



BEYOND THE HANDOVER

The Data Centre Exit: Why Infrastructure Never Becomes Someone Else's Problem

The workloads leave; the cost, the accountability, and the regulator stay — a reflection on cloud migration in the regulated enterprise

Contents

1. Executive Summary	3
2. The morning the racks went dark	3
3. The seductive arithmetic of the exit	4
4. "Someone else's problem" is a category error	4
5. What moved, and what stayed behind	5
6. The regulator does not recognise the handover	6
7. Why the pattern persists	7
8. The strongest case against this reading	8
9. What the exit actually is	8

The Data Centre Exit: Why Infrastructure Never Becomes Someone Else's Problem

Transformation · No. 256

Originally written February 2017 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

"Infrastructure never becomes someone else's problem."

1. Executive Summary

For three years the dominant instinct in enterprise technology has been to get out of the business of running data centres. The logic is close to irresistible: stop pouring capital into racks and cooling, convert a fixed cost into a variable one, and hand the undifferentiated heavy lifting to providers who do it at a scale no single organisation can match. Framed that way, the data centre exit is not a project but a liberation — infrastructure becomes, at last, someone else's problem.

This essay is an argument that the phrase is where the trouble starts. Watching organisations move through their migration programmes over the past two years, a consistent pattern emerges: the workloads leave, and almost nothing else does. The operating model, the accountability, the regulatory exposure, and the habits of a generation of engineers all stay exactly where they were, now pointed at an estate no one can walk into. The exit transfers the operation of infrastructure while leaving the ownership of its consequences firmly in place — and because the business case was written as a disposal rather than a handover, the organisation is rarely staffed, funded, or governed for the relationship it has actually entered.

The point is not that the exit is wrong. In most cases it is right. The point is that "someone else's problem" is a category error that quietly sabotages the very benefits the move was meant to deliver, and that the pattern persists not through stupidity but because the structural forces around the business case, the operating model, and the profession's own self-image all reward the illusion. What follows traces why.

2. The morning the racks went dark

There is a particular kind of ceremony that has become familiar. The final workload is cut over on a weekend, the migration programme declares success on the Monday, and a photograph circulates of the last physical server being switched off — sometimes literally unplugged by a smiling executive. The programme director stands down the war room. The slide that has been shown to the board for eighteen months, the one with the burn-down of "applications remaining in the data centre," finally reaches zero.

Everyone is, briefly and genuinely, delighted.

Then the bill arrives, and the questions begin. Not immediately — the first months are cushioned by migration credits and the sheer relief of no longer firefighting a decaying facility. But somewhere around the second year, a finance business partner asks why the run-rate is higher than the estate it replaced, and no one in the room can answer with confidence. The people who used to know how the infrastructure behaved have moved on, been reorganised, or were never rehired after the facility closed. The estate is now a set of line items on a monthly invoice, and the invoice is the only map anyone has.

I have watched this sequence enough times to stop treating it as bad luck. It is the predictable consequence of a framing. When an organisation tells itself it is exiting infrastructure, it prepares for an ending. What it has actually signed up for is a beginning.

3. The seductive arithmetic of the exit

It is worth stating the case for the exit at its strongest, because it is a good case and pretending otherwise explains nothing. The economics of running your own data centre have genuinely deteriorated. A private facility is a large, lumpy capital commitment sized for a peak you hit rarely and a future you cannot forecast. You pay for the peak all year, you refresh hardware on a depreciation cycle rather than a needs cycle, and you carry a standing team whose deepest expertise — power, cooling, cabling, hardware lifecycle — is precisely the expertise that differentiates your business not at all.

Against that, the public cloud proposition is compelling on almost every axis a chief financial officer cares about. Capital expenditure becomes operating expenditure. Utilisation becomes, in principle, something you pay for only as you consume. Provisioning that once took a twelve-week procurement cycle takes minutes through an interface. And the provider's scale buys a security posture, a resilience engineering capability, and a pace of feature delivery that an internal team of any realistic size cannot rival. When people say the heavy lifting is undifferentiated, they are right: almost no organisation creates competitive advantage by being unusually good at operating uninterruptible power supplies.

So the direction of travel is sound. The mistake is not in choosing to go. The mistake is in believing that the thing you are handing over is the whole of the thing you are responsible for.

