Whitepaper

Navigating the Hybrid Cloud Ecosystem in 2026

A Practical Framework for Workload-by-Workload Platform Decisions in an Era of Vendor Disruption and AI-Driven Infrastructure Change

Phil Hawkshaw CTO EMEA
August 2026

Insight®

How to Read This Paper

This paper is structured in three layers, so readers can engage at the depth their situation requires.

Layer	What it covers	For
Executive Summary	The argument in short: the pattern, the framework, the destinations	C-suite, time-limited readers
At a Glance	The traps, the nine assessment dimensions, and the destination options – mapped for orientation	Readers who want a fast map before the detail
Sections 1–6	The full analysis: background, framework criteria, a worked case study, the options landscape, and the path forward	Infrastructure leads, architects, and programme teams doing the work

Most readers will start with the Executive Summary and then move to the section most relevant to their immediate situation. The document is written to support both a fast read and a working reference.

Executive Summary

Every enterprise still running on-premises virtualisation is being pushed toward the same reckoning in 2026: a hybrid cloud and platform reassessment that vendor consolidation, the closing-off of easy infrastructure choices, and the AI-driven need to rethink data & LLM platforms are making unavoidable. The question is no longer whether to revisit virtualisation and infrastructure strategy. It is when, and on what terms.

The Broadcom acquisition of VMware is the clearest and most immediate example of this shift. Licensing consolidation, the elimination of perpetual licences, per-core pricing, and mandatory multi-year subscriptions have driven significant cost increases across the customer base. The many cost-increase multiples circulating in the press and analyst commentary, whether three to five times previous spend or far higher, are worth treating with a pinch of salt: Broadcom has not confirmed these figures, and no single number should be treated as a reliable planning input. What is not in doubt is the direction of travel, which named public sector procurements and the per-core pricing mechanics evidence clearly enough. The partner ecosystem has been restructured. The Omnissa spin-out has created new interdependencies for organisations running Horizon VDI. Regulatory investigations are underway in both the EU and the UK. For VMware

customers, the urgency is real and immediate. But the underlying pattern, vendordriven cost and commercial change forcing a platform reassessment, is not unique to VMware, and organisations on other platforms should expect to face a version of the same question before long.

That was the first phase. Most organisations have now absorbed the initial shock, negotiated or renewed their contracts, and have a clearer picture of what the new commercial reality means for them. What comes next is the operational phase. VCF 9.1.1 is releasing at the end of June 2026, bringing with it significant upgrade pressure across a customer base where fewer than one in twenty organisations that have purchased VCF are actually deploying it. vSphere 8 reaches end of general support in October 2027. The decisions that organisations deferred during the initial shock now need to be made.

This paper is written for organisations across public sector, regulated industries, and commercial environments. The specific pressures vary: a central government department faces different procurement constraints to a financial services firm or a mid-market manufacturer. But the underlying decisions are the same: what to do with the estate an organisation has, on what timeline, and toward which destination.

There is no single right strategy. The right answer depends on the specific workloads in the estate, their dependencies, their technical constraints, and the organisation's strategic direction. This paper presents a workload-by-workload assessment framework covering total cost of ownership, technical debt, technical constraints, application dependencies, platform interdependencies, business roadmap alignment, skills readiness, data security and sovereignty, and AI and data platform readiness. The framework enables disciplined, defensible decision-making.

Multiple valid destinations exist: staying on VMware where the assessment supports it, migrating to Nutanix AHV, adopting Microsoft Azure Local, moving to public cloud, or, most commonly for large enterprises, operating a mixed estate where different workloads sit on different platforms. The framework does not prescribe a destination. It produces the answer that fits the organisation.

What matters is starting: understanding what an organisation has, assessing each workload on its merits, and building a phased programme that moves at the pace of evidence rather than the pace of panic.

At a Glance: The Shape of the Decision

What follows is a fast route through the argument, before the detailed sections build it out in full. Three things to take in: the traps that derail an otherwise sound decision, the nine dimensions that make up a proper workload assessment, and the destinations actually on the table. Each points forward to where the detail lives.

Three ways organisations respond – one works consistently

Extraction. Leave VMware entirely, as fast as possible, driven by cost frustration rather than assessment. Speed without understanding routinely creates problems as expensive as the ones it was meant to solve. (Section 2)

Paralysis. Do nothing and hope the situation resolves itself. Inertia now carries a hard deadline: vSphere 8 end of general support, October 2027. (Section 2)

The disciplined path. A workload-by-workload assessment, paired with rigorous cost-benefit analysis, destination by destination. The approach that works consistently, and what the rest of this paper sets out. (Section 2)

Five destinations, assessed workload by workload

Option	In short
Staying on VMware	For many, the practical near-term task is upgrading to VCF 9.1.1 before vSphere 8 end of support in October 2027, and negotiating better commercial terms while doing it – not necessarily leaving.
Nutanix AHV	An on-premises, operationally familiar alternative with mature migration tooling. Strongest for organisations that want to stay off public cloud.
Microsoft Azure Local	Hyper-V based, tightly integrated with the Microsoft ecosystem. Newer than VMware or Nutanix, so validate through a proof of concept first.
Public cloud	The right destination for genuinely cloud-ready workloads. Not a universal answer once latency, sovereignty, and cost at volume are factored in.
The mixed estate	The most common outcome for large organisations. Different workloads land on different platforms, each assessed on its own terms.

The framework does not prescribe a destination. It produces the answer that fits each workload. Section 4 shows what this looked like for one organisation that went through the full process; Section 5 walks through each option in detail.

Readers who want to go straight to how Insight runs this assessment in practice can find it here

Nine dimensions, one workload at a time

No single dimension gives an organisation the answer. The full criteria for each sit in Section 3 – this is the shape of the assessment.

1 Total Cost of Ownership
Migration, retraining, and partner-change costs, not just the licence.

2 Technical Debt & Migration Complexity
How entangled with VMware-specific features. [Tactical: upgrade cycles add time pressure.]

3 Technical Constraints
Hard platform dependencies that limit the options. [Tactical: GA announcements can shift options.]

4 Application Dependencies
Certified and supported platform requirements.

5 Platform Interdependencies
Adjacent systems: storage, networking, backup, security.

6 Business Roadmap Alignment
Where the organisation is heading, five to ten years out. [Tactical: end-of-life dates create hard deadlines.]

7 Skills & Operational Readiness
Can the team staff and run the new platform day one.

8 Data Security, Sovereignty & Compliance
Classification, residency, and compliance requirements.

9 AI, LLM & Data Platform Readiness
Does this decision enable or constrain AI later.

1. How We Got Here

I have had some variant of the same conversation dozens of times since early 2024. A CxO calls, and within minutes we are talking about VMware. The specifics vary. The underlying question does not: what do we do now?

The scale of this shift is genuinely unusual. This is not a routine licensing renegotiation. It is a structural transformation of how virtualisation infrastructure, the layer that underpins the vast majority of enterprise computing, is commercialised, supported, and controlled. The VMware story matters here as a worked example of a broader pattern: every enterprise still running on-premises virtualisation will face a version of the same vendor-driven reckoning, and what follows is as much a method for that wider challenge as it is an account of Broadcom.

The acquisition and its logic

Broadcom completed its acquisition of VMware in November 2023 for $69 billion. That figure alone says something about the financial model that was always going to follow. No acquirer spends $69 billion to continue running the business the way it was. Broadcom has a well-established pattern with large software acquisitions: consolidate the product portfolio, move customers to subscription licensing, and concentrate investment in the highest-value accounts. It followed this pattern with CA Technologies, and again with Symantec's enterprise division. VMware has followed broadly the same path, and organisations that understand the pattern are better placed to plan around it.

Within weeks of the acquisition closing, Broadcom began reshaping VMware's commercial model. The changes were swift and, for many organisations, shocking in their speed.

What actually changed

Licensing consolidation. VMware's product catalogue, which previously ran to hundreds of individually purchasable SKUs, was collapsed at the time of the acquisition into just three subscription bundles: VMware vSphere Foundation, VMware Cloud Foundation, and vSphere Essentials Plus Kit for smaller environments. Bundle composition, naming, and the small-environment options have all continued to shift since, so that three-bundle structure is best treated as a historical snapshot rather than the current menu: verify the present line-up and an organisation's own entitlements directly with the account team before assuming any of it still applies. The principle, though, has held. An organisation that wanted vSphere took the bundle; one that wanted NSX networking took the more expensive bundle. The à la carte menu was replaced by a set menu, and organisations were paying for courses whether they intended to eat them or not.

