

Navigating the perfect storm

Rebuilding control, resilience and economic flexibility for the hybrid and AI era

This is the Perfect Storm

Enterprise technology leaders face a difficult combination of pressures, many of which are converging at the same time.

Infrastructure component costs are increasingly volatile. Virtualisation economics have shifted. Public cloud continues to deliver valuable agility, yet consumption, data movement and workload placement can make expenditure harder to predict. Meanwhile, years of platform expansion have left many organisations managing fragmented tools, multiple hypervisors, Kubernetes platforms, public clouds and legacy infrastructure across established technology and organisational silos.

Boards still expect IT to reduce costs, strengthen resilience, enable AI and deliver business value at greater pace.

Together, these pressures create the Perfect Storm

The answer calls for more than another tactical infrastructure decision. Replacing a hypervisor may address a licensing concern, but the wider operational complexity remains.

Moving more workloads into public cloud without a clear placement strategy can also transfer existing challenges into a different commercial and operational model.

A sustainable response means moving beyond static infrastructure planning towards continuous optimisation. It means bringing fragmented technology operations together and making workload placement a deliberate business decision.

Hybrid is no longer an interim step towards a different destination. For most complex enterprises, it is the operating model. The opportunity now is to make that model more consistent, flexible and economically sustainable.

How can we restore control over cost, capacity and workload placement while building an operating model ready to support hybrid IT and AI for the next decade?





Economic shock

Commodity Prices, VM license changes and public cloud costs are rising

Architecture drag

Infrastructure stacks weren't built for hybrid and AI



Operational strain

Fragmented tools, complexity siloed teams, rising risk

AI pressure

To deliver value and keep moving in challenged times

1. The Perfect Storm:

Multiple pressures, one strategic problem

CIOs are managing four interconnected forces: economic shock, operational strain, architecture drag and AI pressure.

Each brings its own challenges. Their combined effect is what demands a strategic response.

A licensing increase can be negotiated. A component shortage can be absorbed. Cloud consumption can be optimised and legacy platforms can be modernised.

Managing all these pressures simultaneously is a different proposition. Organisations find their choices increasingly constrained at the point when the business needs greater flexibility.

The underlying problem is a loss of flexibility

Commercial commitments, established architectures, specialist skills, data gravity, regulation and operational dependencies can all limit the options available.

The priority is to recover that freedom of choice, giving IT leaders the flexibility to respond to changing business requirements without repeatedly revisiting the foundations of their technology estate.

2. Economic shock:

Infrastructure economics are being reset

Infrastructure planning once followed relatively predictable refresh cycles and software licensing models. CIOs could make long-term investment decisions with a reasonable degree of financial certainty.

That certainty has weakened

Commodity volatility is now coinciding with significant changes to enterprise virtualisation economics. Following Broadcom's acquisition of VMware, subscription packaging, product consolidation and per-core licensing have changed the commercial considerations for many organisations.

The financial impact depends on the individual estate and contract. Reported price increases are not a reliable universal multiplier. What is clear is that virtualisation costs and renewal exposure now demand greater executive attention.

The decision extends well beyond VMware

Most large enterprises already operate multiple hypervisors, container platforms, Kubernetes distributions and public clouds. Adding another virtualisation platform may reduce a specific licensing cost, while introducing additional management requirements, skills dependencies and operational overhead.

A broader assessment is needed

What operating model will give the enterprise the right balance of cost, control and flexibility across virtual machines, containers, cloud-native applications, data and AI?

That question creates a more useful starting point than a platform replacement decision made in isolation.