4. "Someone else's problem" is a category error

Here is the distinction the phrase elides. You can outsource the operation of infrastructure. You cannot outsource the accountability for what that infrastructure does. The provider will run the hardware, patch the hypervisor, and replace the failed disk you never see. What the provider will not do is decide how your workloads are

architected, whether your data sits in a jurisdiction your regulator accepts, how much you spend, or what happens to your customers when a region you depend on goes dark. Those remained yours the entire time. They were never on the table.

You can transfer the operation of infrastructure to a provider who does it better than you ever could. You cannot transfer the accountability for what runs on it – to your customers, your board, or your regulator. The exit changes who holds the spanner; it does not change who answers the phone.

This is why the exit so often feels like it has failed even when the migration itself was flawless. The technical programme delivered exactly what it promised – the workloads run, the facility is closed – and yet the organisation feels less in control of its infrastructure than before, not more. That feeling is accurate. Before the exit, the infrastructure was visible, walkable, and staffed by people who could answer questions in a corridor. After it, the same accountabilities are discharged through a portal, an invoice, and a support contract, mediated by a provider whose incentives are aligned with yours only in the general case and never in the specific one.

“We did not dispose of the data centre. We changed its address, dismissed the people who held its map, and told ourselves we had solved a problem when we had merely made it invisible.”

5. What moved, and what stayed behind

The clearest way to see the gap is to look at what a typical migration actually transports. Overwhelmingly, the dominant pattern of the last two years has been "lift and shift" – rehosting existing workloads more or less as they are onto rented virtual machines. It is the fastest route to closing the facility, and closing the facility is what the business case rewards. But a rehosted workload is a data-centre workload wearing a different badge. It is sized for peak and runs around the clock. It assumes the storage, the network, and the failure modes it was built against. It has learned none of the behaviours – elasticity, statelessness, graceful degradation – that make the new environment economical.

Consider a composite that will be recognisable to anyone who lived through one of these programmes. A back-office estate of roughly four hundred servers is rehosted over a year. The business case promised a thirty per cent reduction in infrastructure cost. Eighteen months after cutover, the run-rate is twenty per cent higher than the facility it replaced. Nothing went dramatically wrong; a dozen small things went quietly right for the provider and wrong for the customer. The instances were sized to the old peak and left running overnight and at weekends because no one owned the decision to

switch them off. The commitment discounts available for reserving capacity in advance were never taken up, because taking them required a forecast no one felt able to sign. Duplicating data across regions to satisfy residency requirements added transfer and egress charges that the model never included. And the team that used to sweat utilisation — because a full data centre was a visible, embarrassing constraint — had been disbanded on the grounds that infrastructure was now someone else's problem. The bill arrives monthly and itemised, but the accountability it represents was never a line item anyone agreed to own.

The exit as sold	The exit as lived
A fixed capital cost becomes a lower variable one	A variable cost that drifts upward because no one owns it
Undifferentiated heavy lifting is handed away	Differentiated decisions — sizing, residency, resilience — quietly handed away with it
The standing infrastructure team is no longer needed	A new capability — cloud economics, architecture, provider management — is needed and was not built
Infrastructure becomes invisible, and therefore solved	Infrastructure becomes invisible, and therefore ungoverned

The table is not an argument that the estate should have stayed. It is an argument that the operating model was the thing that needed to move, and it was the one thing the programme was not chartered to touch.

6. The regulator does not recognise the handover

Nowhere is the category error more exposed than in regulated sectors, and the regulatory environment of the last eighteen months has made this impossible to ignore. The invalidation of the Safe Harbor arrangement in late 2015 turned the question of where data physically resides from a footnote into a board-level concern almost overnight. Its replacement is in place but visibly contested, and few practitioners regard the matter as settled. The General Data Protection Regulation, which takes effect next year, raises the stakes of getting data location and processing wrong to a level that concentrates minds. And in the United Kingdom, the vote to leave the European Union has injected a fresh and unwelcome uncertainty into the lawful basis for moving data across borders that, until recently, moved without a second thought.

Set that backdrop against the comforting phrase. An organisation that has told itself infrastructure is someone else's problem discovers, when the regulator writes, that the regulator has never heard of the handover. The accountability for where a customer's data sits, who can access it, and under whose legal jurisdiction it falls did not migrate with the workload. It sits precisely where it always sat: with the accountable executive whose name is on the regulatory return. The provider is a processor acting on instructions; the responsibility for those instructions is stubbornly, legally, non-transferable.