Per-core pricing. The old per-socket licensing model, where a two-socket server required two licences regardless of how many cores those sockets contained, was replaced by per-core pricing with a minimum of 16 cores per CPU and 72 cores per server. For organisations running modern processors with high core counts, the arithmetic was brutal, and the 72-core-per-server floor makes the point sharply. A server with two 16-core CPUs has 32 actual cores, and under the old model was licensed simply at two sockets. Under the new model, the 72-core minimum applies regardless of actual core count, so that same server is now licensed at a minimum of 72 cores: more than double what is physically installed. Organisations with higher-density processors found themselves licensing at ratios that bore little resemblance to their previous costs.

Perpetual licences eliminated. There are no longer any new perpetual VMware licences. Every customer must move to multi-year subscriptions. For organisations that had been running on perpetual licences purchased years ago, paying only annual support and maintenance fees, the shift from a relatively modest annual cost to a full subscription fee represented a sharp increase in expenditure.

The cost impact. The numbers vary considerably with the size of the organisation, the products in use, and the terms of existing agreements, but the pattern is consistent. The multiples that circulate in the market, increases of three to five times previous spend, industry-survey ranges of 3x to 12x, the occasional report of more than 10x, are worth taking with a pinch of salt: Broadcom has not confirmed these figures, and treating any single multiplier as a planning input carries real risk. The firmer evidence sits elsewhere. Large public sector procurements have been particularly exposed: several have seen renewal quotes that dramatically exceeded prior spend, prompting public pushback and, in some cases, parliamentary or press scrutiny of the figures involved. In the commercial sector, Insight has worked with mid-market organisations whose VMware spend has moved from a manageable annual line item to a board-level conversation overnight. Whatever the precise numbers, the direction of travel is clear and consistent across the customer base.

Partner and channel disruption

Broadcom also restructured VMware's partner ecosystem substantially. The broad network of authorised resellers and service providers that had grown around VMware over two decades was reduced to a smaller set of strategic partners. For many organisations, particularly those in the public sector who procure through frameworks and established channel relationships, this created immediate practical disruption. Procurement routes that had worked for years simply no longer existed. The partners they had relationships with were, in some cases, no longer authorised to sell VMware. Commercial organisations faced a different version of the same problem: partners they had relied on for managed services or break-fix support were suddenly no longer in the ecosystem, leaving gaps in ongoing operational cover that were not easy to fill quickly.

The Omnissa question

One complexity that catches organisations off guard is the Omnissa spin-out. Broadcom divested VMware's End User Computing division, the products most organisations know as Horizon (virtual desktop infrastructure) and Workspace ONE, into a separate company called Omnissa, backed by private equity firm KKR. For organisations running Horizon VDI, this creates a dependency that is more nuanced than it first appears, and one the assessment framework returns to in detail later in this paper.

Why this is unlikely to reverse

It is tempting to look at the scale of customer dissatisfaction (surveys suggest that between 72% and 92% of VMware customers are actively evaluating alternatives) and assume that Broadcom will be forced to change course. That is not a safe assumption. Broadcom's Infrastructure Software segment, which is substantially comprised of VMware products and services, generated $21.5 billion in revenue in fiscal year 2024: a 270% increase on the prior year. The annualised booking value from VMware customers reached approximately $15 billion, with over 4,500 of Broadcom's largest customers converted to VCF subscriptions. From Broadcom's perspective, the strategy is working precisely as intended. The customers who leave are, in their model, the customers they were never going to invest in retaining.

This is not a comfortable conclusion, but it is an honest one. The changes are structural, they are deliberate, and they are delivering the financial outcomes Broadcom's shareholders expect. Organisations need to plan on the basis that the new reality is, in fact, the reality.

Broadcom's Infrastructure Software segment, which is substantially comprised of VMware products and services, generated

$21.5billion

in revenue in fiscal year 2024:

a 270% increase

on the prior year.

That was the first phase. The commercial restructuring. What comes next is the operational phase: the upgrade pressure, the adoption gap, and the decisions every organisation now needs to make.

Where we are now

With the initial commercial shock largely absorbed, a new set of pressures is emerging. The immediate question for most organisations is no longer what VMware will cost. It is what they are going to do with what they have bought.

VCF 9.1.1 and the upgrade mandate. VCF 9.1 had significant issues and Broadcom will not be back-patching older versions. VCF 9.1.1 is releasing at the end of June 2026 and it carries a clear message from Broadcom: upgrade first, then adopt. Get onto VCF 9.1.1, and then work toward fully adopting VCF capabilities over time. With around 250,000 customers needing to move, the scale of the upgrade effort across the industry is substantial.

vSphere 8 end of general support: October 2027. This is a hard deadline. Any organisation still running vSphere 8 after October 2027 will lose general support entirely. That is not a distant planning assumption: it is sixteen months away, and for organisations with large, complex estates, sixteen months is not a comfortable margin.

The adoption gap. Fewer than 5% of organisations that have purchased VCF are actually deploying it. The rest have bought it, in many cases under considerable commercial pressure, and are still running legacy vSphere configurations. This suggests that the cost increases of the past two years have not, in most cases, been matched by a change in how the infrastructure is actually used. Organisations are paying VCF prices for vSphere behaviour. That gap is not just an adoption curiosity, it is wasted spend in its own right: licensing capability that sits unused is a direct hit to total cost of ownership, on top of whatever the underlying price increase already cost.

A more pragmatic definition of deployment. Broadcom has acknowledged the adoption gap and is responding. Their current position is that vSphere 9.1 plus vCenter 9.1 plus VCF Ops (with SDDC Manager now integrated into VCF Ops) constitutes a VCF deployment. No vSAN required. No NSX required. Most customers will remain on three-tier architecture. This is a meaningful shift from the original VCF proposition, and it opens up a more accessible upgrade path for organisations that are not ready to fully transform their networking and storage.

Investment in the upgrade path. Broadcom has a dedicated team investing in Glasswing, their migration and upgrade tooling. Their focus is on identifying and fixing issues in the upgrade path continuously. This is a signal that Broadcom is investing in making the transition smoother, though the scale of the task means that complexity will remain for some time.

Make use of the professional services already in the contract. Many VMware contracts now include Broadcom professional services entitlements, such as Service Entitlement allocations, that cover migration and upgrade support. These go unused far more often than they should, usually because the organisation does not realise they are there or has not built them into the project plan. Checking what is included before resourcing an upgrade internally, or before bringing in a partner to cover the gap, is a straightforward way to offset the burden of moving onto VCF 9.1.1. It is, in effect, capability the organisation has already paid for.

The picture that emerges is one of a market that is moving from disruption into execution. The commercial terms are largely set. The deadlines are known. What remains is the work: understanding where each organisation sits against the upgrade timeline, closing the gap between licences purchased and capabilities deployed, and making clear-eyed decisions about which workloads should stay on VMware and which should move elsewhere. The rest of this paper is about how to do that well.

2. The Temptation of Simple Answers

We are now in the next phase of this disruption. The initial shock of the acquisition has settled into a more persistent pressure: VCF 9.1.1 arriving at the end of June 2026, vSphere 8 end of general support set for October 2027, and 250,000 customers needing to move. The urgency is no longer abstract. It has dates attached to it. And if anything, that sharpens the temptation to reach for a simple answer.

When the scale of a licensing change lands on a CxO's desk, particularly one measured in multiples of previous spend, the pressure to act decisively is enormous. Boards want answers. Finance directors want forecasts. Procurement teams want alternatives. And the market obliges with no shortage of confident voices telling organisations exactly what to do.

In the two years since the acquisition, organisations have responded in broadly three ways. Two of them produce, at best, mixed results. The third is more disciplined, and the difference in outcomes has been striking.

The extraction trap

The first response is best described as panic extraction: the decision to leave VMware entirely, as fast as possible, driven primarily by the desire to escape the new cost structure. The impulse is understandable. When a virtualisation bill has tripled or worse, the emotional logic of "we need to get off this platform" is powerful.

But migration at speed, driven by cost pressure rather than strategic assessment, routinely creates problems that are at least as expensive as the ones it was meant to solve.

This plays out in practice. I have seen an organisation commit to a full VMware exit on an aggressive timeline, only to discover mid-programme that critical line-of-business applications had undocumented dependencies on VMware-specific features: vSAN storage policies, NSX micro-segmentation rules, distributed switch configurations that had been built up over years by engineers who had long since moved on. The confidence of the initial board paper (clear savings, clean timeline, new platform) collided with the complexity of what actually existed in production. The programme slowed, costs escalated, and the organisation found itself running two platforms in parallel for far longer than anyone had planned or budgeted for.

