REDLEIF

Evaluation Checklist:

Unified Data & AI Platform

A practical guide for evaluating platforms that unify data engineering, analytics, governance, and AI

Pre-RFX Supplement | 2026

✓ How To Use This Checklist

This checklist helps you run structured market research and vendor evaluation without writing a full, formal RFX. It delivers many of the same decision-support inputs procurement teams typically build during RFX intake: clear requirements, evidence requests, evaluation criteria, and a documented decision trail.

If you later choose to run a formal RFX, you can use the outputs from this checklist as a plug-in for your requirements attachment, vendor demo script, and scoring rubric.

Step 1. **Start with outcomes and constraints.** Capture mission outcomes, risk boundaries, and required timelines before comparing features.

Step 2. **Run a structured vendor screen.** Share this checklist in advance and ask vendors to respond in writing (short, direct answers).

Step 3. **Hold at least one focused screening call per vendor (we recommend about 60 minutes).** Use the "Minimum bar" items in each section to eliminate poor fits early. If needed, schedule a second deep-dive session for the short list.

Step 4. **Require proof for the requirements that matter most.** For high-risk or complex requirements, ask for a demo, architecture evidence, or a reference that confirms the claim. Use judgment on lower-impact items.

Step 5. **Score what matters.** Weight governance, security, and implementation realities more heavily than "nice to have" capabilities.

Step 6. **Translate into procurement language.** Use the "Evidence to request" lines to build an RFX addendum or requirements attachment.

What "done" looks like: You can explain, in plain language, why the selected platform is the best value and/or lowest-risk way to deliver trusted enterprise search and answers across your environment, within your timeline and compliance boundaries.

1. Mission Fit and Modernization Approach

Modernization Philosophy

- [ ] Can deliver value without forcing a rip-and-replace of mission systems in year one
- [ ] Supports incremental modernization (prove value on one workload, then expand)
- [ ] Makes it realistic to reduce the cost and risk created by duplicated data, brittle pipelines, and disconnected tools

Workload Alignment

- [ ] Clear support for your priority workloads (batch ETL, streaming, BI, data science, GenAI)
- [ ] Demonstrates how the platform handles both structured and unstructured data (documents, logs, images, etc.)
- [ ] Can map platform capabilities to mission outcomes, not just technical activities

Evidence To Request

- A sample 90-day plan that starts with one workload and shows measurable outcomes
- Two examples of modernization programs that expanded from an initial pilot

Red Flags

- Time and materials-based “big bang” migrations presented as the default path
- The platform story depends on replacing multiple existing systems before value shows up

2. Platform Architecture and Operating Model

Managed Services vs. In-Environment Operations

- [ ] Can explain what is managed by the vendor versus what runs in your environment
- [ ] Supports deployment patterns appropriate to your environment (including restricted networks and compliance zones)
- [ ] Demonstrates how data stays under your control (identity boundaries, storage boundaries, and network boundaries)

Open Architecture and Portability

- [ ] Uses open formats and interfaces to avoid hard lock-in
- [ ] Provides a clear data portability and exit approach (export, migration, or multi-platform coexistence)
- [ ] Can operate across one or more clouds as needed (or explain limitations clearly)

Evidence To Request

- Reference architecture diagram for a regulated environment
- Written explanation of where metadata, logs, governance enforcement, and audit records reside

Red Flags

- Vague answers about where data, metadata, or governance enforcement lives
- No credible exit or portability story

3. Data Ingestion, Processing, and Reliability

Batch and Streaming Fundamentals

- [ ] Supports both batch and streaming pipelines with consistent operational visibility
- [ ] Provides mechanisms for data quality checks and validation before data is used downstream
- [ ] Can demonstrate how pipeline failures are detected, triaged, and remediated

Engineering Reliability and Operational Burden

- [ ] Can show how the platform reduces “3 AM pager duty” dynamics through monitoring and automation
- [ ] Provides repeatable patterns for building production-grade data products, not just one-off pipelines
- [ ] Can demonstrate performance tuning and cost controls for common workloads

Evidence To Request

- Live demo: a broken pipeline scenario with alerts, logs, and remediation steps
- Example runbooks and an SRE-style operating model for data pipelines

Red Flags

- Reliability is treated as an afterthought or “handled by your team”
- Observability exists only at the infrastructure layer, not the data pipeline layer

4. Data Governance, Lineage, and Auditability

Governance is Enforced, Not Documented

- [ ] Governance controls are built into the platform, not dependent on policy binders and manual processes
- [ ] Role-based controls can be enforced across data engineering, analytics, and AI workflows
- [ ] Supports consistent enforcement for sensitive data (row/column-level controls where applicable)

Lineage & Traceability

- [ ] Can trace data from source through transformations to dashboards and AI outputs
- [ ] Supports audit trails for access, changes, and policy enforcement events
- [ ] Can explain how governance scales as data sources, users, and models increase

Evidence To Request

- Demo: lineage graph from a metric or model output back to raw sources
- Sample audit log exports suitable for oversight or audit requests

Red Flags

- Lineage is partial, manual, or only available for certain tools
- Governance controls break when workloads move between analytics and ML

5. Security, Compliance, and Risk Posture

Baseline Security Controls

- [ ] Encryption at rest and in transit, with clear key management options
- [ ] Clear identity and access integration patterns (including how changes propagate)
- [ ] Network controls appropriate for restricted environments (private connectivity options, segmentation)

