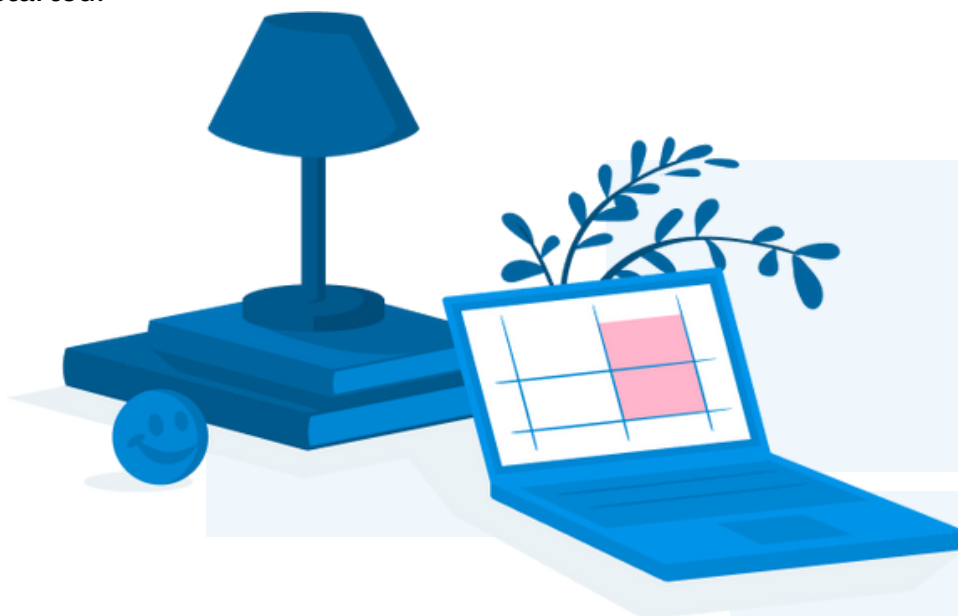
Benefits of data observability include:

- Increasing data trust and adoption
- Mitigating the risks of bad data
- Boosting revenue and data product ROI
- And reducing time and resource costs associated with data quality

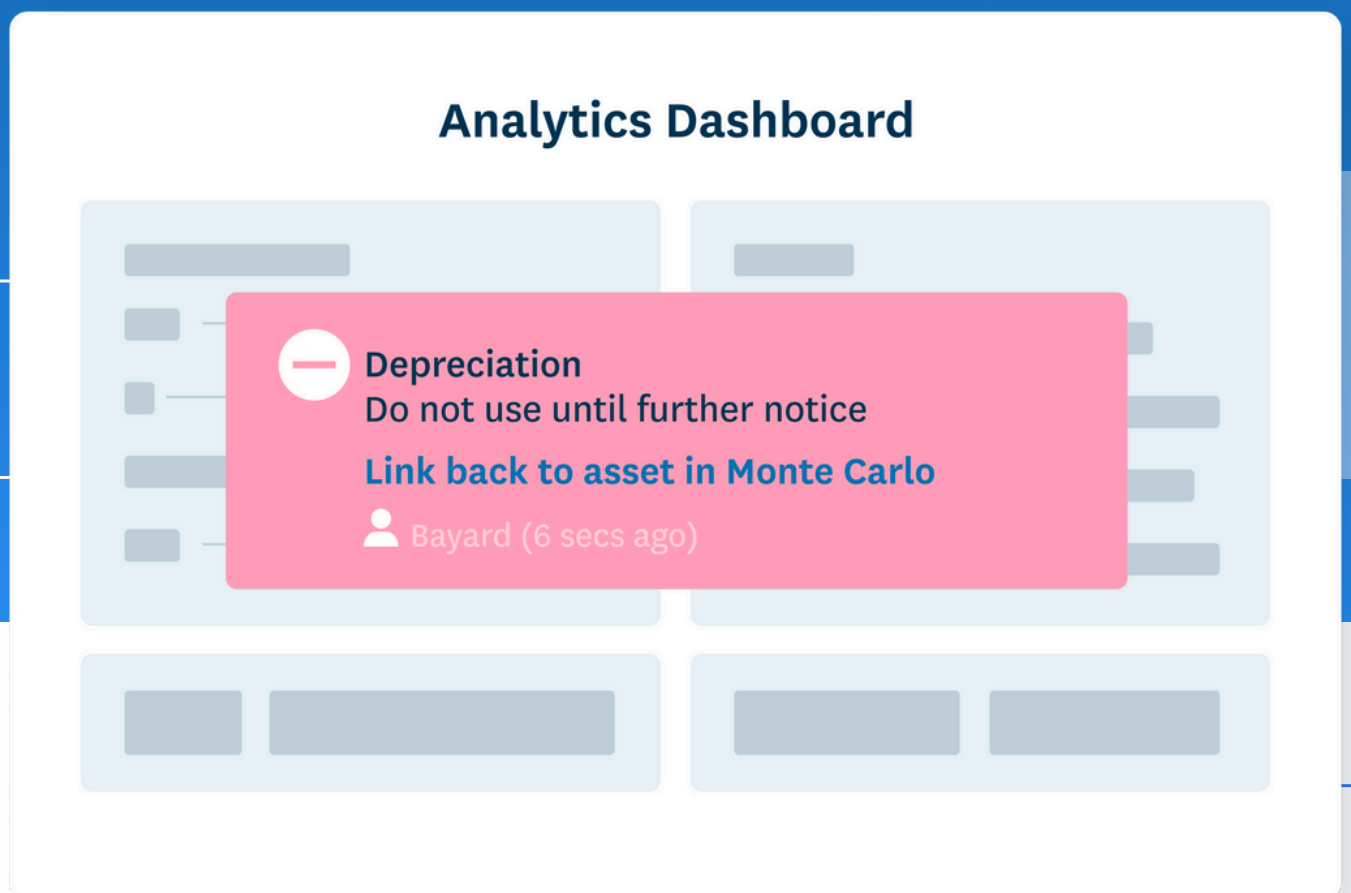
And as the category grows, so too do the alternatives. Following Monte Carlo's creation of the data observability category in 2019, a myriad of alternative data observability tools have entered the market at various levels of maturity.

