FRONTIER.

FRONTIER / NEW ANALYSIS



Smartsheet’s vsActual Bet: AI Can Accelerate Work—or Industrialize Execution Drift

At ENGAGE 2026 Smartsheet unveiled vsActual and new AI agents, integrations, security controls and consumption pricing. The CIO challenge is to turn the crowded work-management stack into governed, provable execution—not faster status theater.

FRONTIER / Dion Hinchcliffe / November 5, 2026

NEW ANALYSIS

Built to change a decision.

At ENGAGE 2026 Smartsheet unveiled vsActual and new AI agents, integrations, security controls and consumption pricing. The CIO challenge is to turn the crowded work-management stack into governed, provable execution—not faster status theater.

EDITION	**New Analysis**
AUTHOR	**Dion Hinchcliffe**
PUBLISHED	**November 5, 2026**

This PDF is a publication rendering of the canonical Frontier News article. The web article remains the live source of truth; analysis, figures, and source citations travel with this edition.

Smartsheet is making an important strategic move at ENGAGE 2026: from managing work to claiming the intelligence layer that tells enterprises whether work still serves the strategy. Its newly unveiled vsActual product sits above existing systems, while new AI agents, integrations, security capabilities and consumption-based pricing aim to modernize the underlying work-management platform. The ambition deserves attention. So does the distance between a persuasive announcement and a proven operating model.[1][2][3]

The Frontier thesis: the next enterprise work platform will not win by (just) managing tasks and projects better or providing prettier executive status reports. It will win by proving, continuously and under policy, how a funded decision becomes an actual outcome. Without that evidence, AI will merely industrialize execution drift rather than eliminate it.

The underlying problem: AI speed is not organizational progress

Work management has always been deceptively difficult. A project can be green while its value proposition is collapsing. A task can be complete while the dependency it was meant to satisfy has changed. Resource commitments, financial forecasts and delivery reports often live in different systems, with different owners and inconsistent definitions of “done.” Add autonomous workers and the enterprise can accelerate the production of inconsistent answers.

This is not merely a vendor-created problem. In a 2025 study of more than 5,800 people, the Project Management Institute found that only about half of projects delivered value exceeding their effort and expense; 13% failed outright. In a separate executive survey, the planning–execution disconnect was the most frequently named barrier to reinvention. Faster task completion cannot, by itself, repair that disconnect.[4]

Smartsheet has earned the right to address this gap. For years, it has been among the more enterprise-capable collaborative work platforms, with program templates and Control Center, portfolio rollups, SAML SSO, directory integration, granular sharing and audit capabilities. It has occupied a useful middle ground between flexible spreadsheets and rigid departmental project systems. Those assets matter: enterprise AI requires an enforceable operating model, not simply a chatbot pasted onto a board.[5][6]

But history provides no immunity. As work becomes executable through agents, Smartsheet must move quickly from managing human updates to governing machine-initiated work across systems it does not own. Its November 4 announcements show a serious attempt—not proof that the transition is complete.

The new architecture: from sheets to a plan of record

vsActual is the strategic wager. Smartsheet says the standalone, AI-native product can ingest strategy documents, budgets and business cases, construct an approved plan of record, connect initiatives to execution in tools such as Jira, ServiceNow, SAP, Oracle and Workday, and flag budget, schedule and capacity drift. Agents would reconcile data, propose scenarios and explain recommendations; people must approve changes to the plan, with revisions recorded. That human-approval boundary is a substantive design choice, especially when systems are increasingly empowered to act.[1][7]

Underlying the pitch is an Execution Graph connecting tasks to projects, programs, people and strategic goals. Smartsheet invokes 20 years of accumulated workflow intelligence; independent reporting also describes its claims about a large historical data corpus. Such context could become a defensible advantage, but an enormous corpus is not evidence of accurate, current causal relationships in any particular customer environment. CIOs should ask how the graph establishes provenance, confidence, change history and identity across conflicting systems, rather than accept scale as a proxy for truth.[3][8]