This is not an unusual story. It is a common outcome for organisations that treat VMware migration as a procurement problem rather than an infrastructure transformation programme. And the risk is not confined to large public sector estates: a mid-market commercial organisation with 200 VMs can face exactly the same problem, because undocumented dependencies do not care about the size of the organisation. They surface mid-migration regardless. Migration complexity is routinely underestimated, new platforms carry their own maturity risks and learning curves, and the team that was already stretched thin managing the existing estate is now expected to run the old environment, build the new one, and execute migrations all at once, often with skills they are still developing.

None of this means migration is wrong. It means migration done hastily, without a thorough understanding of what is actually being moved and what it depends on, is a recipe for a different kind of cost overrun.

A variant of this failure mode is worth naming: organisations that leave on principle, driven by frustration with Broadcom's commercial behaviour rather than by a workload-level assessment of what they are actually migrating away from. The frustration, in many cases, is entirely reasonable. But emotion is not a substitute for evidence, and a decision to leave that cannot survive a TCO test is in the same category as panic extraction, just with a different emotional driver.

There is also a calculus that has shifted in ways the frustration can obscure. Geopolitical uncertainty is pushing many organisations toward greater control over where their data physically sits, and AI workloads are raising the bar on the security and data protection a platform needs to provide. Both of those trends can genuinely strengthen the case for staying on a mature, well-supported on-premises platform. Leaving without a proper assessment can mean missing a platform that has genuinely strengthened.

The paralysis trap

The second response is strategic paralysis: the decision, often unstated, to do nothing and hope the situation resolves itself. Perhaps Broadcom will moderate its pricing. Perhaps a regulatory intervention will force change. Perhaps the renewal deadline is far enough away that this can be next year's problem.

This too is understandable, particularly for organisations with large, complex estates where the sheer scale of any change programme feels overwhelming. But inertia has its own cost, and it compounds. More importantly, paralysis now has a hard deadline. vSphere

8 end of general support in October 2027 means that organisations still running legacy configurations will not simply face penalty uplifts: they will be running unsupported infrastructure. That is a different category of risk entirely.

Broadcom's licensing model is designed to reward early commitment and penalise delay. Organisations that miss renewal windows or allow agreements to lapse face penalty uplifts, reported at around 20% in some cases. Meanwhile, the competitive alternatives landscape is evolving rapidly. Nutanix is investing heavily in VMware migration tooling and training. Microsoft is offering aggressive incentives for Azure Local adoption. The organisations making decisions now are building institutional knowledge and operational capability on new platforms. The organisations standing still are not, and when they eventually do move, they will be starting from further behind, with less leverage and fewer options.

Some organisations have gone further than delay and simply let their support agreements lapse entirely, with no active renewal in motion. Where that describes the estate, it is worth naming the position honestly: "do nothing" is not a sustainable middle ground once support has actually expired. There is a genuine fork in the road. One path is to catch up: renew, absorb the penalty uplift, and bring the estate back under support. The other is to make a deliberate decision to isolate and ringfence the unsupported estate from the rest of the infrastructure, accepting the risk of running without vendor support on a contained, well-understood set of workloads while the rest of the environment moves forward. Both are legitimate strategies. What is not legitimate, in Insight's experience, is treating an expired agreement as a problem that has somehow resolved itself simply because nothing has broken yet.

The disciplined path

The third response, and the one that works most consistently, is neither flight nor freeze. It is the approach taken by organisations that have resisted the pressure to make a single sweeping decision, done the work to understand what they actually have, and then made deliberate choices about what to leave, what to move, and what to transform.

Organisations of all sizes have done this well. What they have in common is a structured methodology: a workload-by-workload assessment that examines each part of the estate on its own terms. That means looking at cost profile, technical dependencies, application requirements, and relationship to the business roadmap. It is, in spirit, analogous to the 6 R's of cloud migration (retain, retire, rehost, replatform, refactor, repurchase): a framework that refuses the temptation of a blanket answer and instead asks the right question for each workload. That assessment is then paired with a rigorous cost-benefit analysis, destination by destination, so that the choice of platform is driven by evidence rather than vendor enthusiasm.

The result, consistently, is a strategy that is both defensible and deliverable. Not a single heroic exit. Not a policy of hope. A phased programme, grounded in a clear-eyed view of the estate, that moves at the pace of evidence. This approach works for commercial organisations as much as it does for public sector: the framework is sector-agnostic.

One honest caveat is worth adding here: not all VMware environments are complex, and not all answers need to be elaborate. For organisations with relatively simple, standardised estates (a modest number of workloads, limited customisation, straightforward dependencies) a full migration to a single alternative platform, or indeed a decision to stay on VMware and negotiate the new commercial terms, may be entirely the right answer. The framework in this paper is designed to handle complexity. Simpler situations often have simpler answers, and it would be a mistake to impose analytical overhead where it is not warranted. The discipline lies in knowing which situation an organisation is actually in before it commits.

The regulatory question

Some organisations are holding out hope for regulatory relief. And there are genuine investigations underway. The European Commission opened a formal antitrust investigation into Broadcom's bundling practices in February 2025. In the UK, the Competition and Markets Authority launched a compliance investigation in March 2025, specifically examining whether Broadcom has breached the commitments it made when the acquisition was cleared: commitments around customer portability, interoperability, and anti-bundling. The CMA has gone so far as to issue a Section 174A notice to Broadcom, with potential penalties of up to 5% of global annual turnover for any breach found.

These are meaningful actions, and they may ultimately bring some change to Broadcom's commercial practices. But regulatory processes operate on their own timelines, measured in years rather than quarters. No organisation facing a VMware renewal in the next twelve to eighteen months can realistically build a strategy around the assumption that a regulator will resolve their commercial problem before their contract expires. Regulatory outcomes are a factor to monitor, not a strategy to depend on.

The noise in between

Between the urge to extract and the temptation to freeze sits an enormous amount of vendor noise. Every alternative platform provider is, quite naturally, marketing hard into the disruption. “Migrate to us” is the message from every direction, and each comes with its own set of favourable benchmarks, migration tools, and customer stories. That is marketing, and it is doing exactly what marketing is supposed to do. But it is not strategy. Strategy requires understanding the estate, its dependencies, its constraints, and the organisation’s own business direction, and making decisions on that basis rather than on the basis of a vendor’s slide deck.

What is actually needed

What sets apart the organisations that navigate this well is less dramatic, but far more effective than either panic or paralysis. They resist the pressure to make a single, sweeping decision. They accept that the answer is not simple, because the reality is not. And they commit to a disciplined, workloadspecific approach that examines what they actually have before deciding where it should go.

The first phase of this disruption was about shock and response. This next phase is about structured decision-making. That third path, the disciplined one, is what the rest of this paper is about. So if the answer is not “get out now” and not “wait and see,” what is it? It starts with understanding what an organisation actually has: workload by workload, dependency by dependency.

3. A Framework for Workload-by-Workload Decision-Making

Almost every organisation believes, at the outset, that it understands its VMware estate. Some have asset registers or configuration management databases, or both. They have architecture diagrams that were accurate at some point in the past. And yet, within weeks of starting a structured assessment, they discover dependencies they did not know existed, integrations that were never documented, and workloads whose relationship with the hypervisor layer is far more intimate than anyone assumed.

That discovery process is often not a failure of governance, but a reflection of how deeply VMware has woven itself into enterprise infrastructure over two decades. The hypervisor is not just a platform: it is the substrate on which everything else sits. Changing it, or deciding whether to change it, requires a structured way of thinking that accounts for the full picture.

What follows is not a vendor comparison matrix. It is a framework for examining each workload on its own merits, understanding its constraints, and reaching a decision that is defensible precisely because it has considered every dimension that matters. Nine dimensions, each illuminating something the others miss. The framework applies whether the organisation is a government department, a financial services firm, or a commercial enterprise. The dimensions are the same; the weighting will differ.

3.1 Total Cost of Ownership

The licensing invoice is the number that gets attention. It is rarely the number that determines the real cost.

Total cost of ownership has to be thought through well beyond the subscription fee. What does it actually cost to move a workload from one platform to another?

The migration itself – the engineering hours, the testing cycles, the inevitable rework when something does not behave as expected on the new platform.

The dual-running period – paying for both the old environment and the new one at once, sometimes for months.

Retraining – the operations team has to learn new tooling, new troubleshooting approaches, new operational rhythms.

Opportunity cost – every engineer working on a migration is an engineer not working on something that moves the business forward.

Organisations regularly fixate on a licensing comparison that shows a 40% saving on the new platform, only to discover that the total programme cost (migration, dual-running, retraining, and lost productivity) exceeded three years of the higher VMware subscription. That does not mean the migration was wrong. It may still have been the right strategic decision for other reasons. But the economics need to be honest, and they need to account for the full picture.