So, with so much choice, how does a data leader decide? In this post, we'll share the must-have features and benefits based on both analyst perspective and our own experience building the category, as well as a sample RFP template you can use to support your own evaluation process.

Let's get started!



# II. The Key Components of Data Observability



# What's data observability again?

Data observability refers to an organization's comprehensive understanding of the health and performance of the data within their systems.

The purpose of data observability tools is to reduce data downtime by automating or accelerating the detection, management, and resolution of data quality issues. The core features of data observability tools were originally defined across five pillars that included four types of machine learning anomaly detection monitors and data lineage:

- Freshness– Did the data arrive when it was expected?
- Volume– Did we receive too many or too few rows?
- Schema– Did the structure of the data change in a way that will break data assets downstream?
- Quality– Are the values of the data itself within a normal range? Has there been a spike in NULLs or a drop in percentage of unique values?
- Lineage– How does the data flow through your tables and systems across your modern data stack? This is critical to both prioritizing incident response and finding the root cause.

Whatever else a tool may have, if it doesn't cover these five pillars, it's not data observability.

Having said that, while the five pillars are essential to the data observability category, they also aren't the only dimensions we can — or should — use to evaluate these tools.

Let's take a look at each of the critical dimensions impact the efficacy, usability, and longevity of a data observability tool.

# Data Observability Evaluation Criteria

At its core, a data observability solution is a programmatic plug-and-play system that offers data health visibility and management through a mix of automated and customizable data quality tools like automated metadata monitors, impact analysis, data lineage, and data health insights—all to help data engineers proactively detect, resolve, and prevent data incidents faster.

Comprehensive data observability tools will integrate across your entire data stack and provide coverage for issues like data freshness, volume anomalies, and schema changes from CI/CD in addition to monitoring your data directly on your most critical tables for things like NULL rates, duplicates, and values outside a normal distribution.

At Monte Carlo, our customers' needs are never far from our mind when we think about developing data observability as a category and our own feature roadmap. Considering the needs of our customers and the data industry at large, here are the key evaluation criteria we feel are the most critical to a comprehensive and future-proof data observability solution:

- Enterprise readiness
- End-to-end coverage
- Seamless incident management
- Integrated data lineage
- Comprehensive root cause analysis
- Quick time-to-value
- AI observability

So, now that we've got those points in mind, let's dive into each of these criteria in a bit more detail.

