ACTIAN
a division of HCLSoftware

CHECKLIST

Evaluate What Matters

What to Look for In Data Observability Solutions

Choose a solution that doesn't just detect data issues.
It prevents bad decisions, protects SLAs, and makes AI safer by design.

Data observability is the difference between confident analytics and silent failure, and between AI that delivers reliable outcomes and AI that quietly amplifies risk. For businesses, when dashboards drive decisions and AI agents take autonomous action, "good enough" monitoring isn't enough.

According to ISG Research, through 2027, more than two-thirds of enterprises will invest in initiatives to improve trust in data by adopting data observability tools to support the detection, resolution, and prevention of data reliability issues.

Yet with so many options on the market, picking the right solution and validating vendor claims can be difficult. One thing is clear: modern observability must do more than flag problems. It should show business impact, assign ownership, enforce go/no-go gates, and trigger remediation.

From Signal to Action

Turning Data Issues Into Rapid Resolution

Modern data observability isn't a passive monitoring layer. It's an operational control system for data. It should continuously validate whether data is fit for purpose, translate anomalies into business impact, and route issues to the right owners with clear next steps. It should also support automated guardrails, so unreliable data can't quietly flow into dashboards, downstream pipelines, or AI workflows without a clear signal to pause, investigate, and remediate.

That's why this checklist focuses on capabilities that separate "issue detection" from decision protection. It explains impact-aware alerts, accountable ownership, enforceable go/no-go gates, and remediation triggers that reduce incidents, not just report them.

