// WHITE PAPER · PUBLISHED 16 SEPTEMBER 2026

CapEx, OpEx, and EBITDA: The Same Dollar, Two Different Stories

Why identical spend earns two different verdicts, and what that does to the number your board watches

Part 1 of 10 in *The Money Language Series*. Each piece stands on its own. The others: Part 2 · Part 3 · Part 4 · Part 5 · Part 6 · Part 7 · Part 8 · Part 9 · Part 10.

Jason Newell · MAX Research Collective · Report WPR-BIZ-002 · Published 16 September 2026
Licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0), https://creativecommons.org/licenses/by/4.0/
Original: https://www.maxresearchcollective.com/articles/capex-opex-and-ebitda.html
Cite as: Newell, J. (2026). CapEx, OpEx, and EBITDA: The Same Dollar, Two Different Stories. MAX Research Collective, WPR-BIZ-002. https://www.maxresearchcollective.com/articles/capex-opex-and-ebitda.html

A very good architecture proposal died in a budget meeting because the CTO could not answer one question: "What does this do to EBITDA?"

The proposal was right. The migration would have cut five-year cost by a third. The savings showed up as a smaller cloud bill three years out, though, and the spend showed up as operating expense starting next quarter. The CFO saw margin compression. The CTO saw a better system. Both were reading the same numbers off different pages, and the person who could read only one page lost.

This series is about reading both pages.

The distinction that drives everything

Money you spend to acquire or build something durable is capital expenditure. It goes on the balance sheet as an asset and gets expensed slowly across its useful life. Hardware depreciates, software and other intangibles amortize, and those are one mechanic under two names.

Money you spend to run the business this period is operating expenditure. It hits the income statement now.

The same dollar lands on a different page, and the page it lands on changes how your company looks to everyone who matters.

The pivot is EBITDA, earnings before interest, taxes, depreciation and amortization. It is the market's rough stand-in for operating cash generation. Private equity buys on a multiple of it, lenders write covenants against it, and boards benchmark margin with it. By construction it excludes depreciation and amortization.

Which produces the fact that trips up most technology leaders the first time they meet it: a move that lowers your total cost can make your EBITDA look worse. Servers you bought sat below the EBITDA line as depreciation. Cloud invoices sit above it as OpEx — better economics, uglier headline. That is arithmetic rather than a trick, and it has shaped more infrastructure decisions over the past decade than any technical argument has.

It runs the other way too. Capitalize aggressively and EBITDA looks strong this year, while you build an amortization tail that outlives your tenure, surfaces in every quarter after you have moved on, and gets normalized straight back out by any buyer who reads the footnotes. Internal debates usually skip that step.

