FRONTIER.

FRONTIER / NEW ANALYSIS



AWS Transform unifies cloud migration economics and execution around a single inventory

AWS has combined migration assessment and wave planning in a single AWS Transform job, part of a broader shift from sequential migration studies toward continuously updated workflows that keep cost assumptions and execution plans aligned.

FRONTIER / Frontier Research / November 20, 2026

NEW ANALYSIS

Built to change a decision.

AWS has combined migration assessment and wave planning in a single AWS Transform job, part of a broader shift from sequential migration studies toward continuously updated workflows that keep cost assumptions and execution plans aligned.

EDITION **New Analysis**

AUTHOR **Frontier Research**

PUBLISHED **November 20, 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.

What changed

On November 19, 2026, Amazon Web Services updated AWS Transform so that customers can conduct migration assessment and migration planning within one master migration workstream instead of creating separate jobs for each activity.[1] The same uploaded infrastructure inventory now supports both the financial business case and the execution plan as an integrated whole.

AWS documentation actually describes discovery, assessment and migration planning as a single job step. Teams may run assessment or planning first, use either workflow independently and rerun each one as their source data improves.[2] This assessment can estimate prospective AWS costs and support a total-cost-of-ownership case. Planning can group infrastructure into applications and organize those applications into migration waves. Because both workflows use the same imported inventory, a change to the underlying data can feed a revised economic analysis, a revised migration plan or both.[3]

This is an important workflow upgrade rather than a new migration methodology. AWS has removed a key product and process boundary that previously separated the business-case job from the planning job. For example, a team can build its initial case, create migration waves and then estimate the cost of a particular wave without transferring inventory between separate AWS Transform jobs. It can also start with wave design and calculate the corresponding economics afterward, maintaining seamlessly traceability and transparency.[1]

AWS says customers can continue from the resulting assessment and plan into network migration, landing-zone deployment and the migration of servers, storage and containers within the same job. The company also says the update is available in every AWS Region where AWS Transform is offered.[1] Those availability and workflow-scope statements are vendor claims; the supplied sources do not include independent product testing.

The shared inventory is not necessarily a reliable inventory, however. AWS's discovery guide warns that RVTools exports omit network connections and database details, while vCenter exports can miss powered-off machines and usage peaks. Without connection data, AWS Transform can plan applications separately even when dependencies exist; without utilization data it uses default sizing assumptions, and missing SQL Server editions prevent licensing analysis.[3] These omissions can corrupt both the cost case and the migration-wave plan. A common job cannot correct the data it never received.

Why it matters

CIOs have often had to govern migration economics and migration execution as related but artificially separate work products. One team develops inventory assumptions and a cost case; another turns application dependencies, infrastructure data and organizational constraints into migration waves. When those activities rely on separate copies of the estate inventory, the approved business case can stop describing the plan that engineers are actually preparing to execute.

AWS Transform now gives customers a more powerful mechanism for keeping those artifacts connected. That does not make either artifact correct, but it makes them more consistent and seamless. It makes disagreements easier to expose because assessment and planning can draw from the same recorded estate. If a newly discovered server, storage requirement or application grouping changes a wave, the organization can rerun the corresponding assessment rather than reconcile separate jobs manually.[2][3] The practical value depends on inventory quality, ownership metadata and the discipline with which teams update the shared job.

The immediate procurement question is therefore not simply whether AWS Transform can produce a credible initial estimate. CIOs evaluating migration tooling must now ask whether economic assumptions remain properly traceable to the infrastructure and application groupings that will actually move. A lower-cost target architecture has limited governance value if the estimate excludes systems that later appear in a wave. Conversely, an executable wave plan may be difficult to fund if its resource assumptions cannot be connected to the approved cost model.

A shared workflow also changes when leaders should release migration funding. Rather than approving an entire program against a static assessment, they can require the cost model to be rerun when material inventory or wave changes occur. The relevant control is not a ceremonial reassessment at every planning edit. It is a defined threshold—such as a major application regrouping or a meaningful change in the infrastructure assigned to the next wave—that forces the financial case and implementation plan back into alignment. The FinOps Foundation distinguishes an estimate of potential costs from a forecast that finance and engineering must revisit as scope, timing and implementation assumptions change.[4]

There is an important platform tradeoff. Using AWS Transform for the business case, wave plan and downstream migration activities can reduce handoffs inside an AWS migration program, according to the workflow AWS documents.[1] It can also make AWS-specific tooling more central to planning before workloads move. Enterprises pursuing multiple clouds, retaining substantial on-premises capacity or requiring vendor-neutral portfolio analysis should decide which system owns the canonical inventory. AWS Transform can be an execution view without becoming the enterprise system of record.

CIOs should also resist treating workflow consolidation as proof of savings. No supplied source independently measures time saved, planning accuracy, migration success or cost reduction. The defensible benefit is narrower: assessment and planning can now reference one imported inventory and be rerun independently. Enterprises still need to test whether that design reduces reconciliation work in their estate and whether the resulting estimates survive comparison with invoices, contracts and observed utilization.