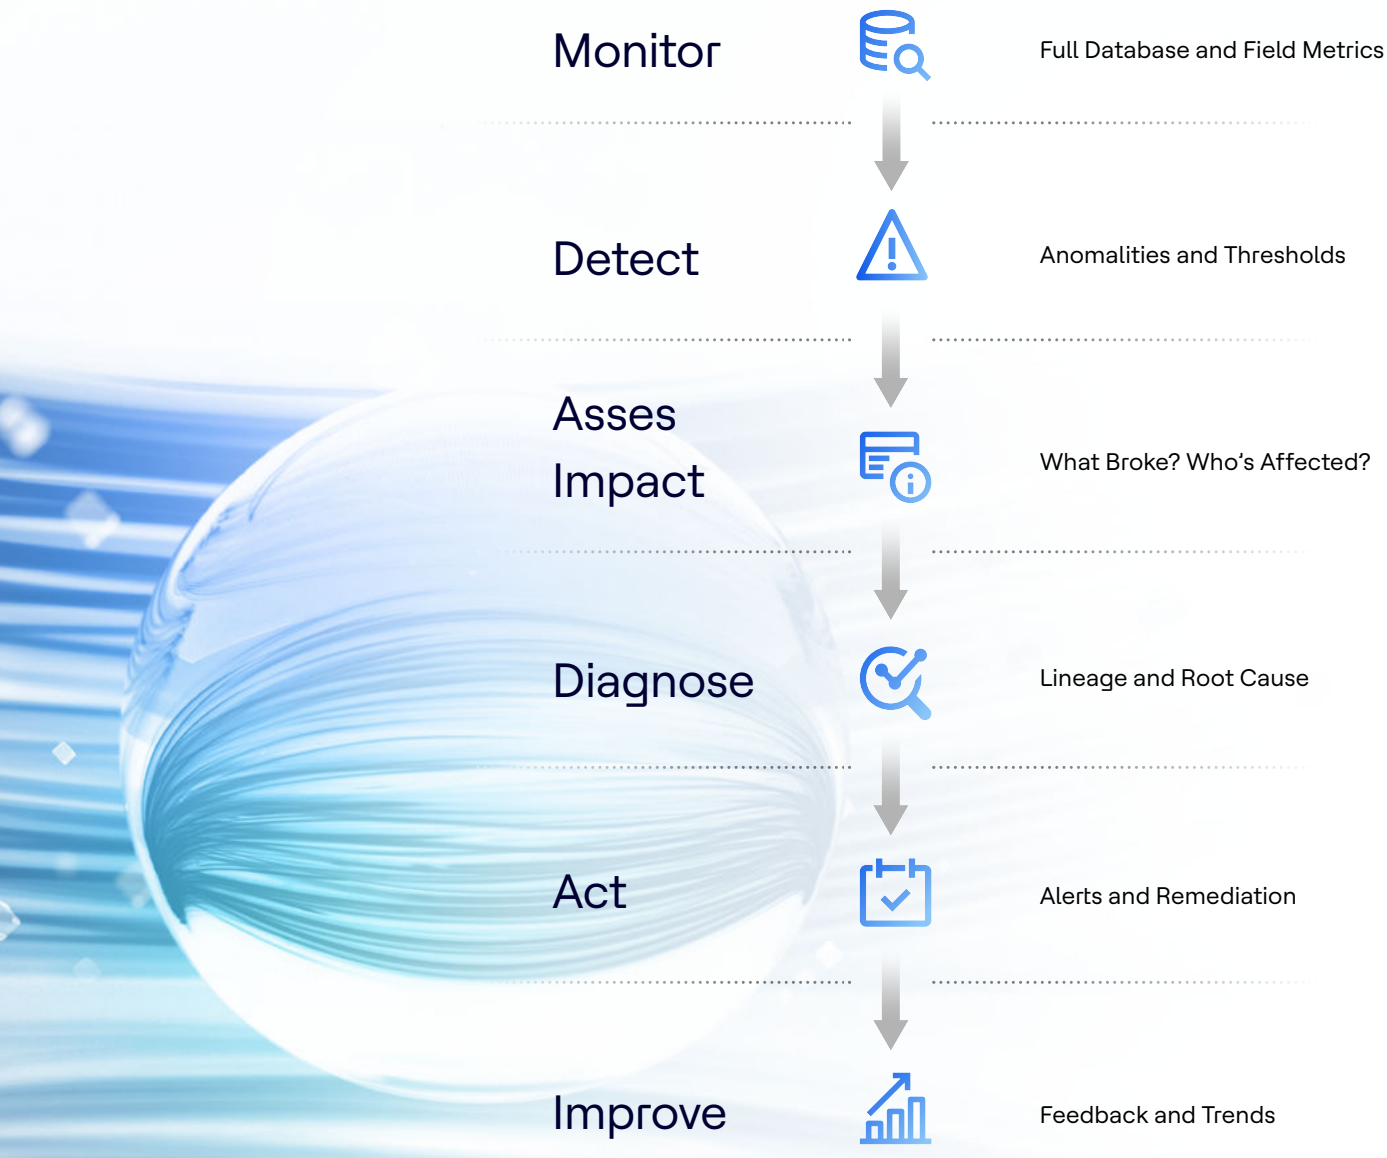

ACTIAN™
a division of HCLSoftware

7 Factors to Consider When Evaluating Solutions

These seven factors separate "yet another dashboard" from a platform that prevents bad data from becoming bad decisions. Use them as your go/no-go thresholds to identify a data observability solution that ensures reliable data at enterprise scale.

#1: Ensure the Right Fit and Coverage

Your observability solution should match your actual ecosystem, not force you to rebuild it to match the vendor's version of an ideal architecture.

Confirm the solution will give you broad coverage without blind spots by asking:

- [ ] Does the solution support our current stack and what's coming next, such as a lakehouse, streaming data, or data intelligence platform?

Action: Inventory your current systems and pipelines, as well as what you'll be adding in the next 18 months. Require proof of support via a demo or pilot in at least two critical environments.

- [ ] Can it monitor data as-is without forcing transformations?

Action: Validate "as-is" data monitoring on existing pipelines. Confirm metrics such as quality, freshness, volume, and schema without creating duplicate copies or re-platforming.

- [ ] Does it support the formats we rely on such as CSV, JSON, or Parquet?

Action: List top formats by volume and criticality, then test monitoring on each. Avoid tools that encourage teams to work around standard processes in order to use preferred formats.

- [ ] Can it provide unified coverage across hybrid and multi-cloud environments?

Action: Map end-to-end data flows across cloud and on-premises systems. Verify consistent alerting, ownership, and reporting.

Business benefit: You establish a single observability layer across all platforms and teams, reducing tool sprawl and accelerating time-to-value, without reshaping your ecosystem to gain visibility into data health.

