Enterprise AI spend is a reallocation, not a new budget line — and build-lite is now real

Ten US enterprise IT budget owners describe **flat software budgets** in which AI has grown from low-single-digits to mid-teens share, funded not by new money but by SaaS consolidation and seat right-sizing. The panel converges hard on "hold flat and reallocate." It splits on one question only: whether AI has changed the build-versus-buy calculus. That split is the story.

BOTTOM LINE

Treat enterprise AI demand as wallet migration, not expansion. Growth comes from displacing incumbent SaaS at renewal, so the defensible positions are suite-native consolidation and narrow, short-payback builds. Point tools at the edges are the funding source, not the market.

Held by 10/10 on flat-and-reallocate; the build-vs-buy split (below) is where the panel actually disagrees.

WHERE THE PANEL DIVIDES · Q3, BUILD VS BUY

- **4 consolidate** into suite-native AI — a custom build is a maintenance liability the vendor should carry
- **4 build-lite converts** — a narrow internal build now beats the vendor SKU on unit economics
- **2 constrained** — FDA/ISO validation and FERPA privacy freeze the calculus regardless of cost

WHAT THE PANEL AGREES ON

- **Budgets are flat; the mix is moving fast.** AI grew from low-single-digits to mid-teens share of the software line — one owner reports 2% to 19% in a year — entirely funded by consolidation and seat cuts, not net-new dollars.
- **Build-lite crossed the line from theory to practice.** One CTO's four-week internal build replaced a vendor SKU at roughly 60% lower unit cost, 90–95% accuracy, with a human-review lane. That is now a repeatable move, not an outlier.
- **Cuts hit the edges first.** Meeting transcription, writing and grammar, BI viewers, forms, and async video compress at renewal. The systems of record — ERP, CRM, HRIS, QMS — are augmented, never replaced.
- **Every AI dollar clears a gate.** 60–90 day KPI-gated pilots (10/10), no-train data guarantees (6/10), suite-native fit and audit/validation (7/10 each), hard spend caps and cost predictability (4/10). Nothing buys on demo alone.
- **The buying story is fixed.** To a CFO the pitch is "hold flat and reallocate," backed by a measured outcome. Growth-by-headcount narratives do not clear procurement.

Signals to track: AI moving to ~15–20% of the software mix; seat compression in notes / grammar / BI-viewer categories; pilot pass rates; vendor-count decline; and renewal-driven switches fuelled by suite bundling. Movement on these is the leading indicator of where the spend lands.

ABOUT THIS STUDY

Synthetic panel, n=10, directional — not a census. Recruited from FishDog's census-grounded population of 340,000 US residents; a regulated FDA/ISO voice and a K-12 operator were recruited deliberately as conservative outliers. 7 open-ended questions, fielded 2026-05-27. The full split, the procurement-gate matrix, and the methodology are in the accompanying **Decision Brief**; every response and every number is in **responses.csv** and **data.csv**. Positions on the build-vs-buy split are analyst-classified from the study's own segment and divergence tables.