The CIO & IT manager guide to private & hybrid cloud

Six insights for a future-proof IT strategy

YOU FOCUS, WE MANAGE.

EPACT

Introduction

The cloud market has changed significantly over the past few years.

Licensing models have been revised, compliance requirements tightened, and the expectation that everything would migrate to public cloud has not materialised for many organisations — or has led to costs and dependencies that are difficult to justify in hindsight.

CIOs and IT managers reassessing their cloud strategy today are doing so in a more complex context than five years ago. The choice is no longer “cloud or no cloud”, but which model suits which workload, which partner offers the right combination of control and managed services, and how to build an environment that is sustainable both technically and legally.

This ebook is about those trade-offs. Not as an introduction for beginners, but as a reference for those who already have foundational knowledge and are looking for sharper insights: on private and hybrid cloud, security and compliance, migration and management. Each chapter stands on its own. Together, they form **a framework for asking the right questions and better evaluating the answers.**

CHAPTER 1

Why private cloud is still relevant for modern IT organisations

For years, a persistent narrative has been told in the IT sector: public cloud is the future, and anyone still thinking about private infrastructure is looking backwards. That picture is wrong, and it is costing organisations that blindly follow it both money and control.

Private cloud is not nostalgic. It is a deliberate strategic choice with clear advantages, but also with real costs and responsibilities that should not be underestimated.

The difference that really matters

Public cloud and private cloud are not variants of the same product. They start from fundamentally different premises.

With public cloud, you pay for shared use: the provider manages everything and you scale effortlessly. That is not just a sales pitch. The flexibility is real, time-to-market is fast and you do not need to keep infrastructure expertise in-house. But that structure also brings limitations.

The key pain points: your data leaves your own environment, you have no control over the underlying hardware and the cost structure quickly becomes complex and difficult to predict as your environment grows.

Private cloud runs on infrastructure exclusively dedicated to your organisation, either in your own data centre or with a partner such as Epact. You retain full control over the network, storage, hypervisor layer and data location.

Security and compliance do not depend on the rules of a large external provider, but on choices you make yourself. But to be fair: private cloud requires a higher initial investment in hardware, demands internal or external infrastructure expertise, and offers less elasticity under unexpected peak loads.

When public cloud falls short

At Epact, we believe that private cloud forms the best foundation for many medium to large SMEs — but not for all. Below, we help you assess when that is the case and when it is not. The three moments when public cloud fails to deliver on its promises are always the same.

1

Cost management

Public cloud is cheap to get started with and expensive to stick with. Licensing costs for specific services, the price of managed databases, network transfers, support contracts, egress costs (cloud providers often charge for data leaving their cloud, while incoming data (ingress) is frequently free) — the costs add up.

Organizations that carefully analyze their public cloud costs regularly find that for predictable workloads, they are paying significantly more than necessary. That said, for irregular or unpredictable workloads, public cloud can be cheaper than owning your own infrastructure.

2

Vendor lock-in

The deeper you integrate into one provider's ecosystem, the greater your dependency becomes. Migrating to another environment, or even bringing part of your workloads back, becomes an architectural and budgetary challenge. That dependency also carries a strategic price: if the provider changes its pricing policy or product offering, you have little negotiating room.

For completeness: private cloud can also lead to lock-in if you deeply integrate with the tooling or hypervisor of a single vendor. The question is therefore not whether you avoid dependency, but from whom and to what extent.

3

Data sovereignty and compliance

For organisations in regulated sectors — such as healthcare, government, financial services or legal services — data location is not a minor detail. GDPR and sector-specific regulations such as NIS2 set concrete requirements for where data is stored and processed, who has access to it and what guarantees apply. With public cloud, that assurance depends on contractual clauses.

With private cloud, it is an architectural certainty. This distinction is strongest in sectors with high regulatory pressure. For companies without strict compliance requirements, this argument carries less weight.

Business impact: three concrete advantages

1 Predictable TCO

With private cloud, you know in advance what your infrastructure costs. no surprises on the invoice, no unexpected additional costs during a usage peak.

That predictability is particularly valuable for multi-year planning, provided your workloads are stable and predictable. if not, the reasoning is reversed: you then pay for capacity you do not always use.

