



ebook:

Navigating the modern cloud hosting landscape: a practical guide

Rethinking hyperscalers and choosing the right hosting strategy for your business.

Author: Ludi Nel, iSphere Cloud Lead, EPI-USE Labs

Ludi brings over twenty years of IT industry experience to his role as iSphere Cloud Lead. Having worked on both sides of the service delivery model, as an end-client and as a service provider, Ludi has an in-depth understanding of the challenges businesses face when moving to the cloud, along with the security and compliance complexities that come with it.

A career spanning Support Engineer, Network Administrator, and IT Manager roles for major mining and business clients, followed by co-founding a cloud provider specialising in SAP®, has shaped Ludi's deep expertise in cloud computing, security, and service delivery. He remains dedicated to advancing cloud strategy and compliance in an ever-evolving digital landscape.



Contents

Introduction

The cloud conversation has changed.....	3
---	---

Chapter 1

The private cloud renaissance: Why enterprises are rethinking hyperscalers.....	4
---	---

Chapter 2

The sovereign cloud imperative: Why data sovereignty is no longer optional.....	6
---	---

Chapter 3

Understanding your cloud options: Private, public, hybrid, and managed.....	8
---	---

Chapter 4

The hidden costs of hyperscalers: What your bill is really telling you.....	11
---	----

Chapter 5

Before you sign: 10 questions every enterprise should ask their cloud provider.....	13
---	----

Chapter 6

What a good migration looks like: From assessment to go-live.....	15
---	----

Conclusion

Making your decision.....	17
---------------------------	----


The cloud conversation has changed

Why enterprise IT leaders are asking harder questions of their cloud providers

The cloud was supposed to make enterprise IT simpler, but for many organisations, it has done the opposite. Bills that seemed reasonable at the start have grown unpredictable. Compliance teams are raising questions about where data actually resides and whether the answer would satisfy a regulator. The 2024 Broadcom acquisition of VMware sent shockwaves through enterprise IT departments, forcing thousands of organisations to reassess their infrastructure strategy almost overnight.

At the same time, the regulatory landscape has tightened significantly. Europe's General Data Protection Regulation (GDPR) has matured from a compliance checkbox into an active enforcement regime. Equally stringent data privacy legislation has emerged across the world. Executive boards and finance directors are asking questions they never asked before: "Where does our data actually live? What foreign laws apply to it? What happens if a jurisdiction changes?"

This is the context in which enterprise cloud decisions are being made today. The market has matured, and the early promises of the hyperscaler era are now being weighed against the lived experience of scale-up costs, support tiers, and the uncomfortable reality of egress fees ('Cloud Lock-ins' – see Chapter 4). And a new category is gaining ground: sovereign, managed private cloud, with infrastructure built around the client's requirements rather than the provider's economics.



The question is no longer whether to use the cloud. The question is which cloud, for which workloads, and under whose jurisdiction.

This guide is for IT leaders, finance directors, and CIOs who are asking difficult questions. Not just "How much does it cost?", but "What happens if I want to leave?", and "Who will I speak to when something breaks at 2am?".

What you will find in this guide:

- Why enterprises are moving away from hyperscalers, and what they are moving to
- Why data sovereignty has become a board-level conversation
- How to understand your true total cost of cloud ownership
- A practical framework for evaluating cloud providers
- Ten questions to ask any provider before you sign
- What a well-managed migration looks like from start to finish.

The private cloud renaissance

Why enterprises are rethinking hyperscalers

The promise versus the reality

The hyperscaler proposition was compelling: infinite compute* on demand, global reach, consumption-based billing, no capital expenditure. For certain workloads – such as high-growth consumer applications, experimental projects and burst compute** – it delivered. Enterprise IT leaders embraced it, and for much of the 2010s, the migration to public cloud was largely unquestioned.

But the economics of the hyperscaler model are not neutral; they are designed to reward consumption. Compute is sold at low margins to draw workloads in; the revenue is recovered through egress fees, support tiers, premium services, and the sheer complexity of the billing model. Organisations that moved significant workloads to Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform (GCP) discovered that costs grew faster than anticipated and were harder to predict than promised.