CEO Rajeev Singh captured the aspiration: “Every company has a strategy. The ones that win are the ones whose work actually delivers on that strategy.” He is right about the competitive distinction. The unresolved question is whether vsActual will expose uncomfortable truth or merely reconcile systems into another confidently green executive dashboard.[1]

Availability matters. vsActual is in early access with enterprise design partners; Smartsheet plans general availability in early 2027 and has not announced pricing. It is positioned to work without requiring customers to migrate their teams onto Smartsheet. This makes a federated architecture credible as a product direction, not yet a proven capability at heterogeneous enterprise scale.[7]

What IT can actually use—and what it must verify

The less theatrical announcements may have greater immediate operational value.

Agents with inherited access, not a new shadow identity system. Available now, Agentic Workflows can schedule recurring summaries and risk checks. The Creation Agent builds sheets, dashboards, reports and automations; the Data Quality Agent surfaces gaps and conflicts; Smartsheet's broader product recap also lists an Automations Agent. These run through Smart Assist or connected AI clients and, according to the product documentation, inherit users’ existing permissions. Smartsheet says changes to sheets require approval. That is the right default, but administrators still need to test delegated authority, service accounts, bulk changes, approval bypasses and revocation timing.[2][8][9]

AI switches an administrator can actually operate. Smartsheet's Admin Center already offers per-connector controls for ChatGPT, Claude, Copilot, Gemini and other MCP connections, alongside Smart Assist and public MCP access. Disabling a Smartsheet-side connector prevents its next authentication; it does not govern what a third-party assistant retains after data has left the platform. This distinction is essential in a multi-agent estate. OAuth scopes, group-based write entitlements, the human identity used for each action and downstream AI-provider retention terms all need documented policy. Smartsheet's MCP documentation says the server enforces existing user permissions; it also warns that external provider data handling follows the customer's separate agreement.[10][11][12]

Security telemetry that reaches security operations. New Security Event Streaming sends activity into SIEM tools instead of relying on periodic Events API pulls. Implementation documentation is unusually concrete: an Enterprise System Admin needs Safeguard or Advance Platinum, a TLS-capable syslog destination, an OCSF-compatible receiver, certificate management and a tested connection. Events typically arrive within minutes. There is no historical backfill, duplicate feeds are possible, and Smartsheet Gov does not support the built-in streaming integration. Event Reporting includes downloads, shares, API calls and administrative changes—but CIOs should test which agent-specific decisions and tool invocations are attributable end-to-end, not infer that every reasoning step is logged.[13][14]

Classification and assurance, with boundaries. New Safeguard Data Classification lets admins define sensitivity labels, users tag sheets, and security teams report on labels and changes. That is useful information governance, but a label is not automatically a policy that blocks every AI read, export or integration. Verify actual enforcement paths. Smartsheet also announced ISO/IEC 42001 certification and an Australian IRAP assessment at OFFICIAL: Sensitive. The former certifies its AI management system, not a customer's deployment; the latter is an assessed security posture, not universal accreditation for all workloads. Request the current certificate, scope, assessor documentation and contract-specific regional coverage. Some older trust documentation still describes ISO 42001 as in progress—a documentation lag that procurement should resolve explicitly.[2][15][16][6]

Integration is the new failure surface. Native integrations for ServiceNow, Dynamics 365, HubSpot and Workday are announced as available today. A natural-language Integration Agent for GitHub, Zendesk, Airtable and Monday.com is in early access or pending EAP rollout. The attractive promise is less ETL maintenance; the risk is automating ambiguous field mappings and granting broad access to badly understood workflows. IT needs integration ownership, source-of-truth rules, schema drift detection, reconciliation evidence and a switch to disable or reverse each flow. A one-sentence integration request is not a one-sentence integration architecture.[2][8]

Smartsheet also previewed Advanced Portfolios, aiming to propagate standardized changes across active projects and surface portfolio-level AI insights. That addresses a long-standing administration problem, but should not be confused with the existing basic Portfolios behavior: current documentation says ordinary template changes affect future projects, not automatically existing workspaces. Demand a feature-specific release and migration matrix before assuming retroactive control is available.[8][17]