2 Low latency

Workloads that require fast response times perform better on infrastructure that is geographically close to the users. This applies to ERP systems, databases and any application where milliseconds are noticeable.

For workloads where latency is less critical (batch processing, archiving, test environments) this advantage is less decisive.

3 Maximum control

From hypervisor configuration to firewall rules and backup policies: in a private cloud environment, you (and/or your IT partner) are the sole decision-maker. That makes audits simpler, incident response faster and architectural choices more autonomous.

That control is also a responsibility: you bear the burden of updates, patches, monitoring and capacity planning yourself, or you delegate it to a partner who takes it on in a managed services model.

The hypervisor layer: a strategic choice that is postponed too long

Those thinking about private cloud typically focus on architecture, data location and management model. But there is a choice that lies structurally beneath all of that, and one which many organisations make too late: which platform drives your virtualisation layer?

For years, that answer was obvious: VMware. But following Broadcom's acquisition of VMware and the significant changes to the licensing model, costs for many organisations have risen sharply in a short period of time.

For CIOs and IT managers, this is therefore no longer a purely technical question. It directly affects the TCO calculation that makes private cloud attractive. Open source alternatives such as Proxmox offer a fully-fledged alternative.

If you want to learn more about this, our ebook comparing VMware and Proxmox makes for interesting reading.

Bekijk ons Ebook

CONCLUSION

Private cloud is not the opposite of innovation. It is the infrastructure layer that enables your organisation to drive innovation on your own terms, provided you have the scale, the stability and the compliance needs to justify it.

For CIOs and IT managers thinking about their cloud portfolio, the question is not “public or private cloud?”, but “which workloads fit where?” A well-considered private cloud environment at the core of your strategy, supplemented with public cloud capacity where that makes sense, gives you the best of both worlds. Not as a compromise. But as a deliberate architectural choice.

CHAPTER 2

Hybrid cloud: bridge between legacy & future-proof IT

The reality of most IT environments is not one of a clean slate. There are legacy systems that support critical processes and cannot simply be replaced. There are applications built for on-premise infrastructure. There are compliance requirements that dictate where certain data may reside. And at the same time, there is the business expectation to operate flexibly, quickly and scalably.

Hybrid cloud is therefore a deliberate — and often necessary — architectural choice that combines continuity, flexibility and control in one coherent model.

What hybrid cloud concretely means

A hybrid cloud environment combines your on-premise infrastructure and, where appropriate, public cloud resources into one integrated whole. The key lies not in the technology itself, but in the integration: workloads, data and identity management should ideally be able to move seamlessly between environments, based on what is most logical functionally and/or strategically.

In practice, this means you can keep your business-critical data and applications on a private, controlled environment, while simultaneously scaling to public cloud capacity for peak moments or experimental workloads. A best of both worlds, in a sense.

The best of both worlds, without the pitfalls

The strength of hybrid cloud lies in the freedom it provides. No forced choice between the flexibility of public cloud and the control of private infrastructure. No dependency on a single provider for all workloads. No sudden migrations that introduce operational risks.

For CIOs and IT managers, hybrid cloud is the model that most closely aligns with the reality of their organisation: gradual evolution rather than disruptive transition, with an architecture that scales as needs evolve. It is not the easiest path, but it is the most responsible one.

What hybrid cloud delivers for your organisation

The business value of hybrid cloud is most tangible in three areas.

1

Cost discipline without capacity anxiety

You pay permanently for what you structurally need, and temporarily switch to public cloud capacity when that makes sense.

This avoids two classic pitfalls simultaneously: the overcapacity that traditionally characterises on-premise environments, and the uncontrolled cloud spend that arises when public cloud is the only option.

The result is a cost structure you can plan and defend, even to a CFO who wants to know where the IT budget is going.

2

Operational resilience

A hybrid architecture is inherently more redundant than a single environment. Critical processes can be virtualised across multiple locations.

If an on-premise component fails, the cloud layer takes over. If a cloud service is temporarily unavailable, core processes continue on the private infrastructure.

3

Speed where it counts

Teams that want to try out new applications, temporarily need extra computing power for a project, or want to scale quickly for a campaign or launch: they no longer need to wait for infrastructure.