# Enterprise readiness

The world of data is always evolving. That's why you need a data observability provider that can serve as a strategic advisor. Any vendors can promise the world, but what can they actually deliver? Will a team of 12 people in a garage still be around in a year to observe your data? These are important questions to answer through customer reference calls to understand a solution's overall maturity.

Some key areas to evaluate for enterprise readiness include:

- Security– Do they have SOC II certification? Robust role based access controls?
- Architecture– Do they have multiple deployment options for the level of control over the connection? How does it impact data warehouse/lakehouse performance?
- Usability– Is an alert just pretty or will it actually save you time—like bulk update incidents or being able to deploy monitors-as-code. Usability can also be subjective in a POC, so it's important to balance this with perspective from actual users.
- Scalability– What are their largest deployments? Has this organization proven its ability to grow alongside its customers? Other key features here include things like ability to support domains, reporting, change logging, and more.
- Support– Data observability isn't just a technology, it's an operational process. The maturity of the vendor's customer success organization can impact your own level of success as can support SLAs (if the vendor doesn't have support SLAs, that's a red flag).
- Innovation history and roadmap– The data world changes rapidly and as we enter the AI era, you need a partner that has a history of first-to-market innovation. Fast followers are often anything but, with comparative features shipped 6 months to a year later. (That's 25 in chief data officer years.) Cloud-native solutions often have an advantage here.

## End-to-end coverage

The true power of data observability lies in its ability to integrate across modern data platform layers to create end-to-end visibility into your critical pipelines.

For years, data testing—whether it was hardcoded, dbt tests, or some other type of unit test—was the primary mechanism to catch bad data.

While still relevant in the right context, the problem with data testing as a complete practice is that you couldn't possibly write a test for every single way your data could break. No matter how well you know your pipelines, unknown unknowns will still be a fact of life. And even if you could identify every potential break (which you can't), you certainly wouldn't be able to scale your testing to account for each one as your environment grew. That leaves a lot of cracks in your pipelines to fill.

Data observability tools should offer both broad automated metadata monitoring across all the tables once they have been added to your selected schemas, as well as deep monitoring for issues inherent in the data itself.

A strong data observability tool will also integrate widely across your data platform, from ingestion to BI and consumption, and enable quick time-to-value through simple plug and play integrations.

Be sure to verify that your chosen solution offers tooling integrations for each of the layers you'll need to monitor in order to validate the quality of your data products, as well as integrations into existing workflows with tools like Slack, Microsoft Teams, Jira, and GitHub.



# Incident management

Most data teams we talk to initially have a detection focused mindset as it relates to data quality, likely formed from their experience with data testing.

The beauty of data observability is that not only can you catch more meaningful incidents, but the best solutions will also include features that improve and accelerate your ability to manage incidents. Bad data is inevitable and having tools to mitigate its impact provides tremendous value. There are a few areas to evaluate when it comes to incident management:

- **Impact analysis** — How do you know if an incident is critical and requires prioritizing? Easy—you look at the impact. Data observability tools that provide automated column-level lineage out-of-the-box will also sometimes provide an impact radius dashboard to illustrate how far a quality issue has extended from its root. This can help data engineers understand at a glance how many teams or products have been impacted by a particular issue and who needs to be kept informed as it moves through triage and resolution.
- **Internal team collaboration** — Once an alert has triggered there needs to be a process for assigning and potentially transferring ownership surrounding the incident. This may involve integrating with external ticket management solutions like JIRA or ServiceNow, or some teams may choose to manage the incident lifecycle within the data observability tool itself. Either way, it's helpful to have the flexibility to do both.
- **Proactive communication with data consumers** — When consumers use bad data to make decisions, you lose data trust. Data observability solutions should have means for proactively communicating with data consumers the current health of particular datasets or data products.


# Integrated data lineage

Lineage is a dependency map that allows you to visualize the flow of data through your pipelines and simplify root cause analysis and remediation.

While a variety of tools like dbt will provide lineage mapping at the table level, very few extend that lineage into the columns of a table or show how that data flows across all of your systems. Sometimes called “field-level lineage,” column-level lineage maps the dependencies between data sets and tables across data products to understand visually how data moves through your pipelines.