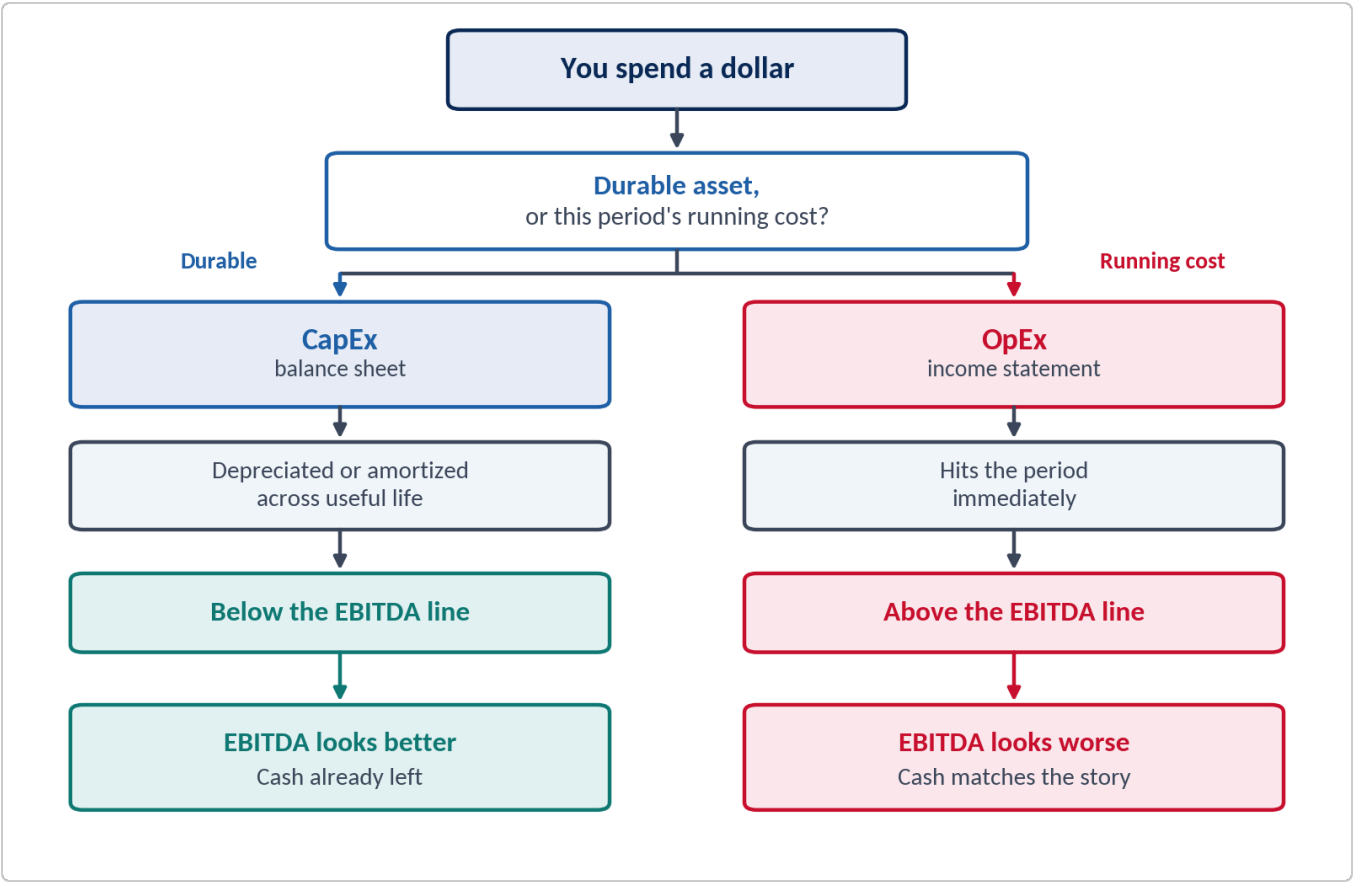


Figure 1. Follow the dollar to the line. The balance-sheet route flatters EBITDA after the cash has already gone.

Three waves, one recurring pattern

Three industry-wide ”this changes everything” cycles have come through since I started doing this work. The technology differed each time. The financial argument stayed almost identical.

The BI and data warehouse era

That era was CapEx-native to its bones. You bought appliances and storage arrays, perpetual licenses with annual maintenance, and a large integration project. Hardware depreciated over three to five years and implementation labor was often capitalized. The pain was lead time and sizing: you provisioned for peak and for three years out, and being wrong cost real money in both directions.

The cloud-first era

Cloud sold itself on elasticity. What it delivered was a structural migration of IT spend from CapEx to OpEx. That is why "cloud-first" so often needed the CFO's sponsorship rather than the CTO's enthusiasm, and why the strategy landed differently at a venture-funded startup (where burning OpEx is the plan) than at a leveraged mid-market business (where every covenant watches EBITDA).

Two very different projects hid under one banner. Lift-and-shift moved the workload and kept the architecture — renting the same oversized boxes you used to own, minus the depreciation shelter, plus egress. Re-architecture redeemed the pricing model, because consumption pricing pays only when the application can scale down.

The correction arrived on schedule. 37signals left AWS after roughly $3.2M a year, spent about $700K on Dell hardware in a colo, reported close to $2M in annual compute savings, then replaced S3 with about $1.5M of storage hardware. Their projection runs north of $10M saved over five years, which is roughly three years of their old AWS bill. Dropbox did the same thing at larger scale years earlier. GEICO's platform leadership described a decade in the cloud where bills rose roughly 2.5x without finishing the migration.

Cloud is a pricing model optimized for variance. Steady, high-utilization, data-heavy workloads pay a premium for flexibility they never use.

The AI era

AI spend splits three ways at once, which is why budgets are a mess right now — one label, three different things underneath it.

AI SPEND TYPE	TREATMENT	FORECASTING DIFFICULTY
Tooling seats and API tokens	OpEx, immediate	High. Usage is behavioral rather than provisioned
Implementation and re-architecture	Often a capitalizable project	Medium. It has a start and an end
Dedicated GPU infrastructure	CapEx, or a capital-like commitment	Medium, and the useful-life call is brutal

The same drama plays out at the top of the market. The big three hyperscalers spent years pushing assumed server life from three or four years out to six, which lowers annual depreciation and lifts reported income. Then in early 2025 Amazon moved part of its fleet back from six years to five, citing the pace of AI development, taking roughly $700M off operating income plus about $920M in accelerated depreciation. Meta went the other way and extended to 5.5 years, cutting expected 2025 depreciation by about $2.9B, which is four times the size of Amazon's operating-income hit in the opposite direction.

Useful life is an estimate rather than a fact, and it is one of the largest earnings levers in technology. Buy GPUs and you will make that call yourself.