Compliance Readiness

- [ ] Can describe the compliance approach relevant to your environment (authorizations, audits, and evidence production)
- [ ] Provides clear logging, retention, and eDiscovery/legal hold considerations (where relevant)
- [ ] Can explain how sensitive data is protected in analytics and AI workflows

Evidence To Request

- Security architecture overview and shared responsibility model
- Sample compliance evidence package (controls mapping, audit artifacts)

Red Flags

- “Trust us” security posture with minimal documentation
- Inability to explain how AI features inherit security and governance controls

6. Analytics and Decision Support

BI and SQL Performance

- [ ] Supports SQL analytics at scale with predictable performance and operational controls
- [ ] Can demonstrate concurrency patterns and governance of shared compute
- [ ] Supports semantic consistency so different teams do not generate conflicting truths

Compliance Readiness

- [ ] Can show how analytics outputs connect to real operational decisions (not just dashboards)
- [ ] Enables near-real-time or event-driven analytics where mission needs require it

Evidence To Request

- Demo: governed SQL analytics on representative datasets with a performance discussion
- Examples of how the platform supports consistent definitions and “single source of truth” semantics

Red Flags

- Analytics is treated as separate from governance, engineering, and AI
- Performance claims without workload-specific proof

7. AI and GenAI Readiness

Responsible AI foundation

- [ ] AI capabilities are grounded in governed data, with traceability back to sources
- [ ] Clear controls for model access, data access, and output logging
- [ ] Can explain how the platform supports model lifecycle management and monitoring

From Prototype to Reduction

- [ ] Demonstrates a clear path from experiments and prototypes to production deployments
- [ ] Monitoring exists for model performance, drift, and operational stability
- [ ] Can support GenAI applications in a way that keeps data under agency control

Evidence To Request

- Demo: end-to-end workflow from governed data to model training or GenAI app evaluation and monitoring
- Example policy: what is logged, who can access logs, and how outputs are reviewed

Red Flags

- AI is positioned as “magic” without a concrete governance and monitoring model
- No clear operational story once models move beyond pilots

8. Data Sharing, Federation, and Interoperability

Federation and Cross-Estate Access

- [ ] Can query or analyze data across the estate without creating unnecessary copies
- [ ] Can explain when federation is appropriate versus when consolidation is necessary
- [ ] Supports secure sharing patterns that do not require complicated ETL or replication

Interoperability

- [ ] Integrates cleanly with your current tools and existing investments
- [ ] Provides APIs and extensibility that allow you to standardize without blocking teams

Evidence To Request

- A reference architecture for federated access across two or more data domains
- A security model for sharing datasets, models, or outputs with internal and external partners

Red Flags

- “Sharing” requires full duplication of data or weak governance boundaries
- Vendor cannot explain the tradeoffs and limits of federation in plain language

9. Implementation, Enablement, and Operational Ownership

Time-to-Value

- [ ] Can be operational on a defined scope in weeks, not quarters
- [ ] Provides a phased rollout plan (pilot, expand, standardize)
- [ ] Clearly defines what you need to provide (people, access, environments, data readiness)

Enablement and Sustainability

- [ ] Includes training, documentation, and knowledge transfer
- [ ] Provides clear operational roles: what your team owns versus vendor responsibilities
- [ ] Has a support model and SLAs that match mission criticality

Evidence To Request

- Implementation plan with milestones, acceptance criteria, and deliverables
- Reference calls that include both technical and executive stakeholders

Red Flags

- Implementation depends on hard-to-staff specialist roles with no enablement plan
- Vendor cannot commit to clear milestones and acceptance criteria

10. Commercials, Procurement, and Value Validation

Pricing Clarity and Cost Controls

- [ ] Transparent pricing model with clear cost drivers (compute, users, workloads, features)
- [ ] Ability to forecast and control spend under real usage patterns
- [ ] Contract terms align with budget cycles and procurement realities

Proof of Value

- [ ] Offers a pilot with measurable success criteria
- [ ] Provides a measurement plan for “time-to-insight,” reliability gains, and governance improvements
- [ ] Can quantify reduction of data duplication, integration overhead, and operational burden over time

Evidence To Request

- Pilot plan with success metrics, data requirements, and an executive readout format
- Three references that can discuss outcomes and tradeoffs, not just deployment

Red Flags

- Opaque pricing and no credible forecasting method
- Value case depends on soft benefits without operational metrics

How To Apply This Checklist

Initial Screening — Use Sections 1, 2, 4, and 5 to eliminate poor fits.

Technical Deep Dive — Use Sections 3, 6, 7, and 8 to validate architecture and workload reality.

Pilot Design — Use Sections 9 and 10 to define a pilot with measurable outcomes and an executive readout.

Procurement Translation — Convert checked items and evidence requests into a requirements attachment, evaluation rubric, and demo script.

Final Guidance

The best platforms will welcome scrutiny. If a vendor cannot explain tradeoffs, show traceability, or demonstrate how governance and security are enforced across analytics and AI, that is a signal.

A unified data and AI platform should reduce integration burden, increase auditability, and make it easier to deliver reliable analytics and responsible AI without destabilizing mission systems.


REDLEIF