#2: Trust Quality Metrics that Are Decision-Ready

Observability should make trust actionable by explaining what's wrong, what's impacted, and whether data is safe to proceed for a specific use case.

Start by asking:

- [ ] Can the solution produce trust metrics for validity, freshness, and contract compliance that teams can act on immediately?

Action: Define your core trust metrics and thresholds. Require alerts to include the reason, owner, downstream impact, and next-best action for an issue, not just a status.

- [ ] Can it return explicit go/no-go outcomes for automated workflows?

Action: Implement quality "trust gates" for critical workflows so stale, broken, or non-compliant data cannot move forward silently.

- [ ] Can agentic workflows verify trust before taking action?

Action: Add a trust-before-action checkpoint for AI assistants and automation. Test that the observability solution can block, escalate, or route issues for remediation.

Business benefit: Unreliable data gets stopped early, before it spreads into dashboards, apps, automated workflows, and use cases like AI.

#3: Enable Data Quality Reporting and KPIs

Data observability should make quality measurable, shareable, and repeatable. This way, teams can track progress, prove compliance, and align on what "good" actually means.

These questions can help:

- [ ] Does the solution automatically generate data quality indicators and health reports from metrics?

Action: Pick a few high stakes data products, such as a finance close report or a customer-facing metric. Ask the vendor to generate a data health scorecard for each that shows what changed, what's at risk, who owns the data, and whether it's safe to use.

- [ ] How customizable are KPIs to match business definitions?

Action: Define KPIs by domain, like completeness for finance and validity for compliance. Confirm thresholds and KPI scoring are governed, versioned, and auditable.

- [ ] Can the solution report consistently across domains so leadership sees trusted numbers?

Action: Standardize three to five enterprise data health metrics. Ensure teams can drill into data ownership, see the root cause of issues, and take action for mitigation.

Business benefit: Reporting becomes continuously updated and decision-ready, while governed KPIs create durable definitions of quality data across domains, use cases, and shifting priorities.

#4: Enforce Rules, Expectations, and Data Contracts

Data observability should help you formalize what "good" looks like and keep it enforced as data, teams, and pipelines change. The goal is to establish durable, testable guardrails that ensure quality data.

Ask critical questions to validate rules and data contract enforcement:

- [ ] Can we define and manage rules, expectations, and contracts for data integrity and reliability?

Action: Begin with high-impact datasets and define contract terms for freshness, completeness, validity, schema, ownership, and escalation. Require versioning and audit history in the observability solution.

- [ ] Can the solution handle both basic data checks and more advanced business rules, such as required fields plus rules that compare multiple fields or datasets?

Action: Start with a small baseline set of checks on your most important datasets. Add business-critical rules your teams already rely on, like "charge code must match service type" or "customer status must align with account type." Confirm the solution can run both basic and advanced checks consistently across large volumes without slowing pipelines.

- [ ] Can teams manage rules in a user interface with ownership, approvals, and change history, without coding, and monitor data continuously?

Action: Confirm rules can be created through a usable workflow, rather than code only, with approval processes and continuous monitoring.

Business benefit: Data users agree on what "good data" means and the solution enforces it. Data quality standards are built into the monitoring process, so when a team changes a dataset or pipeline, you catch problems early instead of discovering them after the data has been used.

#5: Catch Anomalies and Drift Before They Spread

Detect issues early and make them explainable so teams can take quick action. Ask these questions of the observability solution:

- [ ] Does it detect drift and outliers at the data, metadata, and business-metric levels?

Action: Test detection on known incident patterns. Confirm coverage across pipeline events, dataset changes, and KPIs.

- [ ] Can it separate signals from noise to reduce alert fatigue?

Action: Validate issue prioritization by impact and ownership along with suppression for issues that don't require notification or present a problem for the use case. Confirm you can fine tune metrics without concealing risk.

- [ ] Can it provide field and column-level insights to speed root cause analysis?

Action: Plan a controlled anomaly and verify it pinpoints affected columns, records, and use cases.

Business benefit: Silent failures are caught early and triaged quickly because teams can see what changed, where it changed, and what to investigate first.

