

From pilot to production:

The 5 requirements for agent-ready data

EXECUTIVE BRIEF

What's changed for CDAOs isn't analytics or AI — it's autonomy. Once systems are allowed to act on their own, trust, context, policy, and predictability become executive level obligations. Most AI programs don't stall because models fail. They stall because leaders can't approve systems they can't explain, control, or defend. When something goes wrong, no one separates analytics from AI. Leaders ask whether the decision was right — and whether you can prove it.

The core insight

Organizations that consistently move from pilots to production don't succeed by innovating faster or chasing more use cases. They succeed by putting **approval requirements** in place that make autonomous systems governable, explainable, and defensible at scale. These aren't maturity stages or architectural preferences. They're **conditions that must be true before agents can be approved to operate.**

The 5 requirements for agent-ready data

01

Trust

Trust is a release gate, not an aspiration. In agentic systems, trust can't be subjective or retroactive. Leading organizations treat it like uptime or security: every analytics or AI use case ships with explicit standards for accuracy, completeness, freshness, and lineage. If those standards aren't met, deployment is blocked. This shifts the CDAO's role from defending outcomes after incidents to approving systems with confidence upfront. Without provable trust, autonomy is a non starter.

02

Unify data and content

Agents need more than facts — they rely on policies, intent, rules, and rationale. The goal isn't to physically unify structured and unstructured systems; no organization has, and it isn't necessary. What matters is governed retrieval at decision time: standardized, permission-aware access where structured data and unstructured content are retrieved together, with recency, access controls, and citations enforced consistently. This enables organizations to answer once, accurately and defensibly, instead of debating conflicting AI outputs after the fact.

03

Context

A shared understanding ensures the same question yields the same answer everywhere. Without shared context, agents can be factually right but organizationally wrong. Actionable context is grounded in governed definitions, metrics, entities, relationships, and owned data products, giving agents stable, trusted interfaces to reason and act. Most agent initiatives fail on meaning, not access. Context prevents conflicting answers from eroding credibility and executive trust.

04

Policy as code

At agent speed, manual governance fails. Policies must execute automatically at runtime, controlling access, purpose, agent use, and auditability. In agentic systems, policy governs decisions as they happen, not after the fact. This requires agents with identity, authorized plans and actions, and fully provable outcomes. Aligning policies to regulatory and industry standards makes governance explainable, portable, and defensible. Policy as code is what allows autonomy to scale within enforceable guardrails.

05

Deterministic outcomes

As agents act, predictability is non negotiable. Autonomous systems must be observable, bounded, and correctable, using proven techniques like risk scoring, anomaly detection, routing, escalation, and verification to monitor and constrain behavior. The goal isn't day one perfection, but behavior predictable enough to approve and governed tightly enough to correct before failures become incidents. If outcomes can't be anticipated, they can't be approved.

The pathway to agent-ready data

Agent ready data is a pathway from experimentation to scaled autonomy, defined by a small set of requirements. These include trust as a release gate, consistent governance of data and content at decision time, standardized agent access, policy enforced as code, predictable and bounded behavior, and alignment with regulatory standards. Crucially, this can be achieved using shared metadata signals across existing systems rather than a single integrated platform.

Near-term actions CDAOs can take

This isn't an enterprise rollout—it's a focused control demonstration in a single domain. Within one quarter, CDAOs can show real progress by taking three steps: Make one data product agent ready by defining trust (accuracy, freshness, lineage) as a release gate, implement one permission aware retrieval pattern so access and recency are enforced consistently, and enable agent access logging and proof so every action is auditable by agent, purpose, and authority.

Deterministic outcomes come later; they become meaningful only after trust, governed access, and policy enforcement exist. Delivering these first steps proves agents can be governed, earning the right to scale.

The CDAO takeaway

Agent ready data isn't about moving faster. It's about approving autonomy with confidence.

CDAOs who scale analytics and AI responsibly are not more experimental — they are more controlled.

By operationalizing trust, unified access, context, and policy early, autonomy becomes a governance decision rather than a leap of faith.

Talk to a Rocket expert
rocketsoftware.com/talk-to-expert

Assess your readiness against the five requirements and begin your path to agent-ready data.



Modernization. **Without Disruption.**[™]

Visit RocketSoftware.com

© Rocket Software, Inc. or its affiliates 2026. All rights reserved. Rocket and the Rocket Software logos are registered trademarks of Rocket Software, Inc. Other product and service names might be trademarks of Rocket Software or its affiliates.

MAR-18051_CDAOGartnerSummit26_Broch_MichaelCurryPresentation_V3