



BEYOND ACCIDENTAL MULTI-CLOUD

Multi-Cloud Is Not a Strategy — Why Choosing Every Cloud Governs None

How the promise of vendor choice quietly became a tax on control across the enterprise IT estate

Multi-Cloud Is Not a Strategy — Why Choosing Every Cloud Governs None

Transformation · No. 197

Originally written March 2016 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

"Two providers, deliberately governed, is a portfolio."

1. The Question Nobody Could Answer

The moment that exposes most cloud strategies is not a failure. It is a simple question, asked in a finance review, that no one in the room can answer: what does the cloud cost us to run?

Three teams reply. The digital group reports its figure from one provider's console. The data platform team reports another, from a different provider bundled into an enterprise licensing agreement signed two years earlier. A machine-learning group, working on a third platform a data scientist happened to prefer, produces a third number from a spreadsheet. Each figure is accurate. None can be added to the others with any confidence, because each rests on a different billing construct, a different set of tags — or none — and a different definition of what counts as "cloud." The organisation is spending heavily across three providers and cannot state, to the nearest large sum, what it is spending or what it is getting.

This is the condition I want to name plainly, because it is far more common than the strategy decks admit: multi-cloud as most enterprises actually run it is not a portfolio. It is sprawl with a flattering name. And the governing belief underneath it — that having more than one cloud is itself a prudent, diversified position — is close to the opposite of the truth. Choice without a governing spine does not spread risk. It multiplies cost and dissolves control.

2. Multi-Cloud Is Rarely Chosen

Start with how organisations arrive here, because the honest account is uncomfortable. Almost no enterprise sits down, weighs three providers against a coherent set of criteria, and deliberately commissions a multi-cloud operating model. What happens instead is a sequence of entirely reasonable local decisions.

- A product team building something new starts on the provider it already knows, because speed matters and the skills are in the room.
- A data platform lands on a second provider because it arrived bundled inside a wider software agreement, and using it looked like getting something for nothing.

- An analytics or machine-learning effort drifts to a third, because that is where the best tooling for the job sat and the team leading it had a preference.

Each of these is defensible on its own terms. None is wrong in isolation. But their sum was never designed. It is the residue of uncoordinated choices, and it inherits the properties of an accident rather than a plan: no shared identity model, no common control baseline, no single place where cost, risk, and architecture are seen together. "Multi-cloud" then becomes the label we apply after the fact to make the accumulation sound intentional.

The most expensive words in enterprise technology are "we're multi-cloud" spoken as though it were a strategy, when it is in fact the description of a decision nobody made.

3. The Lock-In Argument, Taken Seriously

The strongest defence of running several clouds deserves a fair hearing, not a caricature, because it contains something real. It runs like this: committing to a single provider is imprudent. It concentrates operational risk in one supplier, hands that supplier enormous leverage over future pricing, and leaves you hostage to its roadmap and its outages. Spreading across providers is simply the mature, diversified posture — the same instinct that says not to put a whole portfolio into one holding.

The instinct is sound. The pricing of it is where the argument fails.

Lock-in avoidance is a form of insurance, and like any insurance it should be judged by premium against payout. The premium for genuine multi-cloud portability is paid continuously and in full: duplicated tooling and duplicated controls, skills spread thin across three platforms instead of deepened on one, and — most costly of all — an architecture disciplined down to the lowest common denominator the providers share, so that nothing you build can lean on the differentiated services that were the reason to leave your own data centre in the first place. The payout, meanwhile — actually lifting a material workload off one provider and running it on another — is a claim almost no one ever files. And on the rare occasion someone tries, they find the portability they paid for was never really there, because the moment a workload used anything distinctive about its host, it stopped being portable.

"You pay the lock-in premium every month; you almost never collect on the policy."

So the choice is not the one the argument sets up. It is not single cloud versus multi-cloud. It is governed multi-cloud versus accidental multi-cloud — and the lock-in case, properly examined, is an argument for governing the estate deliberately, not for letting it accrete.