The turning point


Several forces converged to shift enterprise thinking. The Broadcom acquisition of VMware in 2024 drastically increased licensing costs and created uncertainty around on-premise environments. This disruption acted as a catalyst that forced enterprises to urgently re-evaluate their overall cloud strategy. Instead of facing a forced migration on someone else's timeline, organisations began looking toward more flexible, diverse infrastructure approaches to regain control over their systems and budgets.

Simultaneously, the regulatory environment has tightened (see map in Chapter 2). According to [an IAPP analysis in early 2026](#), 179 of 240 analysed jurisdictions have data protection frameworks in place. It's estimated that approximately three out of every four countries are covered by data protection laws, and approximately 80% of the world's total population.

GDPR enforcement action has accelerated across Europe; and the Protection of Personal Information Act (POPIA) has become very active in South Africa. New requirements around data residency, audit trails, and sovereignty have emerged across sectors including financial services, healthcare, and government. Cloud providers headquartered in the United States carry obligations under US law, including the Clarifying Lawful Overseas Use of Data (CLOUD) Act (see Chapter 2) that can create conflicts with European and other regional data protection requirements.

* Compute: a generic term used to reference processing power, memory, networking, storage, and other resources required for the computational success of any programme.

** Burst compute: rapid, flexible scaling of specialised resources to meet short-term, high-demand, intensive workloads efficiently and cost-effectively.



64%

of enterprises report
cloud costs exceeded
budget in 2024

3×

average VMware
licensing cost increase
post-Broadcom

71%

of CIOs cite data
sovereignty as a top
infrastructure concern*

What are companies choosing instead?

Private cloud, once dismissed as the conservative choice for organisations reluctant to change, is experiencing a renaissance. Not because the technology is new, but because the business case has shifted. Dedicated, managed private infrastructure offers something the hyperscalers cannot: predictability. Fixed monthly costs, a single engineer who knows your environment, and infrastructure that lives in a jurisdiction you have chosen and can defend to a regulator.

This does not mean abandoning public cloud entirely. Hybrid strategies are increasingly the norm; such as private cloud for core, sensitive, or regulated workloads, alongside managed public cloud for appropriate use cases.

The shift is to decisions made on the basis of workload requirements and business risk, rather than vendor defaults.

Why private cloud does not mean going backwards

A common misconception is that choosing private cloud means accepting an older technology stack or giving up the operational benefits of Managed Services. Neither is true. Modern managed private cloud providers offer self-service portals, automated monitoring, elastic resource allocation, and the same enterprise-grade redundancy as hyperscalers; with the addition of dedicated tenancy, engineering expertise, and a billing model that does not penalise growth.

References


<https://www.ciodive.com/news/cloud-storage-overspend-wasabi/740940/>

<https://www.cio.com/article/2513389/vmware-licensing-and-pricing-hikes-what-options-do-you-have.html>

<https://www.cio.com/article/3950867/navigating-the-future-of-national-tech-independence-with-sovereign-ai.html>

The sovereign cloud imperative

Why enterprise data sovereignty is no longer optional



Global data privacy laws

[Read about global data privacy compliance](#)

What data sovereignty actually means

Data sovereignty refers to the principle that data is subject to the laws and governance structures of the nation in which it is stored and processed. For large companies, this has two practical dimensions: where your data physically lives, and which legal jurisdiction has authority over it.

These two things are not always the same. Data stored on a server in Ireland which is operated by a company headquartered in the United States may be subject to both EU data protection law and US legal process.

Understanding this distinction is fundamental to making a sound cloud decision and to defending that decision to regulators, auditors, and your board.

The geopolitical dimension

Beyond formal regulation, the geopolitical environment has introduced new uncertainties. Governments in multiple regions have introduced or strengthened powers to compel access to data held by organisations operating in their jurisdiction. For enterprises in sensitive sectors such as financial services, healthcare, energy and government supply chains, the question of which government could legally access their data is no longer hypothetical.

Choosing a cloud provider means choosing a jurisdiction. A provider headquartered in one country, operating infrastructure in another, and subject to the laws of a third creates a complex web of potential obligations. Sovereign cloud where infrastructure, operations, and legal entity are all aligned to a specific jurisdiction eliminates this ambiguity.