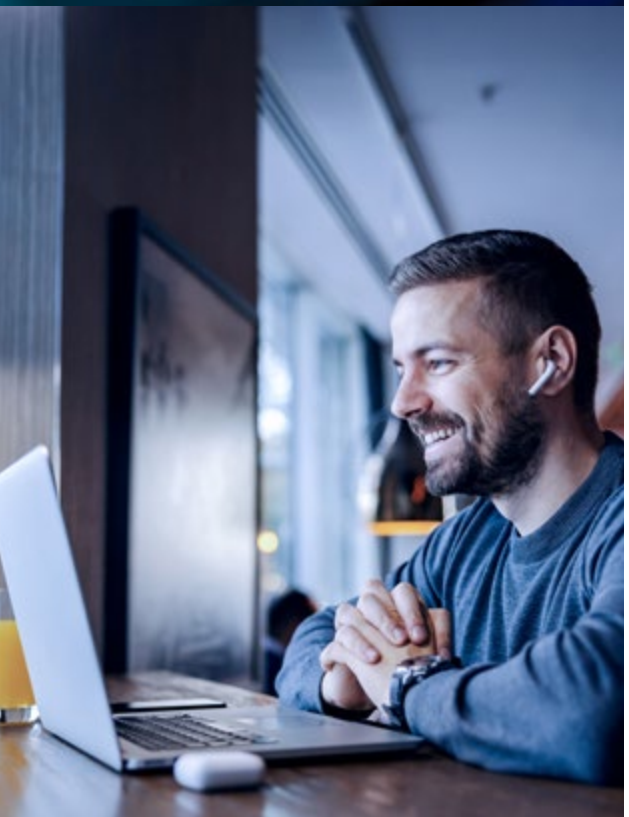

3. Public cloud economics:

Part of the same conversation

Public cloud remains central to enterprise IT strategy. Its agility, scalability and access to cloud-native capabilities continue to create significant value.

The priority is not to retreat from cloud. It is to become more selective about the workloads placed there and the economics that support those decisions.

Cost predictability deserves particular attention

Dynamic consumption, commitment structures, overages, storage and data movement can change the financial picture considerably. Workloads that appeared economically attractive at deployment may no longer deliver the expected value as demand and operating requirements evolve.

Cloud cost management therefore needs to sit within a wider view of technology economics.

FinOps is developing beyond retrospective public-cloud cost control. Its remit increasingly includes SaaS, licensing, private cloud, data centres and AI, helping organisations understand technology expenditure in the context of business outcomes.

Regulation adds another consideration

The EU Data Act introduces cloud-switching and interoperability requirements. These can reduce contractual barriers, but technical dependencies remain. Proprietary services, APIs, data architectures and operating processes can still make migration difficult.

Portability needs to be designed into the architecture from the outset.

For technology leaders, the goal is to maintain access to public-cloud innovation while retaining meaningful choice over cost, placement and future commercial commitments.







4. Architecture drag

Yesterday's architecture constrains tomorrow's business

Many enterprise infrastructure environments were developed around a very different operating model.

Applications primarily resided in corporate data centres. Virtualisation provided the main abstraction layer. Compute, storage, networking and other infrastructure domains could be managed through established teams and processes.

Today's environment is considerably more distributed

Applications, VMs, containers, data and AI now operate across enterprise data centres, colocation facilities, multiple public clouds, edge environments, branches, factories, sovereign locations and isolated infrastructure.

Established architectures have not always evolved at the same pace.

VM-centric practices, accumulated technical debt, ageing equipment and inconsistent platforms can continue to support current operations while making future changes increasingly difficult.

This is architecture drag

An environment may remain stable and functional, yet require substantial engineering effort whenever the business introduces a new service, changes its operating model or expands into a new location.

For CIOs, the more meaningful measure is the organisation's ability to adapt.

Can the existing architecture support changing business priorities without requiring extensive redesign every time a new initiative begins?

Architecture strategy should create room for change, not become another constraint on delivery.