One further cost component is almost always missing from the initial analysis: the cost of changing support partners. When the destination technology stack spans on-premises infrastructure, edge locations, hybrid environments, and public cloud, as it increasingly does, organisations find that their existing partners cannot credibly cover all of those destinations. Not for want of capability, but because the breadth of skills required across a genuinely mixed estate is significant, and not every partner has invested equally across every platform. There is also the question of how well partners work alongside each other: some suit collaborative engagement models, others less so, which adds an overhead to partner management in its own right. The cost of changing partners, or of managing several across a mixed estate (the commercial negotiation, the transition, the overlap period, the relationship investment) is a real TCO component, and one of the most consistently overlooked in the initial analysis.

The cheapest licence is not the cheapest outcome.

3.2 Technical Debt and Migration Complexity

Not all workloads are equal. Some are straightforward virtual machines running a standard application on a standard operating system; they could run on almost any hypervisor with minimal effort. Others are deeply entangled with VMware-specific capabilities: vSAN for storage, NSX for networking, DRS for workload balancing, custom PowerCLI scripts that have accumulated over years of operational refinement. The question for each workload is: how much of what it depends on is VMware-specific, and how much of that dependency is explicit versus hidden? The explicit dependencies are the ones in the architecture documents. The hidden ones are the custom configurations that an engineer implemented five years ago, that nobody documented, and that will surface as a production incident three weeks into the migration.

Technical debt is not a moral judgement. It is a practical reality. Organisations that have run VMware for fifteen or twenty years have built up layers of configuration, integration, and customisation that represent genuine engineering investment. Understanding that investment, and what it costs to replicate or replace, is a prerequisite for any honest migration plan.

There is a further, more specific case worth distinguishing from technical debt built up through customisation: workloads running on unsupported operating systems, ageing middleware, or software stacks that have gone years without an update. These are not simply uncertified on other platforms. They are often genuinely harder to run anywhere else, VMware included, because the constraint sits below the hypervisor layer and no migration removes it. This is, in practice, the estate that tends to stay on VMware by default: not because anyone made a considered case for it, but because there was no credible alternative on the table. Naming that explicitly is worth doing, because it is one of the more concrete, defensible reasons for retaining a VMware, or equivalent on-premises, presence for at least part of the estate.

3.3 Technical Constraints

Some constraints are not about cost or complexity; they are hard architectural limits that no amount of clever engineering can work around. These are the dependencies that should be identified early, because discovering them mid-programme is where migration projects come to grief.

A good example of this is virtual desktop infrastructure, and it is also a good example of how quickly a hard constraint can stop being hard. Organisations running Omnissa Horizon, the VDI platform that was spun out of VMware following the Broadcom acquisition, have, until recently, been tied to VMware's hypervisor for that workload. Omnissa committed publicly to becoming hypervisoragnostic, and as of December 2025, that commitment has started to land in production: Omnissa Horizon 8 on Nutanix AHV reached general availability, with Omnissa publishing formal design guidance for the combination. Support for Microsoft Azure Local is reported to still be in active development rather than generally available, so that leg of the hypervisor-agnostic roadmap is worth checking directly with Omnissa before planning against it.

This does not mean the constraint has disappeared; it means the constraint has moved, and that is precisely the point: assess where things stand today, do not plan against where they stood a year ago. A workload that looked locked to VMware in early 2025 may have a genuine, supported alternative now. An organisation still plans around the dependency, and may still choose to keep a given workload on VMware while migrating others. What it should not do is treat last year's constraint as this year's fact without checking.

There is a commercial dimension worth noting alongside the technical one. Organisations with a VDI estate can also buy VMware vSphere Foundation for desktop directly through Omnissa, considerably cheaper than buying it through Broadcom, and available on a one-year term rather than the multi-year commitments increasingly standard elsewhere. Worth knowing that this route currently only supports vSphere 8, not vSphere 9, so it suits organisations not yet ready to move off the older version rather than those already planning a vSphere 9 upgrade.

This is a constraint that continues to evolve. Omnissa's multi-hypervisor roadmap is progressing faster on some platforms than others, and the timeline for production-ready support across the remaining alternative hypervisors is worth monitoring actively, not assuming.

Horizon is one illustration of a broader principle: any workload that depends on a capability unique to a specific hypervisor needs to be identified early and assessed on its own terms.

3.4 Application Dependencies

The applications running on the virtualisation layer have their own certification requirements, support matrices, and performance characteristics. A platform that works technically may not be supported by the application vendor, and that distinction matters more than it might seem.

This plays out with enterprise applications from major software vendors. The application will run on the new hypervisor. It passes functional testing. But the vendor's support matrix does not include that hypervisor, which means that when an organisation raises a support ticket, the first thing the vendor will say is "move it to a supported platform and reproduce the issue." Unsupported does not mean broken. But it does mean the organisation is on its own if something goes wrong, and for business-critical applications, that is a risk that warrants explicit acceptance at the right level of the organisation.

The assessment here is workload by workload: which applications have hard certification requirements, which vendors are actively expanding their support matrices to include new hypervisors, and which applications are approaching end-of-life anyway and might be better candidates for modernisation rather than migration.

One area where partner experience proves genuinely valuable is in bridging the gap between what a vendor formally certifies and what actually works in practice. Vendors are often slow to update their support matrices; the commercial and testing overhead of certification means that real-world compatibility frequently runs ahead of documented support. A partner with hands-on experience running a specific application on a specific architecture can often provide the confidence that a

support matrix alone cannot. That experience does not replace vendor certification, but in a landscape where not every relevant hypervisor appears on every support matrix, it can be the difference between a wellinformed risk decision and a leap in the dark.

3.5 Platform Interdependencies

Organisations routinely scope a hypervisor migration and then discover, mid-programme, that the real scope is three times larger because of the adjacent systems. No workload exists in isolation. The hypervisor sits at the centre of an ecosystem: storage, networking, backup, monitoring, disaster recovery, security tooling, automation frameworks. Each of those layers has its own relationship with the hypervisor, and changing the hypervisor can cascade into dozens of adjacent decisions.

Consider backup. The current backup solution almost certainly has deep integration with VMware's APIs: vStorage APIs for Data Protection, Changed Block Tracking, snapshot management. Moving to a different hypervisor means validating that the backup vendor supports the new platform with equivalent functionality, or selecting a new backup solution entirely. The same applies to monitoring (does the observability platform support the new hypervisor's metrics?), networking (if NSX is in use, what replaces it?), and security (how does the endpoint protection interact with the new hypervisor layer?).

The point of raising this is not to make migration sound impossible, because it is not, but to ensure that the scope of the decision is properly understood. A hypervisor migration is rarely just a hypervisor migration. It is a platform migration, and the planning should reflect that reality.

There is a further interdependency that is easy to overlook: the gap between purchasing VCF and actually deploying it. The majority of organisations that have bought VCF licences are still running legacy vSphere configurations. They have purchased the future but not yet built it, and are paying for a platform they are not using while the platform they are using moves closer to end of support. Through the total cost of ownership lens, this is not a neutral holding pattern: it is money spent on capability sitting idle, on top of whatever the organisation is still paying to run the legacy estate it was meant to replace. Closing that adoption gap is, for many organisations, the most immediate infrastructure challenge they face, and the most immediate way to stop paying twice.

3.6 Business Roadmap Alignment

Technology decisions should serve business direction, not the other way around. This sounds obvious, but it is remarkable how often the conversation about VMware alternatives proceeds entirely on technical and commercial grounds without reference to where the organisation is actually heading.

If the business is on a trajectory toward cloud-native architectures (containerised workloads, platformas-a-service, serverless) then investing heavily in a new on-premises hypervisor may be building a bridge, not a destination. That is fine, as long as it is acknowledged. A bridge investment and a destination investment carry different expectations for return, different acceptable lifespans, and different levels of acceptable compromise.

Conversely, if the organisation has a ten-year commitment to on-premises infrastructure, whether for regulatory reasons, performance requirements, or strategic choice, then the hypervisor decision is a long-term commitment that warrants a different level of scrutiny. The question is not only "which platform is technically best?" but "which platform serves where this organisation is going over the next five to ten years?"