The public cloud layer increases your business agility without affecting the stability of core systems.

Typical use cases

Burst workloads

Some applications experience predictable peaks: the monthly close in an ERP system, seasonal order processing, annual reporting cycles.

Rather than permanently provisioning overcapacity for those peaks, you can opt for additional computing power from the public cloud.

Business continuity and disaster recovery

A hybrid architecture allows back-ups and failover capacity to be distributed across multiple locations and environments.

If an on-premise system fails, critical processes can be taken over by a cloud layer, with minimal impact on the operation of your organisation.

Data localisation

Not all data may or can be stored in the same place.

Sensitive customer data, medical records or financial records can be kept in a controlled private environment, while less sensitive data or analytical workloads are processed elsewhere.

Phased migration

Hybrid cloud is also the most realistic way to plan a migration trajectory.

Applications can be migrated step by step, with the ability to fall back on the existing environment if a workload is not yet ready for the cloud.

Security and compliance: a separate discipline

A hybrid environment does not inherently increase the risk of attacks, but it does require a well-considered approach. Identity and access management, encryption, network segmentation and monitoring must cover the entire environment, not just the on-premise layer.

Because security in a hybrid context deserves its own strategic depth, we cover this extensively in the next chapter.

Other challenges not to be underestimated

Hybrid cloud is a powerful model, but not a simple one. If you approach it with the right expectations, you avoid the most common pitfalls.

1

Management complexity

managing two environments is not twice as complex as one — it is exponentially more complex. Network configurations, security policies, monitoring and updates must be applied consistently across both private and public layers. Without a clear governance structure and tooling that provides that overview, an operational grey area quickly emerges where no one sees the full picture.

2

Integration quality determines everything

A hybrid environment stands or falls with the quality of connectivity between the private and public layer. Slow or unreliable connections hollow out the promise of seamless workload mobility. This requires attention to network architecture, API design and latency monitoring, especially when applications exchange data in real time across environment boundaries.

3

Lack of internal expertise

Hybrid cloud requires expertise that is not always available in-house: knowledge of both on-premise infrastructure and cloud platforms, combined with insight into security, compliance and architecture design. That combination is scarce. Organisations that underestimate this only notice it when an incident exposes the weakest link. An experienced partner that takes over (part of) the management is not a luxury for many organisations — it is a necessity.

4

Costs spiral without governance

The flexibility of the public cloud layer is valuable, but also tempting. Teams that can independently spin up cloud resources do not always do so with an eye on the bill.

Without central cost control and clear rules about who can deploy what, the predictable TCO of hybrid cloud quickly turns into an unpleasant surprise.

CONCLUSION

Hybrid cloud solves a problem that organisations have known for years but rarely name out loud: the gap between the infrastructure they have and the agility they need.

The model offers a realistic path from where you are now to where you want to go, without throwing everything overboard and without getting trapped in the limitations of a single model.

But hybrid cloud is not a product purchase. It is an architectural choice that requires a clear strategy, a competent team or partner, and a governance model that scales with the environment. If you get this right, you win on all fronts: cost control, operational resilience and the freedom to scale when the business demands it. The question is not whether hybrid cloud suits your organisation. The question is whether your organisation is ready to get the most out of it.

CHAPTER 3

Security in hybrid and private cloud: IAM, network and compliance

A hybrid cloud environment brings more attack surface than a single infrastructure. Not because the technology is inherently less secure, but because complexity increases. There are more environments, more connections, more layers where something can go wrong. Security in a hybrid context is therefore not a feature you switch on afterwards. It is an architectural principle that you incorporate from the very beginning.

We examine three domains that structurally deserve attention in hybrid environments: identity and access management, network architecture, and compliance and data sovereignty.

Identity management: the silent complexity of hybrid cloud

Of all the challenges in a hybrid cloud environment, identity management is the least visible but one of the most critical. The question is simple to state: how do you ensure that the right users, with the right permissions, have access to the right systems, regardless of whether those systems run on your private infrastructure, sit in a public cloud, or somewhere in between?