It’s also important that your data lineage and data incident detection features work as an integrated solution within the same platform. A key reason for this is that lineage grouped alerting not only reduces alert fatigue, but helps tell a more cohesive story when an event impacts multiple tables.

Rather than getting 12 jumbled chapters that may be part of one or two stories, you are getting an alert with the full story and table of contents.



# Root-cause analysis

What is your standard root cause analysis process? Does it feel disjointed hopping across multiple tools? How long does it take to resolve an issue?

Data can go bad in a lot of ways. A comprehensive data observability tool should help you identify if the root cause is an issue with the [data](#), [system](#), or [code](#).

For example, the data can be bad from the source. If an application went buggy and you started seeing an abnormally low sales price from orders in New York, that would be considered a data issue.

Alternatively, a data environment is made up of a panoply of irreducibly complex systems that all need to work in tandem to deliver valuable data products for your downstream consumers. Sometimes the issue is hidden within this web of dependencies. If you had an Airflow job that caused your data to fail, the real culprit wouldn't be the data but a system issue.

Or if a bad dbt model or data warehouse query change ultimately broke the data product downstream, that would be considered a code issue.

A thorough data observability tool would be able to accurately identify these issues and provide the proper context to help your team remediate each at its source.

# Quick Time To Value

Data observability is intended to reduce work—not to add more.

If a data observability tool is providing the right integrations and automated monitors for your environment out-of-the-box, it will be quick to implement and deliver near immediate time-to-value for data teams.

A data observability solution that requires more than an hour to get set up and more than a couple of days to start delivering value is unlikely to deliver the data quality efficiencies that a growing data organization would require to scale data quality long-term.

## AI observability

Building differentiated, useful generative AI applications requires first party data. That means data engineers and high quality data are integral to the solution.

Most data observability solutions today will monitor the data pipelines powering RAG or fine tuning use cases—they are essentially the same as data pipelines powering other data products such as dashboards, ML applications, or customer facing data.

However, the generative AI ecosystem is evolving rapidly and your data observability vendor needs to be not just monitoring this evolution but helping to lead the charge. That means features like observability for vector databases, streaming data sources, and ensuring pipelines are as performant as possible.

# III. Analyst Perspective



**Leader**

**WINTER**

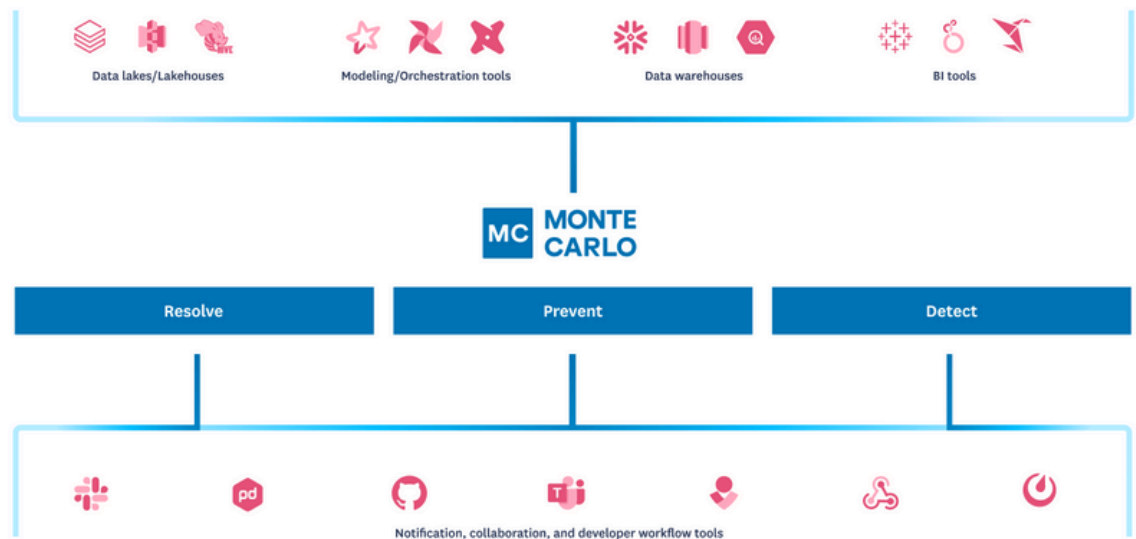
**2024**

# Gartner

While Gartner hasn't produced a data observability magic quadrant or report ranking data observability vendors, they have named it one of the hottest emerging technologies and placed it on the 2023 Data Management Hype Cycle.