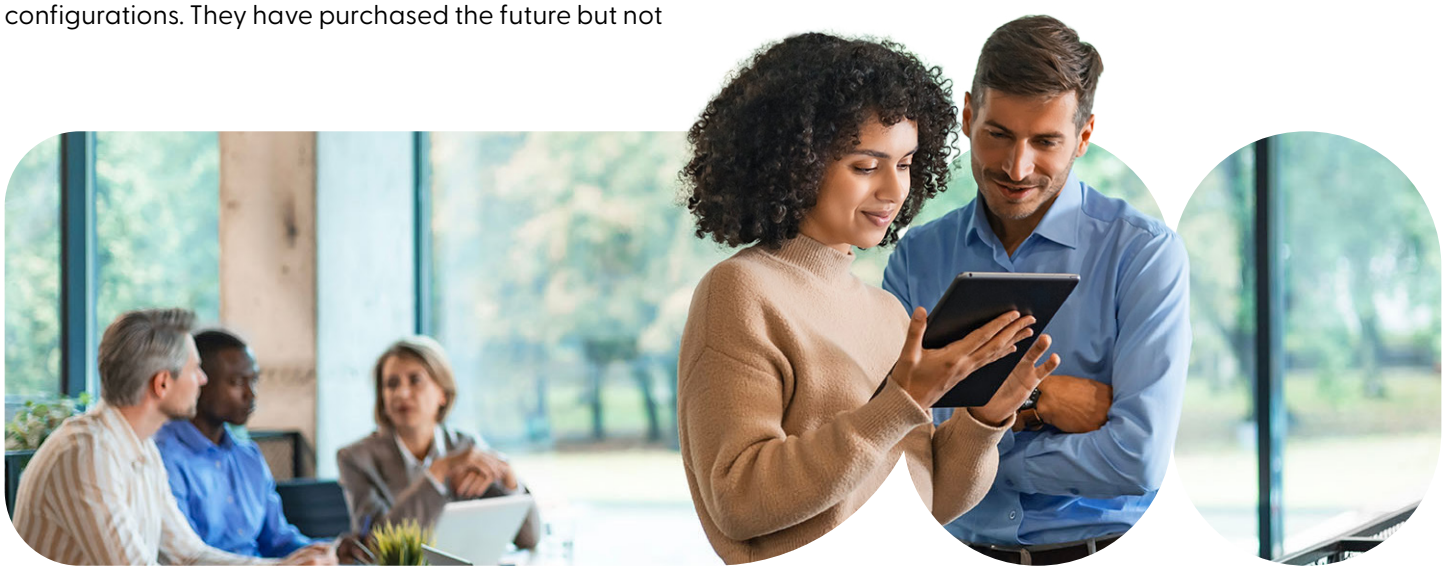

3.7 Skills and Operational Readiness

This dimension does not always make it into the assessment frameworks, but it is often the factor that determines whether a migration succeeds or stalls.

VMware skills are deeply embedded in most enterprise IT teams: two decades of accumulated expertise. Infrastructure engineers know how to troubleshoot vSphere at 3am. They know the quirks, the workarounds, the operational patterns that keep the environment stable. That institutional knowledge has enormous value, and it does not transfer automatically to a new platform.

The alternative platforms, Nutanix AHV, Azure Local, even Hyper-V in its more recent incarnations, have smaller talent pools. Industry data consistently shows that organisations are struggling to hire professionals with skills in these newer platforms. Nutanix has invested in expanding its training programmes, and Microsoft has launched dedicated Azure Local certification pathways, but the supply of experienced practitioners is still catching up with demand. One survey found that over 75% of IT teams reported difficulty finding professionals with virtualisation skills generally, and that figure tends to be worse for the newer platforms specifically.

This is not a reason to avoid migration. It is a reason to plan for workforce readiness as a first-order concern, not an afterthought. Where a migration timeline assumes the organisation can hire or retrain at the pace the project demands, that assumption should be tested early. Skills availability often determines the realistic migration timeline more than any technical factor.

The skills challenge extends to the partner ecosystem too. Broadcom has introduced rigorous certification requirements for partners delivering VCF services, and these certifications are neither quick nor easy to obtain. Even large, well-resourced partners are investing significant time and effort to build certified capacity. For organisations planning a VCF upgrade or adoption programme, it is worth understanding not just the internal team's readiness but their partners' certification status and capacity. The demand for certified practitioners is currently outstripping supply, and that dynamic is unlikely to resolve quickly.

There is a practical funding wrinkle here too. Broadcom provides Service Entitlement, typically around 10% of contract value, to help clients adopt VCF. Clients struggle to take full advantage of it, not because the entitlement is poorly designed, but because the internal teams who would drive adoption are already stretched thin keeping operations running. The entitlement then sits idle while the team that would use it firefights elsewhere, and the funding window is not open-ended, typically around one year after renewal. Lose track of where the estate sits against that window and the entitlement can lapse before it has been meaningfully used. Where Service Entitlement is part of the commercial position, treat internal capacity to use it as a planning item in its own right, not an assumption.

It is also worth recognising that partners can meaningfully augment an internal IT team's skills base. In the public sector this can be more than a convenience: many government and regulated sector organisations operate in environments where the right security accreditations matter, and where having appropriately cleared staff with the relevant certifications is a practical requirement for certain types of engagement. The same principle applies in financial services, healthcare, and any sector handling sensitive customer data, where accreditations and compliance credentials are equally negotiable. A partner that can bring those cleared, accredited resources to bear, particularly during the transition period when the internal team is still building capability on the new platform, can materially reduce the risk profile of a migration programme. That is worth factoring into both partner selection and programme planning.

3.8 Data Security, Sovereignty, & Compliance

For public sector organisations and regulated industries, the platform decision carries an additional dimension that can, in some cases, override all the others: does the platform meet the organisation's security classification, data residency, and compliance requirements?

In the UK public sector context, the Government Security Classification framework's three defined tiers (OFFICIAL, SECRET, and TOP SECRET) each carry different constraints on where data can be processed and stored. Most government workloads sit at OFFICIAL, which permits public cloud deployment with appropriate controls. But workloads at higher classifications require enhanced controls, and in some cases mandate that data remains within specific UK-sited infrastructure with specific assurance models.

This matters for the VMware migration decision in a practical way. An on-premises solution like Azure Local keeps data on hardware the organisation controls, within facilities it manages, while still providing the management and operational tooling of a modern platform. A public cloud solution offers different capabilities but introduces questions about data residency, the jurisdiction of the cloud provider, and the adequacy of the provider's security controls for the specific classification level. Even the management plane, the control layer that orchestrates the infrastructure, should be assessed against the organisation's security requirements.

The same logic applies across regulated industries, though the frameworks differ. Financial services firms operate under FCA and PRA expectations around operational resilience, and those expectations extend to the infrastructure on which critical systems run. Where a platform introduces new dependencies on third-party controlled systems, that needs to be assessed against the organisation's operational resilience obligations. Healthcare organisations handling patient data must satisfy NHS DSPT requirements and the broader obligations around patient data protection. Any organisation holding significant volumes of personal data has GDPR data residency considerations to work through, and those are not confined to the public sector.

There is also a category of risk that applies to any organisation handling commercially sensitive data: proprietary intellectual property, financial transactions, customer records. The platform decision affects where that data sits, who can access the management layer, and what assurances the organisation can give to customers, regulators, and boards. Data sovereignty is not exclusively a public sector concern.

For some organisations this dimension is the most important of the nine. For others it is a box to tick. The framework does not assume which; it asks the question and lets the answer shape the decision.

3.9 AI, LLM and Data Platform Readiness

This dimension sits here deliberately, alongside the other eight rather than as a closing reflection. In conversation after conversation about VMware platform decisions, AI comes up at the end, almost as an afterthought, once the cost, the technical debt, and the migration logistics have been worked through. That ordering gets it backwards. The platform decisions being made today, about where workloads run and how the underlying infrastructure is architected, will shape what an organisation can realistically do with AI over the next five to ten years. That is a present consideration, not a future one, and it belongs in the same assessment as the other eight dimensions, not bolted on afterwards.

The core question is where AI workloads will run, and the honest answer for most organisations is not a single location. Scaling AI use cases into production typically requires some mix of edge, onpremises, hybrid, and public cloud infrastructure, tuned for cost, latency, data gravity, and the regulatory or sovereignty constraints that apply to the data involved. Data management sits at the centre of this. Where the organisation's data lives, how it is governed, and how portable it is across platforms will shape the AI roadmap in ways that are difficult to unpick later, once models and pipelines have been built around a particular platform's assumptions. A platform choice that looks purely like a virtualisation decision today can quietly become an AI architecture decision in twelve months' time.

One specific facet of this deserves calling out on its own terms: how an organisation provisions LLM capacity, not just where. Metered access to LLMs, consumption priced per token through a public API, is the fastest way to start and suits unpredictable or exploratory workloads well. But at production scale, metered consumption becomes a cost line that grows with usage in a way that is hard to forecast and, for some organisations, hard to govern: every prompt and completion passes through a third party's infrastructure, which raises the same data residency and sovereignty questions already covered in Section 3.8. The alternative is unmetered, or meterless, provision: dedicated GPU capacity, owned or reserved outright, that carries a fixed cost regardless of how heavily it is used, and that keeps inference on infrastructure the organisation controls. This is a genuine platform decision, not a procurement detail bolted on afterwards. A hypervisor or platform choice that cannot host dedicated inference capacity at the scale and cost profile an organisation needs is, in effect, committing that organisation to metered consumption by default, whether or not that was ever a deliberate choice.