Finally, the company is testing a deeper economic shift: a platform fee plus pooled consumption credits for AI, automations and public API activity, with Admin Center visibility by person and plan. It could reduce seat-tax friction for occasional collaborators. It could also replace predictable license growth with agent-driven consumption volatility. Smartsheet says early conversions receive overage protection; buyers should verify actual credit conversion, throttles, expiry, unit costs and contract caps before treating "pay for value" as proven. Credits measure platform activity, not necessarily delivered enterprise outcomes.[8]

The crowded market is changing the CIO's buying decision

The enterprise already has Microsoft Planner and Project, Atlassian Jira, Asana, Monday.com, Adobe Workfront, Wrike, Planview, Airtable, Notion, ClickUp and often ServiceNow workflows, alongside ERP, CRM and service-management records. Each owns a partial truth; each is moving toward AI. Forcing every team onto one board would be expensive and operationally disruptive. Buying yet another universal "AI work hub" without a federation policy would be equally naïve.

Smartsheet's opportunity is to become the execution realization layer above this sprawl, not merely to win another departmental application bake-off. But it faces competition from the existing systems of record and their native agents, as well as from integration platforms and enterprise knowledge graphs. A work graph that cannot trace dependencies to a signed decision, current source data and explicit authority is just a more sophisticated dashboard.

There are at least four concrete failure modes: stale or optimistic source updates produce false assurance; inconsistent identifiers join the wrong initiative to a budget; agents inherit permissive human access and amplify the blast radius; and usage-based pricing rewards repeated machine activity even when outcomes do not improve. Human approval of strategic changes helps, but is insufficient if the recommendation itself rests on unverified data.

This is why Smartsheet deserves a hard but fair assessment. The company has real enterprise administration pedigree, has shipped usable agent and governance capabilities, and is treating interoperability as a necessity rather than an afterthought. It has not yet demonstrated that vsActual can reconcile competing systems, prove outcomes and govern agent execution at scale. The graph must become an evidence graph, not a confidence machine.

Exactly three CIO recommendations

1. Establish an execution evidence contract before selecting a strategic work layer. Within 60 days, choose two cross-functional transformation programs and define canonical initiative IDs, authoritative financial and staffing sources, dependency owners, data freshness SLAs, immutable decision history and exception reconciliation. Require any vsActual pilot—or alternative platform—to reproduce the same budget-to-outcome lineage with traceable source records. Decision trigger: advance only when material deviations can be independently reproduced and contested by finance, PMO and IT. Why now: AI-generated executive summaries are arriving before enterprise data contracts are ready.

2. Treat work agents as production identities with explicit blast-radius limits. Pilot the shipped Smartsheet agents inside a least-privilege tenant using named user delegation, separately approved write scopes, bounded schedules, canary sheets and rollback evidence. Connect Security Event Streaming to the SIEM, confirm the Safeguard entitlement, capture before/after events and test connector disablement and external-model retention boundaries. Decision trigger: permit broader autonomous workflows only after failed-write, revoked-access and adverse-data scenarios generate attributable alerts and safe recovery. Why now: scheduled AI activity is already available; no one needs to wait for vsActual to create an agent-control problem.

3. Contract for measured outcomes, not agent activity or a forced suite migration. Keep specialist work tools where they demonstrably deliver, and run a 90-day competitive evaluation of Smartsheet's federation approach against a representative existing stack. Track reconciliation accuracy, time-to-detect drift, decision latency, administrative effort, policy exceptions and cost per accepted, business-verified outcome. Model the proposed pooled credits at low, expected and runaway-agent volumes, with negotiated caps and exit rights. Decision trigger: consolidate or buy the new layer only when verified execution improvement exceeds integration and governance costs. Why now: the market is shifting from seats to automated actions, and contract mistakes could outlast the first wave of AI features.