Four rulebooks, and they do not agree

You do not need to be an accountant. You do need to know that four separate sets of rules apply to your work, and that they answer different questions.

RULEBOOK	GOVERNS	THE THING TO KNOW
ASC 350-40, internal-use software	Software you build to run your business	Rewritten by ASU 2025-06, effective for annual periods beginning after 15 Dec 2027
ASC 985-20	Software you sell or license	Hinges on "technological feasibility." Untouched by ASU 2025-06
ASU 2018-15	Cloud arrangements	Subscription fees expensed. Implementation costs can be capitalized, amortized over the hosting term
IRC Section 174A, Section 41, Section 168(k), Section 179	Tax	A completely separate set of books from GAAP

Two changes from the last eighteen months matter more than the rest.

ASU 2025-06 removed the stage model. Issued 18 September 2025, it strips out the old preliminary, application-development and post-implementation framework that assumed waterfall, and replaces it with a threshold: management has authorized and committed funding, and completion is probable, with a brake for software carrying significant development uncertainty. Most practitioners expect less capitalization and more expensing. That lands on EBITDA. If you have EBITDA-based covenants, that conversation belongs with your lender now rather than in 2028.

Tax reversed, and it now discriminates by geography. The One Big Beautiful Bill Act restored immediate expensing of domestic research costs through new Section 174A for tax years beginning after 31 December 2024, permanently. Foreign research, offshore software development included, stays on a 15-year amortization schedule. That asymmetry is now a first-order input into where you put engineers. Separately, 100% bonus depreciation is permanent for qualified property acquired and placed in service after 19 January 2025, and Section 179 expensing rose to $2.5M with phase-out starting around $4M, up from $1.22M in 2024.

Put plainly: the tax code currently likes it when you buy hardware and when you build software in the United States. That is a different backdrop from 2022 through 2024, when mandatory R&D amortization quietly wrecked cash flow at software companies that had done nothing wrong.

Placement is a workload question

"Cloud-first" and "repatriate everything" are both slogans, and the unit of decision is the workload. The honest version fits in a table.

FACTOR	POINTS TOWARD CLOUD	POINTS TOWARD COLO OR OWNED
Demand pattern	Spiky, seasonal, unknown	Steady, high utilization
Data volume and egress	Low	Large, egress-heavy
Team capability	No platform or SRE bench	Real infrastructure expertise
Capital	Scarce or expensive	Available and cheap
Time horizon	Uncertain product direction	Long-lived commodity workload
Compliance	Easier to inherit certifications	You already run controls

What people leave out of the model, every time, is exit cost: egress, data gravity, dual-running during migration, and rewriting against proprietary managed services. Every migration business case I have read understates the cost of leaving.

Commitment instruments go missing too. Reserved Instances and enterprise discount agreements show up as OpEx on the P&L and behave like capital obligations. A three-year commit signed against optimistic growth is a liability wearing an operating-expense costume.

Newell, J. (2026). MAX Research Collective, WPR-BIZ-002. CC BY 4.0

Waste is getting worse. Flexera's 2026 State of the Cloud Report put self-estimated wasted IaaS and PaaS spend at 29%, the first increase in five years, attributed largely to AI workloads making rightsizing and forecasting harder. That figure comes from a population where 63% of organizations run formal FinOps teams.

The staffing question, asked properly

The tempting formula runs: AI makes engineers faster, therefore fewer engineers. The evidence says slow down.

DORA's 2025 report, covering nearly 5,000 professionals, found roughly 90% AI adoption and, reversing the prior year, a positive relationship between AI adoption and delivery throughput. It also found a persistent negative relationship with delivery stability. Their framing is that AI amplifies. It magnifies what is already there, good or bad, and it exposes weak testing and review downstream.

METR's randomized trial found experienced developers on their own mature repositories took 19% longer with early-2025 tools while believing they had been 20% faster, a 39-point gap between perception and measurement. METR now labels that result historical, and their follow-up cohort with agentic tools showed a much smaller effect with a confidence interval spanning zero. Both readings survive.

Orgvue surveyed 1,000 senior decision-makers and found that 39% had made redundancies after deploying AI, and that 55% of those said they got the role choices wrong. Gartner projects half the organizations that cut customer service headcount for AI will rehire similar roles by 2027. Klarna said publicly that its cost-cutting went too far and restarted human hiring.

My position — and the evidence above is the whole of the reasoning — is to treat AI as a capacity multiplier before a headcount reducer. Ship more roadmap with the same team for two quarters and measure it in your own codebase, then decide. Then budget for the new bottleneck; it is review, verification, testing and security, and it is where a team that has doubled its output without touching its QA function discovers that the output was never the constraint. Throughput gains landing on an unchanged QA function turn into incidents rather than velocity.