Every platform in this paper's landscape has its own answer to this question, and it is worth evaluating each on its own terms rather than assuming any one vendor has solved it. Broadcom's answer, through the VMware portfolio, is Private AI Foundation, built on VKS (VMware Kubernetes Service), which represents its bet on enterprise AI infrastructure for organisations staying within the VMware ecosystem. It is a credible example of how an incumbent platform vendor is positioning for this shift, but it is one example among several, and organisations evaluating Nutanix, Azure Local, public cloud, or a mixed estate should expect, and ask for, an equally clear answer to the same question from each of those options. The practical implication is simple: when an organisation assesses a workload against this framework, it should ask not only where it runs today and what it would cost to move, but where the organisation needs to be in three years' time to take full advantage of AI in production, and whether the platform decision in front of it helps or hinders that destination. The cheapest migration is not the one that limits AI options later.

Summary: The Nine Dimensions at a Glance

Dimension	Core Question	Why It Matters
Total Cost of Ownership	What is the full cost: licensing, migration, dual-running, retraining, opportunity cost, and partner change costs?	The cheapest licence is not the cheapest outcome. Partner transition costs are a frequently overlooked component.
Technical Debt & Migration Complexity	How entangled is this workloadwith VMware-specific features and configurations?	Determines migration effort, risk, and realistic timeline.
Technical Constraints	Are there hard platform dependencies that limit the options?	Some workloads are currently tied to a specific hypervisor. Plan around them, not through them.
Application Dependencies	What do the applications require in terms of certified and supported platforms?	Unsupported does not mean broken, but it does mean uncharted territory. Partner experience can bridge the gaps.
Platform Interdependencies	What adjacent systems (storage, networking, backup, security) are affected?	A hypervisor migration is rarely just a hypervisor migration.
Business Roadmap Alignment	Where is the organisation heading strategically over the next five to ten years?	A bridge investment and a destination investment require different decisions.
Skills & Operational Readiness	Can the organisation staff and operate the new platform from day one?	Often the binding constraint on migration timelines; particularly relevant where security accreditations matter.
Data Security, Sovereignty & Compliance	Does the platform meet security classification, data residency, and compliance requirements?	Non-negotiable for public sector; increasingly critical for regulated industries and any organisation handling sensitive data.
AI, LLM & Data Platform Readiness	Does this platform decision enable or constrain AI and LLM deployment over the next five to ten years?	The cheapest migration is not the one that limits AI options.

No single dimension gives an organisation the answer. Skipping any one of them creates a blind spot. The organisations that navigate this well are the ones that work through all nine, workload by workload, and accept that different workloads will produce different answers.

These nine dimensions will not always agree, and the framework does not pretend otherwise. Total cost of ownership might favour migration while skills availability favours staying; business roadmap alignment might point toward public cloud while data sovereignty rules it out for a specific workload. When that happens, it helps to sort the nine into two kinds of question rather than treat them as nine equal votes. Technical constraints and data security, sovereignty and compliance behave as gates: where either genuinely rules an option out, no amount of TCO or skills advantage elsewhere reopens it. The remaining dimensions, cost, technical debt, application dependencies, platform interdependencies, business roadmap, skills, and AI readiness, are trade-offs to be weighed against each other, and the weighting is a judgement call that belongs to the organisation, informed by where it sits on each one. The framework does not make that judgement for an organisation. It makes sure the judgement is made with the full picture in view, rather than on whichever dimension happened to get raised first in the room.

One thing this framework deliberately does not do is prescribe how to run the discovery and dependency-mapping work itself. That is execution detail, not strategy, and it varies too much by estate, tooling, and existing documentation to generalise usefully here. Insight's workload assessment methodology page sets out how Insight runs that work in practice, for readers who want to go a level deeper than this paper is designed to.

Frameworks are useful. But they become real when seen applied. The next section shows what this looked like for one organisation that went through exactly this process.

4. What This Looks Like in Practice

In this next phase of the VMware landscape, where the initial shock has given way to the harder work of structured decision-making, here is what it looks like when an organisation commits to the disciplined path. Earlier this year, I was involved with an organisation that needed to make a VMware platform decision under exactly the kind of pressure described in the sections above. Their VMware licensing costs had increased substantially following the Broadcom changes, and doing nothing was not a viable option. But they were also clear-eyed about the risks of rushing: they had seen other organisations stumble by committing to a migration before properly understanding what they were migrating from or what they were migrating to.

What made this engagement instructive was not just the outcome, but the discipline of the process.

Stage 1: Assessment and Decision

The first stage was about understanding, not deciding. The team, a joint effort between the organisation's own infrastructure architects and Insight's consulting engineers, worked through the assessment framework systematically. They mapped every workload in their VMware estate. They catalogued dependencies: the explicit ones and, inevitably, quite a few that nobody had documented. They assessed each workload against the nine dimensions I have described: cost, technical debt, constraints, application requirements, interdependencies, business direction, skills, security, and AI and data platform readiness.

What I find interesting, looking back, is that this stage was never purely internal. Assessment work at this scale rarely is. Hardware vendors and hypervisor platform vendors both have a legitimate role in an assessment process: they understand their own technologies in ways that no third party fully can, and engaging them early tends to surface constraints and opportunities that would otherwise take longer to identify. At this stage, the destination had not been decided. We were talking to vendors as part of the process of forming a view, not validating a conclusion we had already reached. That distinction matters. Vendors who are engaged as contributors to an assessment, rather than as salespeople for a predefined destination, tend to provide more useful input, and the organisation retains its own independent judgement throughout.

What emerged from the analysis was not a foregone conclusion. Several platform options were genuinely viable. The decision to pursue on-premises infrastructure using Azure Local on Dell hardware was driven by a combination of factors that came out of that analysis: the organisation had a significant existing investment in the Microsoft ecosystem, which meant operational familiarity and licensing synergies. Azure Local met their technical requirements for on-premises workloads while providing a management model aligned with their strategic direction toward Azure services more broadly. And Dell's hardware platform offered the performance and supportability characteristics their engineering team needed.

The important point is that the destination was not predetermined. It emerged from the assessment. Had the workload profile or the organisation's strategic direction been different, the answer could just as easily have been Nutanix, or staying on VMware and optimising, or a public cloud migration. The framework produces the answer that fits the organisation, not the answer that fits the vendor.

Stage 2: Production Proof of Concept

Here is where honesty about risk becomes important. Azure Local was, and at the time still is, a relatively new solution. The team identified platform maturity as a genuine risk in the risk register: not as a footnote, but as a high-probability, high-impact item that required active management. This was not a weakness in the decision-making process. It was evidence of disciplined decision-making. The organisation knew what it did not yet know, and it designed a programme structure that accounted for that uncertainty.

Rather than committing to a full-scale migration on the basis of analysis alone, they ran a production proof of concept. Not a lab test. Not a synthetic benchmark. A production POC, with real workloads, in their actual environment, testing whether the solution worked under the conditions that mattered: their network topology, their storage profile, their operational processes, their security requirements.

This is where the engagement became genuinely interesting from a delivery perspective. The POC brought together engineers from three organisations (Insight, Microsoft, and Dell) working at the same table, solving problems collaboratively with a technology stack that was, in some respects, still maturing. When an issue arose with a specific hardware configuration, the Dell engineering team was in the room to address it. When a question emerged about Azure Local's behaviour under a particular workload pattern, the Microsoft engineers could provide direct input rather than routing through support channels. And Insight's team, having led the assessment and designed the architecture, provided the continuity between what was planned and what was being built.

This is also where the groundwork from Stage 1 paid off. Because Dell and Microsoft had been engaged as part of the assessment process, contributing technical input before the destination was confirmed, the POC did not feel like a cold start. There was already a shared understanding of the constraints, the risk areas, and the questions that most needed answering. That prior engagement made the collaborative problem-solving in the POC considerably more effective than it might otherwise have been.

That kind of collaboration does not happen by accident. It happens because the organisation chose to work with a Solutions Integrator that has deep engineering relationships across its vendor partners: the kind of relationships where the right people can be brought into the room quickly, regardless of whose technology is involved. For a production POC of a relatively new platform, that access mattered. It was the difference between a problem taking days to resolve through normal support channels and being addressed in hours with the people who built the technology.