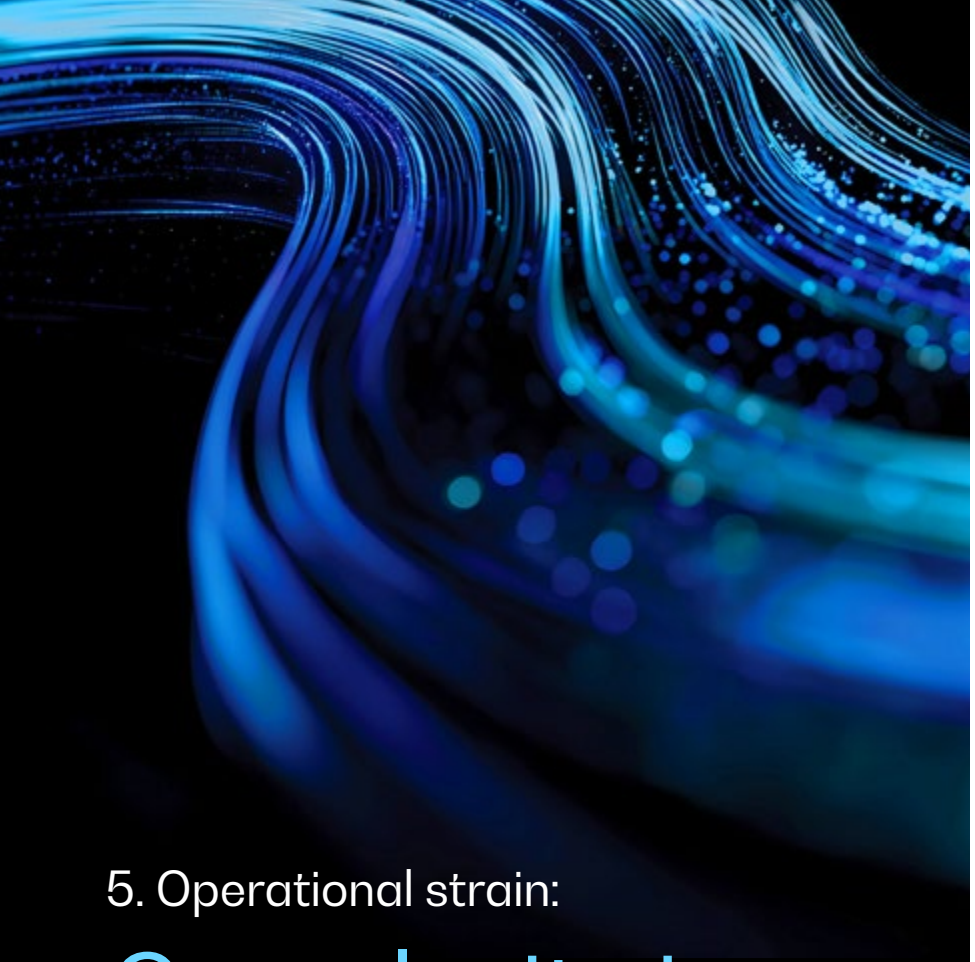

5. Operational strain:

Complexity is a risk multiplier

Fragmented architectures often bring fragmented operations. Compute, networking, storage, cloud, Kubernetes, observability, protection, disaster recovery, security and cost management may each have their own platforms, tooling and operating processes.

Fragmented platforms and processes

Compute, networking, storage, cloud, Kubernetes, observability, protection, disaster recovery, security and cost management may each have their own platforms, tooling and operating processes.

Complexity slows delivery

Over time, individual technology domains develop specialist expertise, separate dashboards, different controls and independent escalation paths. This can make even routine changes more demanding than they should be.

Rising costs and operational risk

Teams spend time coordinating across platforms and organisational boundaries. Duplicated tooling and specialist skills increase operating costs. Inconsistent governance, configuration, patching and recovery policies create additional exposure.

The impact reaches beyond the infrastructure team

Longer delivery cycles, competing operational priorities and limited visibility across the estate can affect the organisation's ability to respond to new business requirements.

Simplification is therefore an opportunity to improve productivity, strengthen resilience and reduce operational risk.

The aim is not necessarily to remove every platform difference. It is to establish consistent management, governance and operational practices wherever they deliver meaningful value.

6. AI pressure:

Turning existing challenges into urgency



AI brings a new level of urgency to infrastructure modernisation. Unlike many traditional enterprise applications, AI relies on the close interaction of compute, data and connectivity.

Models may be developed in one environment, enterprise information may reside in another and inference may need to run close to users or operational systems.

AI could become the most hybrid workload an enterprise manages

The challenge reaches beyond access to models or compute capacity. Technology leaders must consider time to value, the pace of model and tooling innovation, data privacy and control. Infrastructure readiness and governance can determine how successfully AI initiatives move from experimentation into production.