Technical versus contractual sovereignty

True sovereignty requires both technical and contractual assurance. Technically, it means your workloads run on hardware physically located in your mandated jurisdiction, with no data traversing public internet or crossing jurisdictional boundaries without explicit authorisation. Contractually, it means your service agreement includes a data residency clause specifying which physical locations your data may reside in, along with audit trail capabilities to support regulatory compliance.

When evaluating a cloud provider, ask for both. A contractual commitment without technical enforcement is insufficient. Technical isolation without contractual documentation will not satisfy a regulator.

Contractual sovereignty	Technical sovereignty
✔ Service agreement residency clause Includes a legally binding data residency clause in the agreement.	✔ Local physical hardware Workloads run on hardware physically located within your mandated jurisdiction.
✔ Defined physical locations Explicitly specifies which physical locations your data may reside in.	✔ No unauthorised public internet Ensures no data traverses the public internet without explicit authorisation.
✔ Audit trail capabilities Provides comprehensive audit logs to support regulatory compliance requirements.	✔ Strict jurisdictional boundaries Prevents data from crossing jurisdictional boundaries without explicit authorisation.
✔ Regulatory documentation Provides the necessary legal documentation to satisfy a regulator.	✔ Technical enforcement Provides the actual physical and technical isolation of environment and workloads.

Understanding your cloud options

Private, public, hybrid, and managed: which is right for your organisation?

There is no universally correct cloud strategy. The right approach depends on your workloads, your regulatory environment, your tolerance for variable costs, and the internal expertise available to manage your infrastructure.

Outlined below is a practical overview of the main models, and a framework for matching them to your requirements.

The four models

Dedicated private cloud

Infrastructure dedicated exclusively to your organisation. No shared compute, no shared storage, no shared network. Physical hardware may be owned by you or by the provider, but it is not shared with any other customer.

This model offers the highest level of security, performance consistency, and sovereignty. It is the appropriate choice for sensitive workloads, regulated industries, and organisations with strict data residency requirements.

Leveraged (shared) private cloud

Private cloud delivered on shared infrastructure owned by the provider. All services are consumed as-a-service via a self-service portal.

This model offers many of the operational benefits of private cloud including Managed Services, predictable billing and dedicated support. As this comes with a lower cost base, it is well suited to medium-complexity workloads where dedicated hardware is not a regulatory requirement.

Managed public cloud

Organisations use public cloud infrastructure (Azure, AWS, GCP) but engage a managed service provider to operate it on their behalf.

This preserves access to hyperscaler scale and services while offloading the operational burden. It is appropriate for workloads that benefit from hyperscaler-specific capabilities, or where existing investments in a particular platform make migration impractical.

Hybrid cloud

A combination of private and public cloud, with workloads distributed according to their requirements. Core and sensitive workloads run on private infrastructure in a defined jurisdiction; workloads that benefit from scale or hyperscaler-specific services use managed public cloud.

This is increasingly the default strategy for large enterprises, and is enabled by a managed service provider who can operate across both environments.

SAP-specific considerations

SAP workloads, particularly SAP HANA and S/4HANA, have specific infrastructure requirements that generic cloud providers are not always equipped to meet. HANA's in-memory architecture demands NVMe storage, high-memory compute, and low-latency networking.

Organisations running SAP must ensure they choose a partner whose underlying infrastructure stack is fully SAP-certified and architected strictly around SAP design recommendations. You need to:

- Ask for evidence of experience with your specific SAP version and landscape
- Understand how the provider handles SAP landscape optimisation and rightsizing.

The role of Managed Services

Regardless of which cloud model you choose, the quality of Managed Services will determine your day-to-day experience. Managed Services, covering Operating System (OS) monitoring, patch management, networking, security, backup, and disaster recovery, remove the operational burden from your internal team.


The alternative is maintaining in-house expertise across all of these disciplines, which carries its own cost and risk.

When comparing providers, examine their Managed Services scope in detail: what is included, what is optional, and what is out of scope entirely.



Evaluating cloud options for SAP

Complexity in choosing the right cloud partner and model



The hidden costs of hyperscalers

What your AWS or Azure bill is really telling you

The hyperscaler sales pitch focuses on compute costs, which are genuinely competitive. The real cost of hyperscaler cloud is spread across a much wider range of line items, many of which are not visible until you are already committed. Understanding this cost structure is essential before making any cloud decision.