They say data and analytics leaders should, "Explore the data observability tools available in the market by investigating their features, upfront setup, deployment models and possible constraints. Also consider how it fits to overall data ecosystems and how it interoperates with the existing tools."

We anticipate Gartner will continue to evolve and add to their guidance on data observability tools this year.



# Ventana

The Ventana Research Buyers Guide does a good job capturing the essence of these tools saying, “data observability tools monitor not just the data in an individual environment for a specific purpose at a given point in time, but also the associated upstream and downstream data pipelines.”

They also used standard dimensions of SaaS platforms in how they ranked vendors:

- Adaptability
- Capability
- Manageability
- Reliability
- Usability
- Customer Experience
- TCO/ROI
- Validation

But, product capability is the highest weighted at 25% of the evaluation. Here, Ventana really hit the nail on the head saying that the best data observability solutions go beyond detection to focus on resolution, prevention and other workflows:

“The research largely focuses on how vendors apply data observability and the specific processes where some specialize, such as the detection of data reliability issues, compared to resolution and prevention. Vendors that have more breadth and depth and support the entire set of needs fared better than others. Vendors who specialize in the detection of data reliability issues did not perform as well as the others.”

# G2 Crowd

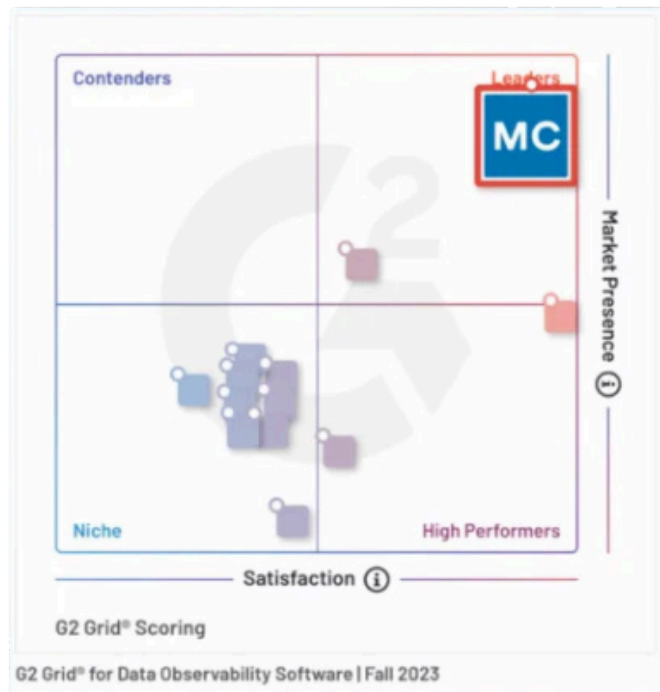
G2 was one of the earliest non-vendor resources to put together a credible list of data observability vendors and a definition for the category. They say:

To qualify for inclusion in the G2 Crowd data observability category, a product must:

- Proactively monitor, alert, track, log, compare, and analyze data for any errors or issues across the entire data stack
- Monitor data at rest and data in motion, and does not require data extraction from current storage location
- Connect to an existing stack without any need to write code or modify data pipelines

Vendors are evaluated by verified users of the product across a list of organizational and product specific capabilities including:

- Quality of support
- Ease of admin
- Ease of use
- Integrations
- Alerting
- Monitoring
- Product direction
- Automation
- Single pane view