The organisations navigating this well are the ones that never believed the phrase in the first place. They treat the provider relationship as a supervised dependency, not a disposal. They know which workloads can sit anywhere and which are pinned to a jurisdiction. They have someone whose actual job is to hold the provider to account. The ones struggling are those that dissolved exactly that capability because they mistook the exit for an ending.

7. Why the pattern persists

If the category error is this visible, the interesting question — the one an essay owes the reader rather than a quick prescription — is why it recurs so reliably. It is not ignorance. Most of the leaders who run these programmes understand perfectly well, if asked directly, that accountability does not transfer. The pattern persists because several structural forces push in the same direction, and each is individually rational.

The first is the shape of the business case. A capital-avoidance and facility-closure story is legible to a finance committee in a way that an operating-model-transformation story is not. "We will close two data centres and move to variable cost" approves itself. "We will fundamentally change how we architect, fund, and govern infrastructure, and the savings will come later if we also change our behaviour" is a harder sell, so it does not get made. The organisation therefore buys the version of the exit it can approve, not the version that would actually pay off.

The second is that the exit has a natural end date and the operating-model change does not. A migration can be finished; a photograph can be taken. The disciplines that make the new environment economical — right-sizing, reserving capacity, architecting for elasticity, managing the bill — never finish. They are a permanent operating capability, and permanent operating capabilities are chronically under-resourced relative to time-boxed programmes with a launch to celebrate.

The third is the profession's own self-image. For a generation, infrastructure expertise meant knowing the physical estate intimately. When that estate leaves, the instinct is to conclude that the expertise has left with it — that there is simply less to manage. The opposite is true. The new environment demands more judgement, not less: about architecture, about cost, about jurisdiction, about how hard to trust a supplier. But that judgement looks nothing like the old expertise, so the capability gap is invisible until the bill or the regulator makes it visible.

"A migration can be finished. An operating model cannot. We keep funding the thing that ends and starving the thing that never does, and then we are surprised when the benefits behave the same way."

Underneath all three sits the seduction of the phrase itself. "Someone else's problem" is not merely inaccurate; it is comforting, and comfort is precisely what an organisation

two years into a hard, expensive migration wants to hear. The framing survives because it relieves an anxiety, and framings that relieve anxiety are the hardest of all to dislodge.

8. The strongest case against this reading

The most serious objection to everything above is worth meeting head-on, because a reflection that never lets itself be challenged is just a longer assertion. The objection runs like this: what I have described is not a failure of the exit strategy at all, but a failure of execution. Of course lift-and-shift without re-architecting disappoints. Of course an organisation that dissolves its cost and architecture capability overspends. None of that indicts the decision to leave the data centre; it merely says the leaving was done badly. Fix the execution — re-architect the workloads, keep the capability, manage the bill — and the category-error framing dissolves into ordinary programme discipline.

There is real force in this, and I want to concede most of it. The remedy genuinely is to treat the operating model as the thing being transformed. But the objection proves the essay's point rather than refuting it, because it assumes the execution and the framing are independent, and they are not. The reason the execution is so reliably poor is because of the framing. An organisation that believes it is exiting infrastructure will, entirely logically, decline to fund a permanent capability for a thing it thinks it is getting rid of. It will write the business case as a disposal because that is what it believes it is doing. The words are not decoration on top of the execution; they are the instruction set the execution follows. Change "we are exiting" to "we are entering a supervised dependency we must now govern for years," and the funding, the staffing, and the architecture all follow differently. That is why this is worth writing about as a framing and not merely as a checklist of migration good practice.

9. What the exit actually is

So what is the more honest description? The data centre exit is not a disposal. It is the exchange of one relationship with your infrastructure for another — trading a relationship you owned outright, with all its capital burden and visible constraint, for one you rent, mediate through a supplier, and must actively govern precisely because you can no longer see it. The physical estate leaves. The responsibility for cost, architecture, resilience, and jurisdiction does not so much as shift its weight.

Write the business case for what the exit really is: not the closure of a facility but the opening of a dependency you will manage for as long as you run the workloads. Fund the capability that does not end. Keep the people who can hold the supplier to account. Move the operating model, not just the machines. The organisations that do this get the benefits the others were promised.

The reflection that stays with me is how much of this was avoidable with a single change of vocabulary. The organisations that struggled were rarely less capable than

those that thrived; they were the ones that took the comforting phrase literally, prepared for an ending, and were caught unprepared for the beginning that followed. Infrastructure never becomes someone else's problem. At best, it becomes a problem you now solve together with someone else — and the whole art of the exit is knowing which parts of it remained, all along, entirely your own.