Frontier take

AWS Transform’s important move is not that two screens or tasks now sit in one product. It is that the migration business case can now become a living control tied to the same inventory used to decide what moves and when. That is the operating model CIOs should demand from migration platforms, whether AWS Transform is the primary tool or one component of a broader portfolio.

A migration estimate assembled before application grouping and dependency work is inherently provisional. Once planners divide the estate into waves, they make choices that can alter sequencing, temporary coexistence, network requirements and the set of infrastructure included in a given funding period. A workflow that permits assessment after wave design—or a fresh assessment after the inventory changes—allows financial approval to follow the plan's current state rather than an earlier snapshot.[2]

The same principle works in the other direction too. A wave may be technically coherent yet fail an enterprise funding constraint. Because AWS Transform allows teams to assess a single wave from the shared inventory, CIOs can use the cost view to challenge sequencing and scope before authorizing execution.[1] This turns migration planning into an iterative negotiation among architecture, finance and delivery rather than a one-way handoff from assessment to implementation.

The strong conclusion from all this is that enterprises should stop accepting migration business cases that cannot be reproduced from the current execution inventory. That standard is more important than choosing a particular vendor. If AWS Transform holds the operational view, the enterprise should preserve source ownership, transformation rules and exports needed to challenge its output. If another platform owns the estate inventory, teams should define how and when that system synchronizes with AWS Transform.

AWS has provided the product structure for this tighter, more integrated loop, but inventory checks are still no substitute for validating dependencies and licenses with application owners.[3] CIOs must measure whether the combined job actually reduces reconciliation work, shortens funding approvals and limits variance between estimated and executed scope. The new Transform integration is a highly useful new governance capability, but not yet evidence of guaranteed migration savings.

Three moves for CIOs

1. — Make inventory lineage a funding control Require every migration cost estimate and wave plan to identify the inventory version, import date and material exclusions behind it. If AWS Transform is used, preserve the source export and document any mapping or enrichment applied before import. Make finance, architecture and migration delivery accountable for approving the same inventory baseline rather than circulating separate spreadsheets.

- Decision trigger: Apply the control before approving a migration tranche, renewing a systems-integrator statement of work or authorizing a wave to enter execution.
- Why now: AWS Transform can now derive assessment and planning from one imported inventory and rerun either workflow.[2][3] That makes unexplained differences between the financial case and execution scope less defensible, while still requiring the enterprise to prove that the shared source is complete.

2. — Set thresholds that force economic reassessment Define changes that require the migration business case to be rerun: additions to a wave, substantial application regrouping, discovery of previously unrecorded infrastructure or a material shift in the next wave's estimated AWS footprint. Tie release of contingency funds to the refreshed assessment rather than allowing the original program estimate to absorb scope changes invisibly.
   - Decision trigger: Invoke reassessment whenever a planning change crosses the enterprise's approved variance threshold or changes the infrastructure assigned to a funded wave.
   - Why now: Assessment and planning can run in either order and be rerun independently in the same AWS Transform job.[2] CIOs can use that capability to connect financial approval to current execution scope without restarting the entire planning process after every routine edit.
3. — Separate the canonical inventory from the execution view Decide explicitly whether AWS Transform will own the enterprise migration inventory or consume a governed subset from a configuration, discovery or portfolio system. For hybrid or multicloud programs, retain a vendor-neutral canonical record and define synchronization, reconciliation and export requirements before AWS Transform becomes embedded in downstream migration activity.
   - Decision trigger: Make this decision before using AWS Transform across multiple business units, extending it into landing-zone and workload migration, or signing a long-term delivery agreement that depends on its job data.
   - Why now: AWS says customers can continue from the shared assessment and plan into network, landing-zone, server, storage and container migration work.[1] As one job spans more of the delivery chain, unclear inventory ownership becomes harder to reverse and can weaken independent comparison of AWS estimates with other deployment options.

Sources

- AWS What's New: AWS Transform combines migration assessment and planning in one experience [1]
- AWS Transform documentation: Discovery, assessment, and planning [2]
- AWS Transform documentation: Discovery and shared inventory [3]
- FinOps Foundation: Forecasting [4]

SOURCES / CITATIONS

Trace the argument.

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

01 **AWS What's New: AWS Transform combines migration assessment and planning in one experience**
https://aws.amazon.com/about-aws/whats-new/2026/11/aws-transform-assessment-and-planning/

02 **AWS Transform documentation: Discovery, assessment, and planning**
https://docs.aws.amazon.com/transform/latest/userguide/transform-migrations-discover-assess-plan.html

03 **AWS Transform documentation: Discovery and shared inventory**
https://docs.aws.amazon.com/transform/latest/userguide/transform-vmware-discover-source-data.html

04 **FinOps Foundation: Forecasting**
https://www.finops.org/framework/capabilities/forecasting/

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.