There is also a direct financial consideration

Unused capacity, unnecessary licensing commitments, dormant cloud resources and duplicated operational tooling tie up budget that could support AI investment.

Modernisation and optimisation can therefore play an important role in funding the next phase of innovation.

The opportunity is to create an infrastructure and operating model that supports AI development and deployment without losing control of the wider technology estate.



7. Data gravity:

Changing the AI placement equation

Moving large volumes of sensitive information to an external model or service is not always practical, commercially attractive or consistent with regulatory obligations.

Healthcare information, financial data, government records, manufacturing intellectual property and security-sensitive data may all require specific controls over location, access and processing.

> For AI, the defining infrastructure question increasingly becomes: **where is the data?**

For many organisations, bringing compute closer to the data offers a more appropriate approach.

That consideration changes the strategic role of private infrastructure.

Private cloud can provide an execution environment for predictable workloads, sensitive information and data-adjacent AI, operating alongside public-cloud services.

The placement decision should reflect the demands of the workload, the value of the data and the level of control required.

AI strategy and infrastructure strategy can no longer be considered separately.



8. Sovereignty:

A resilience requirement

Infrastructure choices increasingly carry implications for organisational control as well as technology performance and cost.

> Sovereignty is resilience under constraint.

Digital sovereignty extends beyond the physical location of information. It requires a broader view across four areas: control, residency, auditability and autonomy.

Organisations need confidence that they can govern their data, demonstrate compliance and retain appropriate operational independence as commercial, regulatory and geopolitical conditions evolve.

Sovereignty should therefore influence architecture and operating-model decisions from the beginning of a programme. Adding controls late in delivery may leave underlying technical and operational dependencies unresolved.

A sovereign-by-design approach provides a more coherent foundation for workload placement, governance and resilience.

9. Loss of control:

The underlying strategic problem

The most significant challenge is not selecting a cloud provider or replacing a hypervisor. It is maintaining sufficient control over a technology estate operating in increasingly unpredictable conditions.

Three areas deserve particular attention.



01 Capacity control

Traditional procurement models often force organisations to make difficult trade-offs between excess capacity and future demand.

Overprovisioning ties up capital in resources that may remain unused. Underprovisioning introduces the possibility of service disruption and operational risk.

AI adds another variable, with demand patterns that can be difficult to forecast through conventional planning cycles. A more responsive approach to capacity helps organisations balance availability, utilisation and financial commitment.



02 Cost control

Technology costs now change across several interconnected areas. Software licensing, public-cloud consumption, infrastructure components, data services and AI each introduce commercial considerations that require ongoing attention.

Without a consistent view, expenditure becomes harder to forecast, attribute and optimise. The priority is to connect financial visibility with operational decisions, giving technology leaders more control over both current spending and future commitments.



03 Workload-placement control

Workloads may remain in environments selected for historical rather than current business reasons. Changes in economics, performance requirements, compliance obligations, data proximity, resilience and strategic priorities can all affect the most appropriate location.

Regular placement reviews help ensure infrastructure decisions continue to reflect business value. The objective is to retain the flexibility to change those decisions as requirements develop.A more responsive approach to capacity helps organisations balance availability, utilisation and financial commitment.

10. Workload placement:

Moving from cloud-first to workload-right

The next phase of enterprise infrastructure strategy is not defined by a return to the data centre or an indiscriminate move to public cloud. A more useful principle is to place each workload where it delivers the strongest combination of business value, operational suitability and economic return.

Importantly, organisations should retain the ability to review and change that placement.

01 More selective public-cloud placement:

Use public cloud where elasticity, experimentation, geographic reach or cloud-native services provide clear advantages. Assess the commercial and operational implications alongside those benefits.

02 AI closer to enterprise data:

Bring compute closer to governed information when privacy, sovereignty, latency or data gravity makes moving that information less appropriate.

03 More distributed private environments:

You feel the estate every day, in tickets, in deployments, in the machines that never quite behave. Prioritising by experience cuts the noise at source. You refresh the devices generating the most friction, not the ones a schedule happens to surface.