4. Why the Governance Fails

When organisations do try to get a grip, they usually reach for the wrong hand. Multi-cloud is treated as a procurement problem: which providers, on what commercial terms, at what discount. Vendor management runs it. And so the governing effort concentrates on the one dimension — the commercial one — where fragmentation hurts least and control matters least.

Consider what the fragmentation does to the economics. A consolidated commitment to a single provider — the kind that unlocks the deeper discount tiers — might earn an effective discount in the region of a quarter off list. Split the same total spend across three providers, and no single relationship is large enough to reach those tiers; the effective discount thins toward single digits. The organisation set out to use competition between vendors as leverage and instead surrendered the only leverage that scale confers, paying close to full price to each because it is a large customer to none.

But the deeper failure is not commercial, and this is the crux. The binding constraints of a multi-cloud estate are architectural and operational, and they are exactly the things a procurement lens cannot see.

1. Identity. Three providers mean three access-management models. A single control — say, enforcing that no storage is publicly reachable — now has to be authored, tested, and maintained three times, in three dialects, and it drifts out of alignment between them. Work that was a day's effort once becomes a week's, and the result is less consistent, not more.
2. The network and its edges. Connectivity, segmentation, and traffic inspection must be reasoned about across three fabrics that do not share primitives.
3. Cost attribution. Without a common tagging and account discipline spanning all three, no one can answer the finance question we started with — which is why we started with it.
4. Observability. Three sets of logs and metrics, in three formats, mean that seeing the whole estate at once requires building something none of the providers gives you.

And in this particular moment there is a fifth constraint, one that has moved from the technical margins to the boardroom: where the data physically sits, and under whose law. The collapse of the Safe Harbor arrangement last autumn removed, almost overnight, the settled basis on which data moved between Europe and the United States. The Privacy Shield framework announced only last month is meant to replace it, but it is not yet in force and few practitioners regard its durability as settled. The forthcoming European data protection regulation promises to raise the stakes further still. In that climate, "which cloud, in which region, holds which data" is no longer an infrastructure detail. It is a compliance question a board is entitled to ask — and in an accidental multi-cloud estate it is frequently a question no one can answer with confidence, because the data landed wherever each local decision happened to put it.

5. The Governing Spine

The alternative is not centralised control that strangles the delivery speed the cloud was meant to give. It is a thin, deliberate spine that lets teams move fast within lanes someone actually drew. Three elements carry most of the weight.

1. A placement policy — what belongs where, and why. Not a decree that everything runs on one provider, but an explicit, written rationale for which classes of workload and data belong on which platform, including the data-residency constraints that settle some placements outright. Placement becomes a decision made against criteria, not an accident of who started first.
2. A minimum common control plane. A small, non-negotiable set of controls that spans every provider: one identity model teams federate into rather than three they each reinvent; a mandatory tagging and account standard so cost and ownership are attributable everywhere; a security baseline expressed once and enforced consistently. Everything above this line can vary by team and by platform. This line cannot.
3. A single accountable owner for the portfolio as a portfolio. Someone whose remit is the whole cloud estate — cost, risk, architecture, and compliance seen together — rather than three teams each optimising their own corner. This is the role whose absence the opening finance review exposed.

None of this requires choosing fewer clouds. It requires governing the ones you have as a coherent whole rather than as three private kingdoms that happen to send their invoices to the same company.

6. Governed, or Accidental

I have watched capable organisations spend two years and considerable money assembling a multi-cloud estate, and then spend the third year discovering they had bought the costs of every provider and the leverage of none. The failure was never the number of clouds. Two providers, deliberately governed, is a portfolio. Two providers that accreted from uncoordinated choices, ungoverned, is an exposure — and a third and a fourth only deepen it.

So the question worth putting to any organisation proud of being multi-cloud is not how many clouds it runs. It is a quieter one, and a more revealing one: who decides what runs where, on what basis, and can they tell you — today, without a week's reconciliation — what the whole estate costs and where the data sits? If the answer comes back confident, the choice is real and worth having. If it comes back as a pause and three different spreadsheets, then what looks like a strategy of choice is the absence of one, and the chaos is not a risk to be managed later. It is already the operating model.