In practice, this means your directory services — Active Directory, Entra ID, or a combination — must work consistently across all environments. Single Sign-On (SSO) and Multi-Factor Authentication (MFA) are not options but minimum requirements. Role-Based Access Control (RBAC) must be managed across environment boundaries, not environment by environment.

Where this goes wrong (and it regularly does) is with exceptions. The application that does not support SSO and therefore requires a local account. The legacy service still running on NTLM authentication. The external partner who temporarily needs access and receives an account that is never cleaned up afterwards.

A hybrid environment forces you to treat identity management as an architectural principle, not an operational task. That starts with an identity audit before the migration and is anchored in a governance model that tracks who has access, why, and until when.

Network architecture: the connection determines the value

A hybrid environment is only as good as the connection that holds the parts together. Traditional network architectures are built for a world where everything sits within the corporate network. In a hybrid context where users work everywhere, applications run in multiple environments and data moves between private and public layers, that model is no longer sufficient.

Modern alternatives such as SASE (Secure Access Service Edge) integrate network connectivity and security into one platform. Instead of routing all traffic back to a central point for inspection, security is applied at the point where the user or system is located. The result is better performance, less complexity and a more consistent level of security across all environments.

As a CIO or IT manager building or modernising a hybrid architecture, the choice of the underlying network model is a strategic decision best made early in the trajectory, not only once performance or security issues arise.

Three minimum requirements for the network layer in a hybrid environment

Encryption of data in transit
All communication between private and public layers must be encrypted, including on internal connections that seem “secure enough”.

Network segmentation
Different environments, applications and user groups belong in separate network segments with explicit access rules. This limits the impact of a security incident to the affected segment.

End-to-end monitoring
Network monitoring that stops at the boundary of the on-premise environment creates blind spots. Monitoring must cover the entire hybrid environment, with central correlation of events across all segments.

Compliance and data sovereignty: structural guarantees vs contractual promises

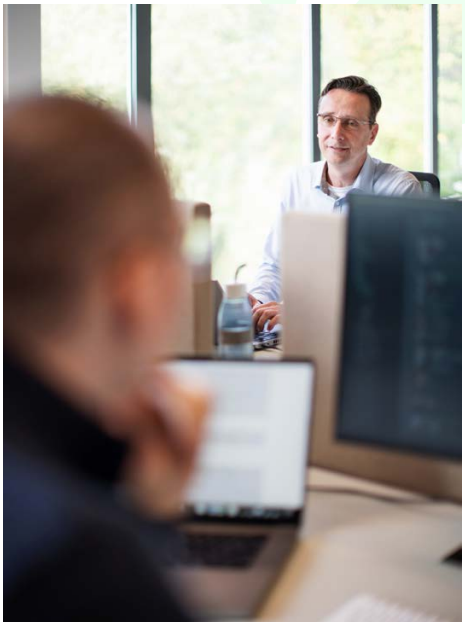

For compliance — whether it concerns GDPR, NIS2 or sector-specific regulations — it is crucial that the hybrid architecture is traceable and auditable. Which data is where? Who had access? How is the data flow monitored?

NIS2 adds an additional layer for organisations in essential sectors: the directive sets concrete requirements for the security of the entire supply chain, including cloud service providers. As CIO, you are therefore not only responsible for your own environment, but also for the choices of the partners that form part of your IT chain.

Data sovereignty deserves particular attention. With a Belgian or European cloud hosting partner, control over data location is structurally guaranteed: architecturally, not merely contractually. That distinction has become relevant in light of the Schrems II ruling and the ongoing tension around the extraterritorial reach of American legislation such as the CLOUD Act.

Cloud services from American providers physically running in Europe do not offer absolute protection against access by American authorities. For organisations with sensitive data or strict compliance requirements, this is not a theoretical risk but a real legal issue that must be explicitly considered when choosing a cloud partner.

CONCLUSION

Security in a hybrid cloud environment is not a technical endpoint but an ongoing process. Organisations that do it well treat IAM, network architecture and compliance not as three separate projects, but as three dimensions of one coherent security strategy — one that also scales as the environment evolves.

The partner you choose for managing your hybrid or private cloud environment is also your partner in that security strategy. The quality of that collaboration — the transparency, the certifications, the local presence — helps determine how solid that foundation is.