04 Hybrid as the operating model:

Recognise that heterogeneous environments are an established part of enterprise IT. Build the governance, skills and processes needed to operate them consistently.

05 Hybrid by design:

Select platforms and operational controls developed to work across different environments. Make integration, visibility and flexibility part of the architecture rather than exceptions managed later.

Run each workload where it makes the most business sense, and retain the ability to change that decision.

11. Continuous optimisation:

Replacing static infrastructure planning

Traditional infrastructure planning followed a familiar sequence: size, purchase, deploy, operate and refresh.

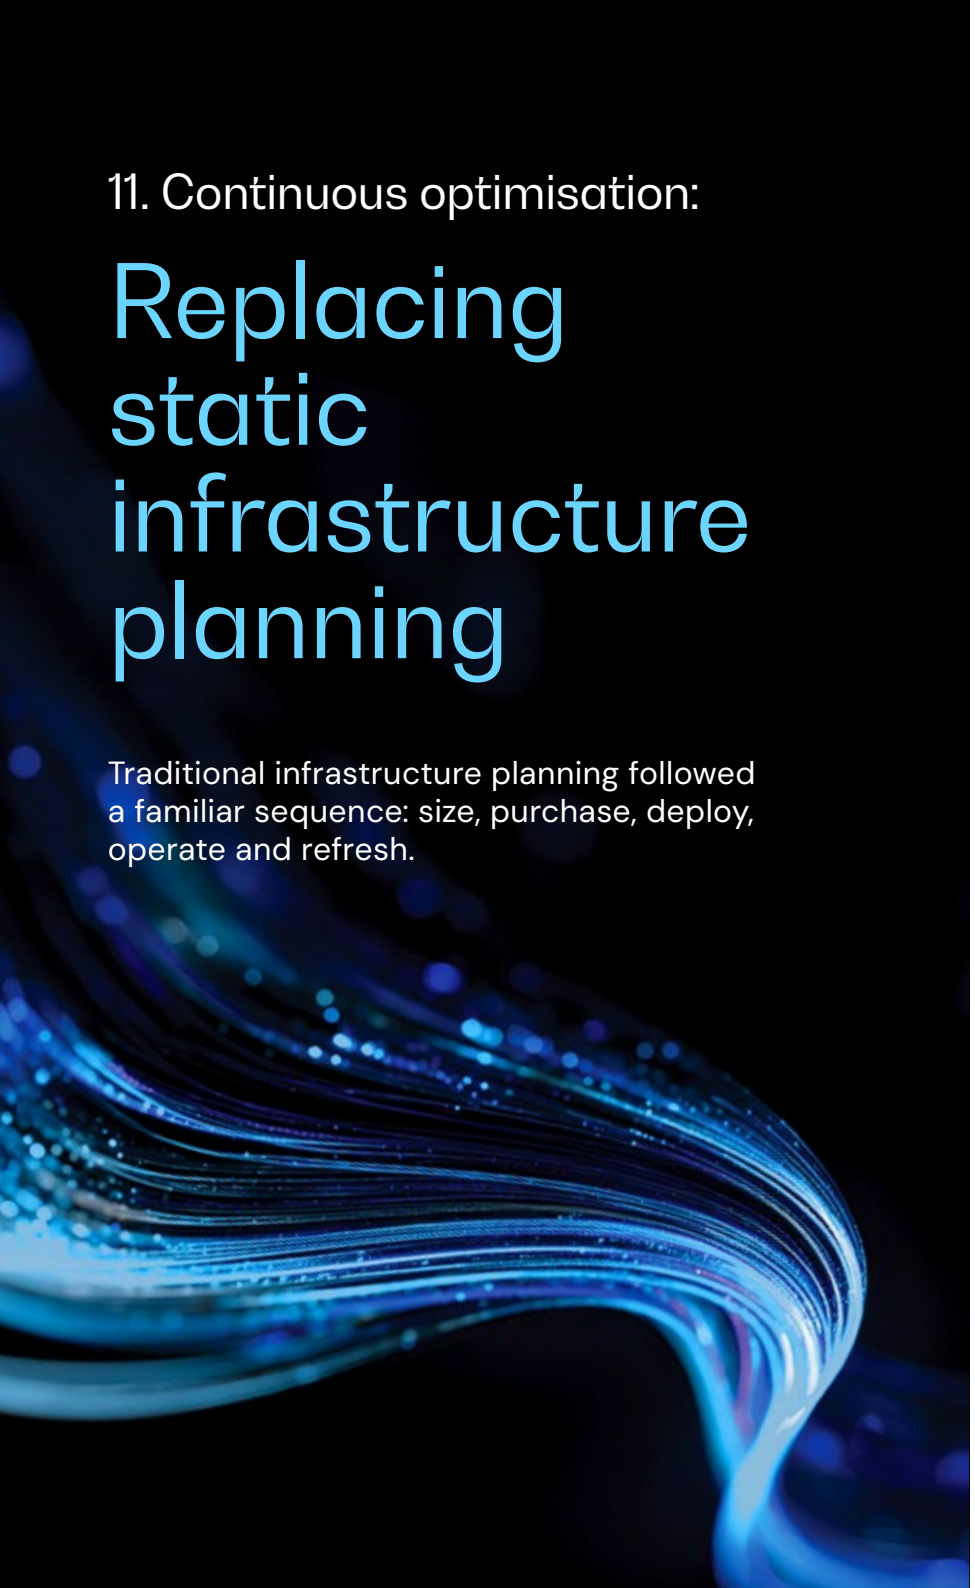

That approach relied on relatively predictable demand and long investment cycles. Today's combination of dynamic workloads, changing licensing models and AI requirements calls for a more responsive approach. A continuous cycle of observing, placing, consuming, measuring, optimising and rebalancing provides greater flexibility.

Static planning	Continuous optimisation
Buy for estimated peak	Use measured demand and buffer capacity
Refresh on long cycles	Optimise and rebalance continuously
Manage platforms separately	Manage technology value across the estate
Place once	Review placement as economics and obligations change
React after the bill	Introduce financial context before deployment

This represents a broader change in infrastructure management

Capacity and workload placement become ongoing portfolio decisions rather than activities reserved for procurement or refresh cycles.

FinOps, enterprise architecture, infrastructure, operations, security and application teams all have a role in making that approach effective.

Shared visibility and coordinated decision-making help ensure technical and financial priorities remain aligned.

12. Strategic response:

A practical roadmap for CIOs

Addressing these pressures calls for a structured plan that creates near-term value while building towards a more flexible operating model. The work should begin with a clear understanding of the existing estate, followed by targeted decisions on cost, placement, operations and future investment.