Smartsheet is trying to move fast enough to stay central as the enterprise workplace changes from people updating tasks to agents executing work. Its instincts are increasingly sound. The CIO's job is to make sure speed buys verified delivery, not a faster illusion of control.

Sources

- Smartsheet: vsActual launch announcement
- Smartsheet: ENGAGE 2026 agents, integrations and security announcement
- Smartsheet: ENGAGE 2026 detailed product recap
- Smartsheet: vsActual product and availability FAQ
- SiliconANGLE: independent ENGAGE 2026 reporting
- PMI: 2025 strategy–execution research
- Smartsheet: agent capabilities and approvals
- Smartsheet: AI Management in Admin Center
- Smartsheet: MCP connector administrator controls
- Smartsheet: MCP server architecture and permissions
- Smartsheet: SIEM streaming setup and restrictions
- Smartsheet: Event Reporting coverage
- Smartsheet: ISO/IEC 42001 scope
- Smartsheet: Australian IRAP assessment
- Smartsheet: enterprise trust and security controls
- Smartsheet: existing Portfolios behavior
- Smartsheet: enterprise plan entitlements

SOURCES / CITATIONS

Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

01 **Smartsheet vsActual launch announcement**
https://www.smartsheet.com/content-center/news/smartsheet-unveils-vsactual-ai-powered-future-work-execution-engage-2026

02 **Smartsheet ENGAGE 2026: agents, integrations and security**
https://www.smartsheet.com/content-center/news/smartsheet-adds-new-ai-agents-integrations-and-security-advancements-engage

03 **SiliconANGLE independent report, November 4 2026**
https://siliconangle.com/2026/11/04/smartsheet-ai-powered-platform-bridges-gap-between-planning-and-results/

04 **Project Management Institute 2025 execution research**
https://www.pmi.org/about/press-media/2025/new-pmi-research-reveals-strategy-execution-gap-is-undermining-transformation-and-how-to-close-it

05 **Smartsheet enterprise pricing and entitlements**
https://www.smartsheet.com/pricing

06 **Smartsheet security and trust**
https://www.smartsheet.com/trust/security

07 **Smartsheet vsActual early access and roadmap**
https://www.smartsheet.com/vsactual

08 **Smartsheet ENGAGE 2026 product recap**
https://www.smartsheet.com/content-center/product-insights/product-updates/engage-product-innovation

09 **Smartsheet agents permissions and approvals**
https://www.smartsheet.com/platform/agents

10 **Smartsheet Admin Center AI Management guide**
https://help.smartsheet.com/articles/2483722-ai-management-overview

11 **Smartsheet MCP connector controls announcement**
https://www.smartsheet.com/content-center/mcp-connector-controls-admin-center-now-available

12 **Smartsheet MCP security architecture**
https://www.smartsheet.com/content-center/product-insights/product-updates/introducing-smartsheet-mcp-server

13 **Smartsheet SIEM event streaming setup**
https://help.smartsheet.com/articles/2483761-set-siem-integration

14 **Smartsheet Event Reporting guide**
https://help.smartsheet.com/articles/2479296-event-reporting-monitor-how-your-organization-uses-data

15 **Smartsheet ISO/IEC 42001 scope**
https://www.smartsheet.com/trust/compliance/iso-42001

16 **Smartsheet IRAP assessment press release**
https://www.smartsheet.com/content-center/news/smartsheet-completes-irap-assessment-official-sensitive-level

17 **Smartsheet Portfolios feature documentation**

https://help.smartsheet.com/articles/2483673-portfolios-overview

FRONTIER.

TAKE THE NEXT STEP

Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

RESEARCH / BRIEFINGS

inquiries@thisisfrontier.ai

Briefings, advisory work, executive discussions, and research inquiries.

REPRINTS / LICENSING

reprints@thisisfrontier.ai

Reprints, enterprise distribution, excerpts, and licensing.

FRONTIER

www.thisisfrontier.ai

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

SEE FURTHER. MOVE SOONER.