#6: Drive Fast Action and Remediation After Issue Detection

Detection becomes valuable when it leads to a resolution and issues are mitigated quickly. Ensure remediation capabilities by asking:

- [ ] Does the solution support workflows across detection, impact, root cause analysis, and action?

Action: Run an end-to-end incident test to ensure the solution provides details that cover an alert, impact, owner, root cause, remediation, and verification. Require clear audit trails for incidents.

- [ ] Can it support lineage for fast tracing at the system and column levels?

Action: Validate lineage depth on critical flows. Confirm tracing to job, dataset, and field levels during incidents.

- [ ] Can we route bad data for remediation or quarantine it to prevent re-contamination?

Action: Define your quarantine controls, such as blocking a job or routing exceptions. Verify the solution can automatically trigger controls using established policies.

Business benefit: Teams spend less time firefighting issues and see fewer repeat incidents, resulting in more reliable analytics and safer AI inputs.

#7: Scale Securely and Drive Solution Adoption

A data observability solution won't deliver its full value if it overwhelms infrastructure or isn't embraced by your data teams.

Here are questions to consider:

- [ ] How does it scale with data volume, velocity, and variety?

Action: Test a real-world scenario. Pick a busy time, like month-end close, and your most important data feeds. Ask the vendor to show how the solution performs and what it costs at that peak, then ask what happens when you double the number of datasets, add a new source, or increase update frequency.

- [ ] Does it rely on sampling? If so, is that acceptable for your quality needs?

Action: Find out if sampling applies and what it may miss. Confirm it still catches the issues you care about or consider a solution that monitors 100% of your data.

Business benefit: Data monitoring and mitigation scale without performance surprises, and adoption drives ROI because data trust isn't limited to a small pool of experts. Strong security and governance also strengthen audit readiness

5 Red Flags During a Data Observability Demo

If you hear these answers during a demo, keep asking questions or keep walking.

#1: We integrate with anything.

Red Flag: No proof of working integrations in your core stack without heavy custom work.

Ask For: A live demonstration of out-of-the-box connectors plus a clear path for adding sources.

#2: Sampling is fine for most organizations.

Red Flag: Sampling is the default and they can't pinpoint exact issues for remediation.

Ask For: Full monitoring where needed, with column or record-level evidence.

#3: You'll just get alerts in email or Slack.

Red Flag: Alerts can't be routed by team, prioritized by impact, or used for automation.

Ask For: Flexible routing plus API and webhook support to trigger tickets, gates, or responses.

#4: The root cause is in the dashboard.

Red Flag: Vague charts with no lineage, field-level diagnostics, evidence, or history.

Ask For: A workflow that covers alert to impact to root cause to fix, with data lineage and diagnostics.

#5: An issue stops all data flows.

Red Flag: No ability to isolate bad data while keeping good data moving.

Ask For: Quarantine and isolation controls to minimize unnecessary downtime.

Turn Data Trust into a Repeatable Advantage

When you evaluate data observability solutions, don't stop at "Can it detect issues?" Ask what actually matters, which is "Can it prevent bad decisions before they happen?"

The strongest solutions produce decision-ready trust signals, enforce data contracts, catch drift early, and take clear, explainable action. This way, reliability becomes repeatable. If a solution can't support go/no-go decisions and remediation at scale, it's not protecting outcomes. It's simply reporting on risk.

With Actian, your teams can pinpoint root causes faster, understand issues in plain business language, and apply tailored validation rules that help prevent repeat failures. You can then move from reacting to incidents to building lasting confidence in your data and AI.

About Actian

Actian empowers enterprises to confidently manage and govern data at scale. Organizations trust Actian data management and data intelligence solutions to streamline complex data environments and accelerate the delivery of AI-ready data. Designed to be flexible, Actian solutions integrate seamlessly and perform reliably across on-premises, cloud and hybrid environments. Learn more about Actian, the data and AI division of HCLSoftware, at actian.com.

Tel +1 512 231 6000 Fax +1.512.231.6010 710 Hesters Crossing Road, Suite 250, Round Rock, TX 78681 **Actian.com**

ACTIAN™
a division of HCLSoftware