01 Establish the real baseline

Create a workload-level view of utilisation, runtime licensing, public-cloud consumption, storage growth, operational tooling, application dependencies, data location, resilience, AI demand and sovereignty constraints.

02 Attack stranded economics first

Identify overprovisioned infrastructure, unused licences, dormant cloud services, oversized instances and unnecessary tooling overlap. Use released budget to create capacity for modernisation.

03 Segment workloads rather than platforms

Classify workloads by cost, elasticity, data gravity, latency, regulation, resilience, portability and AI requirements so a hypervisor decision does not accidentally become the enterprise architecture strategy.

04 Build a cleaner hybrid foundation

Reduce unnecessary platform variation while standardising management, observability, automation, security policy and resilience. Retire legacy platforms progressively rather than attempting wholesale replacement.

05 Introduce consumption discipline across the estate

Apply FinOps-style measurement to cloud, software, licensing, private infrastructure and AI. Measure against business value and workload unit economics, not only aggregate spend.

06 Design sovereignty and portability in

For every strategic workload, determine whether it can continue, move or recover if commercial, geopolitical, regulatory or supplier conditions change.

07 Change the commercial model alongside the technical model

Align capacity, software and infrastructure commitments with the flexibility expected from the technology.
A dynamic architecture constrained by static economics will not deliver its intended value.







13. The strategic destination:

Control without losing cloud velocity

The objective extends beyond reducing infrastructure costs. It is to create greater strategic freedom while retaining the speed and capability that modern cloud technologies offer.

> The winning question is not “What should replace my hypervisor?” It is “How do I stop the next licensing, commodity, cloud or AI shock from dictating my strategy?”