# IV. RFP Template



# Data Observability Request For Proposals

Section	Key Capabilities	Criteria
Company Details	Vendor experience	What is your experience in the industry? What references and case studies can you provide from similar projects?
	Innovation and roadmap	What is your vision for the future of data observability? What are your planned enhancements and new features?
Security	Deployment options	Is your product delivered as a Software-as-a-service (SaaS) offering? Can you deploy the agent in our AWS/GCP/Azure instance if desired?
	System Integration and data handling	Can you provide proof of SOC2 Type II certification? What options do you provide for authentication? Do you have role based access controls? Do you provide an API for retrieving security events to import into a SIEM? Are third-party penetration tests available? What data is exported from our environment? Is it encrypted?
Configuration and Management	APIs	What functionality is available via API? Is the API REST?
	CLI & SDK	Is a command line interface available to simplify API interactions? Are any SDKs available, e.g. for use in Python scripts or Data Science Notebooks?
	Airflow Operator	Can monitoring and alerting be configured from within your Airflow jobs, without breaking a workflow?
	Monitors and notifications as code	Can monitoring and alert routing be defined in a file placed under version control? Can multiple monitors be created at once using abstractions, or must monitoring be instantiated for each table and field one by one?
	Performance	How will your product impact our data warehouse/lake/lakehouse performance and compute costs?

## RFP Continued

Section	Key Capabilities	Criteria
Integrations	Warehouse /Lake /Lakehouse	What is your experience in the industry?What references and case studies can you provide from similar projects?
	Other databases	What is your vision for the future of data observability?What are your planned enhancements and new features?
	BI tools	See above.
	Integration, transformatio, orchestration	See above.
	Collaboration tools	See above.
	Data catalogs	See above.
	Streaming data sources	See above.
	Code repositories	See above.
Support & Use	Product support	Do you provide web-based self-support resources? Do you offer 24/7 support (24 hour SLA) What kind of support is provided during major releases? Are your releases backwards compatible? Do you charge additional fees for providing product support?
	Usability and onboarding	What training, onboarding, and ongoing support is available? How are deployment and product best practices shared amongst users? How is the time required to execute tasks minimized by the UI and key workflows? Is administration and management of the platform and its capabilities low-code or no-code?
Pricing Structure	Licensing model	What is the basis of licenses for the product? Is on-demand/usage based pricing available?

## RFP Continued

Section	Key Capabilities	Criteria
Monitoring	<p style="text-align: center;"><b>Pipeline monitoring</b></p>	<p>What kinds of intelligent features are implemented that allow for finding anomalies without the need for manual input? To what extent is anomaly detection manually configurable or tunable? To what extent is tuning required? Does the product detect data freshness and row count issues automatically for all tables in the data warehouse or a selected dataset? Does the product automatically provide schema change detection for selected datasets?</p>
	<p style="text-align: center;"><b>Data quality monitoring</b></p>	<p>Does the product provide monitoring for percents and/or counts of null or missing values? Does the product provide monitoring for uniqueness, duplicates? Does the product provide monitoring for min, max, average, stddev, variance, skew, percentiles? Does the tool offer dimension tracking capability?</p>
	<p style="text-align: center;"><b>Data consistency</b></p>	<p>Does the product support custom business logic that involves more than one table? More than one database?</p>
	<p style="text-align: center;"><b>Custom monitoring</b></p>	<p>In addition to automated tests and checks, what ability does the user have to create custom checks and tests? Can these tests be scheduled and/or ran manually? Are sample SQL templates provided within the tool for defining a custom monitor? Can custom SQL rules be generated with AI from within the platform?</p>
	<p style="text-align: center;"><b>Preventative actions</b></p>	<p>Does the tool have feature of circuit breaker which can stop the pipelines when data does not meet a set of quality or integrity thresholds? Can the tool show the impact of a pull request or dbt model change on downstream assets such as specific dashboards?</p>
	<p style="text-align: center;"><b>Data profiling</b></p>	<p>How does the service profile data shape statistically? (what stats are provided/considered for various data types)</p>
<p style="text-align: center;"><b>Query usage and performance monitoring</b></p>	<p style="text-align: center;"><b>Cost and runtime optimization</b></p>	<p>How does the service monitor query usage and efficiency for the objects under observation? What sort of feedback or recommendations does the tool provide? How is the warehouse monitored for both usage and storage efficiency? (particularly for BI tools?)</p>