Egress fees: the silent budget killer

Hyperscalers charge for data transferred out of their infrastructure. This is called an egress fee, and it can represent a significant and unpredictable proportion of your total bill. Costs are typically charged per gigabyte of outbound transfer, and they apply whether you are transferring data to your users, to another cloud provider, to your on-premise systems, or in the process of migrating away.

For data-intensive workloads like SAP systems generating large reports, backup and replication traffic, and analytics pipelines, egress fees can easily reach tens of thousands of dollars per month. They are deliberately structured to make leaving expensive.



Hyperscalers draw organisations in with low compute costs and trap them with egress fees and consumption volatility. The total cost of ownership is rarely what was projected at the outset.

Support tier lock-in

Basic hyperscaler support plans provide access to documentation and community forums. Meaningful support – the kind that helps you resolve a Production issue – is only included in premium support tiers, and can add 10–15% to your monthly bill. Even with premium tiers, you are interacting with a support queue; not an engineer who really knows your environment. For organisations running business-critical systems, the true cost of support is rarely accounted for at the point of decision.

Consumption volatility

Finance teams responsible for cloud cost management cite consumption volatility and SKU overheads as part of their top operational challenges.

Consumption-based billing means your invoice varies from month to month based on actual usage. For workloads with predictable patterns, this is manageable. For complex enterprise environments with variable processing loads, backup jobs, and data transfer events, it creates real budget uncertainty.

An additional cost burden compounding this issue is SKU overheads. In many cloud environments, workloads are sold in fixed SKUs, or 'T-shirt sizes', which frequently forces clients to pay for resources they do not actually need. Because there is no granular upgrade path for individual resources, upgrading your memory means you must purchase it in incremental chunks that are invariably tied to a vCPU upgrade as well. This lack of flexibility leads to unnecessary overspending on unused capacity.

Rate of exchange variances

Hyperscalers typically charge in dollars; and exchange rates in regions such as Africa, South America and Asia can vary significantly. This can lead to considerable fluctuations in infrastructure costs.

Building a true TCO comparison

A meaningful total cost of ownership comparison between cloud models should include the following:

- Compute and storage costs
- Egress and data transfer fees
- Support and account management
- Licensing (OS, monitoring tools, backup software)
- Internal staff time for management and operations
- Cost of any professional services or migration work
- Cost of exit if you choose to change provider

When these factors are included, managed private cloud frequently offers a comparable or lower five-year TCO than equivalent hyperscaler deployments, particularly for stable enterprise workloads where the consumption model offers no benefit.



Before you sign

10 questions every enterprise should ask their cloud provider

The right provider will answer these questions clearly, and in writing. Vague answers, deferred responses, or reluctance to commit to specifics are informative in themselves.

1

Where exactly will my data live?

Ask for the specific physical locations, including data centre name, city, country. Understand whether data is ever replicated or processed across jurisdictions, even temporarily. A contractual commitment to data residency should specify physical locations, not just regions.

2

Who has legal access to my data?

Which government can compel your provider to hand over your data? Where is the provider's legal entity incorporated? Understand the interplay between your data residency requirements and the provider's legal obligations in their home jurisdiction.

3

What happens to my costs if my usage grows?

Is billing fixed, consumption-based, or a hybrid? Are there egress fees? Are there tier changes that apply above certain usage levels? Get a clear picture of what your bill will look like at twice your current workload.

4

What does your SLA actually guarantee?

Uptime SLAs should specify what constitutes downtime, how it is measured, and what compensation applies. Understand whether the SLA covers your full stack, or just the physical infrastructure. Ask what the historical uptime has actually been.

5

How do you handle a Production outage?

Who do you call? What is the response time commitment? Is there a direct line to an engineer, or a support queue? Ask for a worked example of how a P1 (Priority 1) incident was handled for an existing client.

6

Is my infrastructure shared with other clients?

Dedicated single-tenant and shared multi-tenant environments have materially different security, performance, and compliance profiles. If the answer is shared, understand what isolation mechanisms are in place and whether they satisfy your compliance requirements.

7

What does migration in and out look like?

Understand the migration methodology, timeline, and risk management approach for onboarding. Then ask the same questions about offboarding. What formats is your data exported in? Who owns the migration process? What are the exit costs?