The Phased Commitment

The two-stage model (decide, then validate) normalises something that organisations often feel uncomfortable about: incremental commitment. There is a pervasive expectation in enterprise IT that decisions should be made once and then executed. But in a situation where the technology landscape is genuinely evolving, where new platforms are maturing rapidly and the competitive dynamics between vendors are still shifting, a more considered approach is to commit in stages.

The assessment gave this organisation a defensible decision. The POC gave them confidence that the decision worked in practice, not just in theory. Only then did they move toward full-scale migration planning. At each stage, they had the option to course-correct: to adjust the scope, change the sequencing, or revisit the platform choice if the evidence warranted it.

This approach is as relevant for a commercial enterprise evaluating its options as it is for a public sector body navigating procurement frameworks. The principle is the same: commit incrementally, validate along the way.

This is what disciplined programme delivery can look like in the current VMware landscape. Not paralysis. Not panic. A structured process that moves at the pace of evidence, with the right expertise available at every stage.

This is one organisation's path. Another's will be different. The framework is the same, but the answers will be its own. What matters is that an organisation starts, and that it starts with a clear view of the landscape ahead.

5. The Landscape of Options

The platform landscape available to organisations reconsidering their VMware strategy is broader than it has been at any point in the past decade. That is both an opportunity and a source of noise. Every alternative vendor is, predictably, shouting about migration, and some of what they are saying is genuinely useful. But a vendor's enthusiasm for winning an organisation's workloads is not the same thing as a strategy. What follows walks through the primary options with the same honesty one would apply in a conversation with a trusted colleague.

Staying on VMware

This option comes first deliberately, because leaving VMware is not always the right answer. Switching hypervisor at scale is slow and expensive, and for a large share of organisations the practical near-term question is not whether to leave but how to manage the upgrade to VCF 9.1.1 before vSphere 8 goes end of general support in October 2027.

Section 1 sets out the shape of this in detail: Broadcom's upgrade-first message, VCF 9.1.1's softened definition of what counts as a deployment (vSphere 9.1 plus vCenter 9.1 plus VCF Ops, no vSAN or NSX required), the adoption gap where fewer than 5% of organisations that have bought VCF are actually running it, and Broadcom's continued investment in Glasswing to smooth the upgrade path. Worth adding here is the practical constraint that gap creates: certified adoption partners with the depth to run these programmes are in short supply, even among large systems integrators, and organisations should expect the process to require hands-on support rather than a self-service exercise.

It is also worth being specific about what VCF 9.1.1 actually delivers, because the case for staying is stronger than "the alternative is disruptive." VCF 9.1.1 makes meaningful, practical improvements: upgrades and patching are designed to have zero impact on the application layer, which removes one of the more disruptive aspects of the old VMware patch cycle, and the release includes new security capabilities aimed at reducing exposure if an attacker does gain a foothold in the environment, limiting lateral movement rather than relying solely on perimeter controls. These are genuine, technical reasons to stay and scale on VCF, not just inertia dressed up as a decision. For organisations whose workload assessment points to staying, VCF 9.1.1 is a credible platform to invest in, not merely the least-worst option.

For workloads with hard platform dependencies (Omnissa Horizon being a prominent example) the picture has shifted recently. Omnissa Horizon 8 on Nutanix AHV reached general availability in December 2025, so VMware is no longer the only production-proven hypervisor for that workload. Support for Microsoft Azure Local is still reported to be in active development rather than generally available. Where a VDI estate depends on Horizon, the right move is to check the current state of Omnissa's multi-hypervisor support directly before assuming the workload has to stay on VMware.

There are also organisations for whom the new licensing economics, while painful, are manageable relative to the cost and risk of migration. Broadcom's stated strategy is to focus on its largest customers, approximately 600 strategic accounts globally. Organisations within that tier may find that direct engagement yields terms that change the calculation. That is not guaranteed, and the negotiating dynamics have shifted, but it is a possibility worth testing before assuming the worst.

If the answer is to renew or stay, the decision should not stop at the platform choice. It should extend into the commercial terms, because a renewal negotiated passively tends to lock in the status quo for years. At a minimum, organisations should be pushing for:

- **Price-escalation caps**, typically in the 3–5% annual range.
- **Exit rights** that do not leave the organisation t rapped if circumstances change.
- **Favourable treatment of non-production environments,** so development and test workloads are not priced as full production-tier deployments.

None of this is exotic. It is the kind of protection any organisation should expect to negotiate into a multi-year commitment, and it is worth treating as a checklist item rather than an afterthought.

One further point worth flagging here: regional portfolio composition and pricing terms are not static, and they vary by geography in ways that are easy to miss. What is bundled, what is priced, and on what contract minimum has continued to shift across regions, and the detail of any given offer can change between the time this is read and the time an organisation negotiates. Rather than working from a generic assumption about what

is or is not available, organisations should verify their current entitlements and contract terms directly with the account team before going into a negotiation. That verification step is itself part of building the protections described above.

Staying on VMware is not a failure of strategy. In this phase, it means actively managing the upgrade path, closing the adoption gap, and building the operational capability to run VCF effectively. For many organisations that is itself a significant programme of work. And where the workload mix genuinely needs the broader VCF capability set (private AI and application platforms, security, VKS, software-defined networking, and operations management together) adopting the full bundle rather than vSphere alone is frequently the better total-cost-of-ownership answer, not simply the unavoidable one. The economics work in VCF's favour precisely when an organisation is using what it is paying for.

VMware in the Hyperscaler Cloud

There is a third path that sits between staying on-premises and refactoring for public cloud: running VMware itself inside a hyperscaler. Azure VMware Solution, Amazon's Elastic VMware Service, and Oracle Cloud VMware Solution all let an organisation keep its VMware-based operating model, tooling, and skills while moving the underlying infrastructure onto hyperscaler capacity. For organisations part-way through a data centre exit, or those that want to reduce on-premises footprint without a full application refactor, this is a genuine bridge option rather than a compromise between the two positions either side of it. It carries its own economics and lock-in considerations, and it does not remove the workload-by-workload assessment; it simply adds another destination to weigh each workload against.

Nutanix AHV

Nutanix has emerged as one of the more visible beneficiaries of the Broadcom disruption, and the numbers are notable. In a single quarter in 2025, Nutanix reported 130 net new competitive wins against VMware, with annual recurring revenue reaching over $2 billion. AT&T, Coventry Building Society, and Sydney Airport are among the named organisations that have made the switch.

Price-escalation caps, typically in the – 3–5% annual range. – Exit rights that do not leave the organisation trapped if circumstances change. Favourable treatment of non-production environments, so development and test workloads are not priced as full production-tier deployments. The appeal is reasonably straightforward. Nutanix AHV is an enterprise-grade hypervisor with a hyperconverged architecture that many VMware administrators find operationally familiar. Nutanix Move, their migration tooling, supports live migration with minimal downtime, and the assessment tools are mature enough to give organisations a credible view of their estate before committing. For organisations that want to remain on-premises with a VMware-like operational model, Nutanix is a serious option.

The constraint worth acknowledging is skills. Nutanix expertise is growing rapidly (Nutanix University expanded its free training programmes in 2024 specifically to address the demand) but the talent pool is still considerably smaller than VMware's two-decade head start. Any organisation planning a Nutanix migration should factor retraining and, realistically, a period of dual-skilling into its programme timeline.

Microsoft Azure Local

Azure Local, the rebranded Azure Stack HCI, represents Microsoft's play for organisations reconsidering their on-premises virtualisation strategy. It is Hyper-V based, managed through the Azure portal, and tightly integrated with the broader Microsoft ecosystem: Azure Arc for hybrid management, Azure Virtual Desktop, Azure Kubernetes Service.

Microsoft is not being subtle about its ambitions here. The migration incentive programmes are significant: Azure Hybrid Benefit, free Extended Security Updates for migrated workloads, dedicated FastTrack engineering support, and partner-funded migration assessments through the Azure Migration and Modernisation Programme. For organisations already invested in Microsoft licensing and tooling, the integration story is compelling.

The maturity dimension deserves an honest note. Azure Local is a relatively new platform compared to VMware or even Nutanix. It is improving rapidly (the update cadence from Microsoft is significant) but for mission-critical workloads, the prudent approach is to validate through a proof of concept in the specific environment before committing the rest of the estate to it. That is not a criticism of the platform; it is the same discipline any new technology warrants. The case study in Section 4 followed exactly this model: assessment first, then a production POC to validate the technology before wider commitment.

For public sector organisations, Azure Local has one particular characteristic worth noting: data stays on the organisation's premises, on UK-sited hardware, while the management plane integrates with Azure. That addresses a meaningful subset of data residency

and sovereignty requirements without sacrificing the operational benefits of cloud-style management.