For enterprise technology leaders, that means developing the ability to:

- **Consume capacity without committing permanently to unnecessary overprovisioning.**
- **Move workloads without redesigning the wider operating model.**
- **Place data in line with sovereignty and governance requirements.**
- **Run AI closer to the information that creates business value.**
- **Manage heterogeneous environments through consistent policy and visibility.**
- **Reduce unnecessary dependency on a single runtime or cloud.**
- **Understand and optimise technology economics continuously.**
- **Protect budget and organisational capacity for innovation.**

Platforms such as HPE GreenLake, HPE Morpheus, HPE OpsRamp, HPE Zerto and HPE Private Cloud have a role within this wider strategy.

Their relevance extends beyond individual product capabilities

Together, they can contribute to an operating model designed to improve common management, resilience, portability and economic flexibility across heterogeneous environments.

Technology selection should support the operating model the business needs, with each component contributing to measurable outcomes.

Conclusion:

Use disruption to reset the operating model

Several assumptions that shaped infrastructure strategy over the past decade are now being challenged simultaneously.

Commodity economics are volatile. Virtualisation licensing has changed materially. Public-cloud costs demand closer management. Infrastructure continues to become more distributed and operational complexity remains a significant constraint.

At the same time, sovereignty has moved higher on the executive agenda and AI introduces new demands on compute, data and connectivity.

These developments create difficult decisions for organisations balancing investment commitments, operational continuity and the need to innovate.

They also present an opportunity to reconsider the wider operating model.

Protecting the existing estate more efficiently may deliver some immediate benefits. A more sustainable response is to use the current disruption to address architecture drag, simplify operations and regain flexibility.

The strategic shift is clear.

Move from	Move to
Static capacity planning	Continuous optimisation
Cloud-first	Workload-right
Infrastructure silos	Common operations
Vendor dependency	Architectural choice
Sovereignty as compliance	Sovereignty as resilience
AI experimentation	An enterprise AI operating model

The organisations best positioned for change will not necessarily be those that predict every market development. They will be those with the visibility, control and flexibility to adapt their technology decisions as conditions evolve.

The advantage lies in building an operating model that can respond to change without compromising business priorities.

The executive question

Consider the next significant disruption.

It could be another major licensing change, a commodity shock, a new sovereignty requirement or a substantial increase in AI demand.

Could your current architecture absorb that change while preserving control over cost, resilience and innovation? The answer is a useful measure of how ready the operating model is for the next decade.

The winners will not predict every change. They will build enough control, flexibility and observability to adapt when change arrives.

Sources and notes

1. SCC & HPE, “Navigating the Perfect Storm” webinar deck, version 1.0, 15 October 2026. Primary source for the Perfect Storm, hybrid operating model, sovereignty, modernisation and response narratives.
2. IDC, “VMware Cost Increases – How Broadcom VMware Product Offerings Are Evolving”, 7 August 2024. Provides context on licensing, subscription and product-bundle changes.
3. ColocationPlus, “VMware Price Increase in 2025: What You Need to Know”, updated 30 June 2025. Provides reported ranges of customer cost increases. Individual outcomes vary by estate and contract.
4. FinOps Foundation, “State of FinOps 2026”. Covers the expansion of FinOps into AI, licensing, private cloud and data-centre economics.
5. FinOps Foundation, “FinOps Framework 2026”. Covers the shift towards executive strategy alignment and cross-technology value management.
6. Regulation (EU) 2023/2854, the EU Data Act, including cloud-switching and interoperability provisions applicable from 12 September 2025.
7. UK Competition and Markets Authority evidence on cloud egress and switching programmes, Appendix N, published 2025.

Editorial note: Examples and reported cost increases are directional. Each organisation should validate them against its contracts, architecture, utilisation and regulatory obligations before making investment decisions.

Copyright © SCC UK 2023-2026 SCC. All rights reserved.

Specialist Computer Centres plc
Registered address: James House, Warwick Road, Birmingham, West Midlands, B11 2LE
Company reg number: 01428210 | VAT reg number: 313651680 | DUNS number: 227720521