8

How are security patches and OS updates managed?

Unpatched Operating Systems (OS) are one of the most common vectors for enterprise security incidents. Understand whether OS patching is included in your managed service, who is responsible for scheduling, and how the provider handles emergency patches for critical vulnerabilities.

9

Do you understand SAP environments?

If you run SAP workloads, this is non-negotiable. Ensure your specific needs are understood and accounted for, with your specific SAP version and landscape.

10


Who will I actually speak to when something goes wrong?

The person you speak to at 2am matters. Is there a named engineer assigned to your account? Will they know your environment, or will you be explaining your architecture from scratch each time? The answer to this question tells you more about the provider than any SLA document.

What a good migration looks like

From assessment to go-live, a framework for low-risk cloud migration

The five stages of migration



Migration is the stage where poorly planned cloud decisions become expensive. Rushed migrations introduce risk to Production systems. Under-resourced migrations create organisational disruption. Over-scoped migrations delay the benefits that justified the decision in the first place. A well-managed migration follows a structured approach with clear deliverables at each stage.

Stage 1: Assessment

Before any technical work begins, your provider should carry out a thorough assessment of your current environment: workload inventory, dependency mapping, performance baselines, and compliance requirements. This is also the stage at which your provider should identify any workloads that require special handling such as SAP Production systems, workloads with strict latency requirements, or data subject to specific regulatory controls. The output of the assessment stage is a clear migration architecture proposal and a risk register.

Stage 2: Architecture and design

Based on the assessment, your provider designs the target environment: compute allocation, storage configuration, network topology, security controls, backup schedule, and disaster recovery setup. This design should be reviewed and signed off before any migration work begins. Changes made after migration starts are expensive; changes made on paper are not.

Stage 3: Test migration

Non-production environments migrate first. A test migration validates the technical approach, identifies any issues with data formats or system dependencies, and gives your team a rehearsal of the cutover process. Any surprises should surface here, not in a Production cutover window.

Stage 4: Production cutover

Production cutover is executed in a planned maintenance window with a full rollback plan in place. Your provider's engineers should be available and engaged throughout. The cutover window is agreed in advance, as is the rollback trigger (the specific conditions under which the migration is reversed rather than completed). Never accept a migration plan that does not include a clearly defined rollback procedure.

Stage 5: Go-live and ongoing support

A migration is not complete at go-live. The weeks following cutover are when issues surface, which can include performance anomalies, backup job failures or user connectivity problems. Your engineering team should be in heightened monitoring mode and available for rapid response. This is also the phase in which your team and the provider's engineers align on ongoing operational procedures: monitoring thresholds, escalation paths, change management processes.

What to watch out for

- Providers who skip the assessment stage and move straight to architecture
- Migration plans without a defined rollback procedure
- Cutover windows scheduled without adequate engineering resource
- Post-go-live support that reverts to standard SLA terms immediately after cutover
- Lack of documented Disaster Recovery (DR) test results prior to Production migration



Making your decision

A summary of the key considerations


The cloud decision your organisation makes in the next 12 to 24 months is likely to shape your infrastructure for the next decade. The framework in this guide won't make that decision for you, but it will help you to ask the right questions of the right people.

The primary considerations

- Sovereignty is not optional. Understand where your data lives, under which jurisdiction, and what legal access could be compelled by which government.
- Total cost of ownership is not the same as compute cost. Model egress fees, support, licensing, and operational overhead before comparing providers.
- Managed Services quality matters more than infrastructure specs. The people managing your environment are as important as the hardware they manage.
- SAP workloads require experienced providers. Do not take compatibility claims at face value; verify experience for your specific region and workload type.
- Migration is where risk is realised. A structured, assessed, staged approach is the only acceptable approach for Production enterprise workloads.
- Ask the ten questions. The answers will tell you more than any sales presentation.

Ready to run your cloud on your terms?

Our engineering team will analyse your current environment and deliver a clear cost, performance, and sovereignty comparison with no obligation.



As a global software solutions and managed services company, EPI-USE Labs helps you to maximise the performance, management and security of your SAP® and SAP SuccessFactors® systems. Our clients tell us every day how we have transformed their business operations. Contact us to find out how we can help you solve your business challenges.