What I actually do

1. Write the capitalization policy before the project starts, with the CFO in the room. Retroactive capitalization arguments fail audits and burn credibility. Under the new rules your funding-approval artifacts and project governance become accounting evidence.
2. Decide placement workload by workload. Hybrid is the equilibrium, and it reflects judgment rather than indecision.
3. Model five years including exit.
4. Separate one-time from run-rate in every business case. A migration is a project and hosting is a run-rate, and mixing them is how boards get surprised.
5. Negotiate provider incentives before committing, then keep the credit away from the architecture decision.
6. Put the EBITDA impact in the proposal as a named line. Someone will ask, and the CTO who has the answer gets the budget.
7. Leave EBITDA alone as an engineering target. A capitalization tail outlives the person who created it.
8. Say out loud that EBITDA is not cash. It excludes the capital you actually spent.

Glossary

TERM	DEFINITION
Amortization	Spreading the cost of software and other intangibles across their useful life
ASC 350-40	US GAAP rules for software you build for internal use
ASU 2025-06	The 2025 rewrite of ASC 350-40 that removed the project-stage model
CapEx	Spend on something durable. Goes to the balance sheet, expensed over time
Depreciation	Spreading the cost of physical assets across their useful life
EBITDA	Earnings before interest, taxes, depreciation and amortization. A proxy for operating cash generation
Egress	Charges for moving data out of a cloud provider
FinOps	The practice of managing cloud spend as a shared engineering and finance discipline
OpEx	Spend on running the business now. Hits the income statement immediately
Section 174A	Tax code section restoring immediate deduction of domestic research costs
Useful life	How long you estimate an asset stays productive. An estimate rather than a fact

References

- FASB ASU 2025-06 summaries: AICPA & CIMA, Forvis Mazars, Baker Tilly, Grant Thornton, PKF O'Connor Davies
- FASB ASU 2018-15: Deloitte Heads Up, RSM US

Newell, J. (2026). MAX Research Collective, WPR-BIZ-002. CC BY 4.0

- Section 174A and OBBBA: RSM US, Kaufman Rossin, Anchin, Bennett Thrasher
- Bonus depreciation and Section 179: CLA, CBIZ, Texas Society of CPAs
- Hyperscaler useful life: Amazon and Meta 10-K disclosures; theCUBE Research; SiliconANGLE
- Cloud repatriation: DataCenterDynamics on 37signals; The New Stack on GEICO; a16z survey data
- Flexera 2026 State of the Cloud Report
- DORA 2025 State of AI-assisted Software Development
- METR, Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity, and the February 2026 design update
- Orgvue annual workforce research; Forrester Predictions 2026; Gartner

Best next read

Part 2, Reading the Financials: The Primer for Everyone Who Never Had To Care. If balance sheet, gross margin and Rule of 40 are phrases you nod along to rather than use, start there. If you already speak that language fluently, skip to Part 3, Capitalize or Expense: Your Jira Board Is Now Audit Evidence, which is where the new rules bite.

The full series

PART	TITLE	WHAT IT ANSWERS
1	CapEx, OpEx, and EBITDA: The Same Dollar, Two Different Stories, which is this one	Why does the same spend earn two different verdicts?
2	Reading the Financials: The Primer for Everyone Who Never Had To Care	What are they actually looking at?
3	Capitalize or Expense: Your Jira Board Is Now Audit Evidence	Which engineering hours become an asset?
4	Depreciation and Useful Life: The Chip Outlives Its Economics	How long does this asset really last?
5	Placement Economics: The Right Answer, Applied Everywhere	Where should this workload live?
6	Lift-and-Shift or Re-Architect: The Two That Modernized Later	When is the cheap migration the expensive one?
7	Vendor Economics: The Email in November	Where does recurring spend quietly compound?
8	AI Spend: Three Budgets Pretending to Be One	What part of the AI bill is a project?
9	AI and Headcount: Somebody on the Board Did the Arithmetic	Can you responsibly plan a smaller team?
10	Outsourcing After Section 174A: Eleven Months to Relearn It	What changed now that geography is taxed?

This series is an introduction for technology leaders. It is accounting and tax context rather than accounting or tax advice. Treatment is fact-specific and elections have deadlines. Model it with your CFO and tax advisor before you commit.

Newell, J. (2026). MAX Research Collective, WPR-BIZ-002. CC BY 4.0

CAREER & LEADERSHIP WPR-BIZ-002

// ABOUT THE AUTHOR

Jason Newell leads organizations and people in technology, and is a researcher, deep technologist and builder — from a CRAY to agentic AI. He founded MAX Research Collective to build AI for regulated industries that runs where its owner has control, with healthcare as its strongest focus. About Jason → Connect on LinkedIn →