## RFP Continued

Section	Key Capabilities	Criteria
<b>Incident Management</b>	<b>Alerting and ownership</b>	<p>What alert channels are supported? How are alerts routed to specific channels and audiences?</p> <p>What kind of guidance or initial information do system alerts provide?</p> <p>How is ownership of incidents and data assets tracked within the platform?</p>
	<b>Lineage and impact analysis</b>	<p>Can the product automatically infer the priority of a detected issue by factors such as table popularity and consumption at the repository and BI layers?</p> <p>Does the product use lineage to navigate users to likely root cause tables in data pipelines with cascading issues?</p> <p>Can the product auto-generate table and column-level lineage within the warehouse/lake/lakehouse?</p>
	<b>Asset metadata /self-service data discovery</b>	<p>How accessible is the history of data incidents or failed checks on an object? Can those incidents be annotated with custom notes?</p> <p>Is key metadata from across the modern data stack, including object tags, surfaced in one pane of glass for each table or data asset? Does it include asset type, schema, table ID, users, usage, logs, and other relevant documentation?</p>
	<b>Consumer notification channels</b>	<p>What service/ticket management/workflow integrations exist? Will they enhance the current levels of collaboration across the data team?</p> <p>How can data consumers be proactively alerted to incidents or the current health status of a dataset?</p>
	<b>Incident tracking</b>	<p>Does the product document detected issues, alert feedback, users involved, and resolution notes as reference for future issues?</p>
	<b>SLA support</b>	<p>Can SLAs or other alert grouping methods be defined? Can alerts be directed to different teams? How flexible is the grouping system?</p>

## RFP Continued

Section	Key Capabilities	Criteria
Resolution/ Root Cause Analysis	Insights	<p>Does the tool detect insights to facilitate in the discovery of the root cause of a particular data incident?</p> <p>Does the product use lineage to help users understand how downstream tables and reports will be impacted by an issue or improvement?</p>
	Data-level RCA	<p>Can the product automatically surface row-level previews of data that likely contributed to a detected anomaly for use in debugging?</p> <p>Does the tool provide segmentation analysis capabilities to further hone in on specific segments where the data is anomalous?</p>
	System-level RCA	<p>Does the product integrate broadly across our data systems to surface alerts behind a single pane of glass?</p>
	Code-level RCA	<p>Does the product correlate changes to query code (no code transformation tools) to specific incidents? Does the tool monitor query performance and execution?</p>
Reporting	Customization	<p>Does the product offer users the ability to create custom dashboards?</p>
	Data health trends reporting	<p>Does the product offer pre-built dashboards showing data health metrics over time at the organizational, domain, and data product level?</p>
	Data health status	<p>Does the product offer pre-built dashboards showing the current health of tables or data products?</p>
	Data operational response reporting	<p>Does the product offer pre-built dashboards showing the operational metrics (time to response, time to fixed) of different teams?</p>

# V. The Future of Data Observability



# What's next for data observability

There's one critical feature that we didn't mention earlier, that plays a huge role in the long-term viability of a data observability solution.

And that's category leadership.

Like any piece of enterprise software, you aren't just making a decision for the here and now—you're making a bet on the future as well. When you choose a data observability solution, you're making a statement about the vision of that company and how closely it aligns to your own long-term goals. "Will this partner make the right decisions to continue to provide my organization with adequate data quality coverage in the future?"

Particularly as AI proliferates, having a solution that will innovate when and how you need it is equally as important as what that platform offers today.

Not only has Monte Carlo been named a proven category leader by the likes of G2, Gartner, Ventana, and the industry at large; but with a commitment to support vector databases for RAG and help organizations across industries power the future of market-ready enterprise AI, Monte Carlo has become the de facto leader for AI reliability as well.

There's no question that AI is a data product. And with a mission to power data quality for your most critical data products, Monte Carlo is committed to helping you deliver the most reliable and valuable AI products for your stakeholders.



# Just starting your data observability journey?

Contact our team today to find out how Monte Carlo can help your team save time, reduce costs, and maximize your data resources with our category-creating Data Observability solution.

Check out more helpful resources on data and AI trends and best practices, including:

- [Data Downtime Blog](#): Get fresh tips, how-tos, and expert advice on all things data.
- [O'Reilly Data Quality Framework](#): The first several chapters of this practitioner's guide to building more trustworthy pipelines are free to access.
- [Data Observability Product Tour](#): Check out this video tour showing just how a data observability platform works.
- [Data Quality Value Calculator](#): Enter in a few specifics about your data environment and see how much you can save with data observability.