For commercial enterprises already invested in Microsoft 365 and Azure, the licensing synergies and operational familiarity can make Azure Local a natural extension of their existing technology strategy, regardless of sector. Azure Local is not the only route into hyperscaler-managed on-premises infrastructure. AWS Outposts and Oracle Compute Cloud@Customer occupy broadly the same space: hyperscaleroperated infrastructure and control plane, physically sited on the organisation's premises. Each ties an organisation to that hyperscaler's ecosystem in the same way Azure Local ties into Microsoft's, so the choice between them tends to follow whichever public cloud relationship the organisation already has, rather than being decided on the merits of the on-premises layer alone.

Public Cloud

For workloads that are genuinely cloud-ready, public cloud (whether AWS, Azure, or Google Cloud) remains a valid destination. The UK Government's Cloud First policy explicitly encourages this direction, and for applications that can be refactored or that already operate in a cloud-native architecture, the economics and operational model are reasonably well understood.

Commercial organisations face similar considerations, though the drivers are more often cost optimisation, speed to market, and the desire to shift capital expenditure to operating expenditure.

But public cloud is not a universal answer, and treating it as one is precisely the kind of blanket decision this paper has cautioned against throughout. Latency-sensitive workloads, data sovereignty requirements for higher security classifications, the cost of running compute-intensive applications at high volume, and the realities of egress pricing all shape where public cloud is genuinely the right choice. The workload assessment framework applies here as rigorously as it does to any onpremises alternative.

Other Alternatives

Proxmox also deserves acknowledgement, having attracted attention as an open-source alternative: Beeks Financial Cloud, Zen Internet, and the Finnish telecom Elisa are among those that have adopted it. Proxmox has a role, particularly for organisations with strong Linux engineering teams and workloads that do not require enterprise vendor support contracts. For most large enterprises, especially in regulated industries, the absence of formal certification programmes and the reliance on community support tend to make it a less straightforward choice. But the market is moving, and it is worth watching.

The Mixed Estate Reality

Here is the truth that most vendor marketing tends to gloss over: the majority of large organisations navigating this transition will likely end up operating multiple platforms. Some workloads will stay on VMware. Others will move to Nutanix, or Azure Local, or public cloud. The assessment framework does not produce a single destination; it produces a workload-specific answer for each part of the estate.

That is not a compromise. It is a reflection of the genuine complexity of enterprise infrastructure. The organisations that navigate this well are, in Insight's experience, the ones that choose with their eyes open, accepting the operational overhead of a mixed estate because the alternative (forcing every workload onto a single platform) creates its own risks and constraints.

It is also worth saying that for organisations with simpler estates, the mixed estate model is not automatically the right answer. The overhead of managing multiple hypervisors and the associated tooling, training, and support relationships is real. Where the workload profile is relatively homogeneous and the dependencies are limited, a single-platform destination may genuinely be the right outcome: cleaner to operate, easier to staff, and cheaper to support. The framework should produce that answer where it is appropriate.

Every option on this list has trade-offs. None of them is perfect. The question is not which platform has no limitations, but which limitations an organisation can most sensibly live with for each specific workload.

6. The Path Forward

The first phase of this disruption was commercial: new licensing, new pricing, new partner dynamics. This next phase is operational: upgrade pressure, adoption planning, platform decisions, skills investment. The landscape will continue to evolve beyond that. What is less likely to change is the need for a disciplined approach.

These are the principles that matter most when closing out this assessment.

Start with assessment. No organisation can decide where to go until it knows what it has. Not what it thinks it has, but what it actually has: workload by workload, dependency by dependency, constraint by constraint. Almost every organisation discovers things during assessment that change its assumptions. Undocumented integrations. VMware-specific configurations that nobody remembers creating. Application support matrices that limit which hypervisors are viable. The assessment is not a preliminary step before the real work begins; it is the foundation on which every subsequent decision rests.

Resist blanket decisions. The right answer for an organisation's VDI workloads may be fundamentally different from the right answer for its database tier or its development environments. The temptation to choose one platform and standardise everything onto it is understandable (simplicity is appealing) but simplicity purchased at the cost of forcing workloads onto unsuitable platforms is not simplicity at all. It tends to be hidden complexity that surfaces later, usually at the worst possible moment.

Plan for phased delivery. Not everything needs to move at once, and trying to move everything at once is a fairly reliable path to programme failure. A phased approach, starting with the workloads that are most straightforward to migrate, building confidence and capability, then addressing progressively more complex estates, reduces risk, builds organisational muscle, and allows for course correction as the organisation learns. The two-stage model described earlier (assessment followed by proof of concept before scaling) is one expression of this principle. There are others. The common thread is incrementalism over ambition. Accept the mixed estate. The endpoint of this process may well be multiple platforms coexisting within the organisation: VMware for some workloads, Nutanix or Azure Local for others, public cloud for those that are ready. That is not a failure to commit; it is often the natural consequence of making workload-specific decisions based on evidence rather than defaulting to a single answer because it feels tidier.

Recognise that the landscape is still moving. The CMA investigation, the European Commission's antitrust probe, Omnissa's multi-hypervisor roadmap, the maturation of Azure Local, Nutanix's expanding capabilities: all of these will likely shift the calculus over the coming months and years. Build a strategy that can adapt. Lock in where the evidence is clear, but preserve optionality where the ground is still shifting.

Start skills planning now. Not when the migration begins, but now. Whether an organisation is moving to Nutanix, Azure Local, public cloud, or managing an optimised VMware estate, the skills dimension will shape what is achievable and how quickly. Retraining takes time. Hiring in a market where every large enterprise is chasing the same talent takes longer. The organisations that invest in workforce readiness early will tend to have a material advantage over those that treat it as something to address once the technology decisions are made.

Consider the AI horizon. This is still the dimension most often missing from the conversation, which is why it belongs alongside the other eight in Section 3.9 rather than as an afterthought here. Treat it as front-of-mind from the start of the assessment, not a question to come back to once the platform decision has already been made.

Factor in the hard deadlines, and be realistic about what they mean. vSphere 8 reaches end of general support in October 2027. That date is a support cliff, not an operational one: infrastructure does not stop working in October 2027, it stops being supported, which is a different and more manageable category of risk, though not one to take lightly for anything business-critical. It is also worth being honest about the arithmetic. A small number of organisations, generally those with modest, well-documented estates already part-way through a migration, will complete a full workloadby- workload assessment and act on its conclusions before the deadline. Most organisations with a medium or large estate will not, and planning as though they will is itself a version of the blanket thinking this paper has cautioned against. The answer is not to choose between finishing the assessment and hitting the deadline. It is to run both in parallel and let the emerging assessment steer the order of the upgrade work: prioritise upgrading the workloads the assessment is pointing toward staying, since those need to be current and supported regardless of the platform decision, and let workloads that look

likely to migrate sit on their current version for longer, since upgrading them is effort spent on infrastructure that is leaving the platform anyway. That sequencing steadily narrows what is still unsupported after October 2027, while continuing to make progress on the right-workloadright- platform goal rather than putting it on hold until the deadline has passed.

Whatever the strategy, whether an organisation is upgrading within VMware, migrating to an alternative, or operating a mixed estate, the deadline belongs in the programme plan, worked backwards from. And timing matters as much as the planning itself: an organisation that starts six months before a renewal is negotiating from weakness, because everyone in the room knows there is no real time left to walk away. One that starts eighteen months out can credibly threaten to take its business elsewhere, and that credibility is what gets better terms. The deadline is fixed; how much leverage an organisation has when it reaches it is not, and that is determined almost entirely by how early it started.

And start. That may be the most important principle of all. The situation is genuinely complex, and it is tempting to wait for more clarity before acting. But the clarity does not arrive on its own; it emerges from doing the work. From mapping the estate. From testing assumptions against the available platforms. From building the internal capability to make these decisions with confidence.

Insight has spent the past two years working through these questions with organisations across Europe, across different industries, different scales, and different starting points. The specifics vary every time, but the pattern is consistent: the organisations that move deliberately, armed with a cleareyed understanding of their estate and a framework that respects the complexity of the problem, come out with strategies they can defend and programmes they can deliver. And the pattern is not VMware's alone. It is the shape of the reckoning every enterprise still running on-premises virtualisation will face as its own platform comes up for the same kind of vendor-driven reassessment, which is why the method set out here outlasts the Broadcom story that prompted it.

The methodology exists.

The experience exists.

The partnerships across the vendor landscape exist.

What matters now is making a start.

Make the Right Platform Decision

Insight®

WHITEPAPER | Navigating the Hybrid Cloud Ecosystem in 2026