CHAPTER 4

Managed private & hybride cloud: what you need to know as a CIO or IT manager

There is a paradox in the way many organisations manage their cloud infrastructure.

They invest in a private or hybrid environment to gain more control, but then spend a disproportionate share of their IT capacity on the day-to-day management of that environment. Rolling out updates, following up on alerts, monitoring performance, handling incidents, validating back-ups: these are tasks that are essential, but add little strategic value.

Managed cloud resolves that paradox. Not by taking away control, but by delegating the management of the infrastructure layer to a partner with the right expertise, so that your internal team can focus on what truly matters.

The real difference: capacity vs depth

For most medium-sized IT organisations, the choice between managed and self-managed cloud is not a matter of principle, but a matter of capacity. Not whether it can be done internally, but whether it can be sustained internally.

Infrastructure management requires current knowledge across a wide range of domains: hypervisors, storage, networks, security, monitoring, compliance. Combining that breadth with sufficient depth to resolve incidents quickly and correctly (including at night and on weekends) is structurally unachievable for a small IT team. Not because the people are not competent, but because the responsibility is too broad.

An additional dimension that is often overlooked: the learning curve during platform transitions. Organisations migrating from VMware to an alternative hypervisor layer — a trajectory that more and more CIOs are deliberately considering due to sharply increased licensing costs — underestimate how much internal knowledge is required. A managed partner who has guided that transition multiple times and manages the new environment daily significantly shortens that curve and reduces operational risk in the vulnerable period just after migration.

The most tangible impact of managed cloud on an IT team is not the time savings on paper, but the peace of mind it provides. When monitoring, back-up validation and incident response have been transferred to a partner, the daily workings of the team change fundamentally. Projects that have been on the backlog for years come back into sight. The energy that went into reactive management can be redirected towards migrations, modernisation or the rollout of new services.

KPIs that matter

A good managed cloud contract is not a vague agreement about “good service”. It is a set of measurable commitments that are monitored and reported.

Uptime and availability

An SLA of 99.9% sounds high, but means in practice less than nine hours of downtime per year. Know what is in your SLA, and also know what falls outside it: planned maintenance, force majeure, incidents caused by the client itself. These nuances determine the real value of the guarantee.

Performance

Are the agreed performance parameters being met? Are deviations proactively signalled and corrected, or only reactively after a complaint? A good managed service partner reports structurally, not only when something goes wrong.

Security

How quickly are patches rolled out? How are security alerts responded to? Are penetration tests and vulnerability scans scheduled? The maturity of your managed service partner in terms of security also determines the maturity of your own environment.

Three absolute points of attention for CIOs and IT managers

1 The gap between SLA and actual responsiveness

Evaluate a managed service partner not only on the written promises of the SLA, but also on practice. Ask concrete questions: how quickly is an incident really picked up? Is there a local presence or does everything work remotely? Are escalation procedures clear and documented? What do existing customers say? Local support — someone you can call and who can come on-site if necessary — is a distinguishing factor that is not visible in most SLAs, but genuinely makes a difference in crisis situations.

2 Exit strategy and portability

A question rarely asked when choosing a managed partner: what if we ever want to switch? Vendor lock-in is not only a risk with public cloud. Managed private cloud can also lead to dependency if the managed environment relies on proprietary tooling or closed platforms. An environment built on open standards gives you structurally more portability than one tied to the toolset of a single vendor. Make portability an explicit requirement when selecting a platform and partner, not an afterthought.

3 Transparency about the underlying infrastructure

Managed cloud does not mean that as CIO you need to know nothing about the layer being managed on your behalf. It means you do not need to deal with it daily. That is an important distinction. Ask every managed service partner for insight into the underlying infrastructure: where is your data physically located? On shared or dedicated hardware? In which data centre? What redundancy is built in? And how does that relate to your compliance requirements? A partner that offers that transparency — and ideally also the associated certifications such as ISO 27001 — gives you the basis to answer those questions to your own management and auditors.

CONCLUSION

The ROI of managed cloud is real, but requires an honest calculation. On the cost side is the partner's invoice. On the return side are the internal hours freed up, the higher uptime that guarantees operational continuity, the faster incident response that limits damage, and the risk that is not realised because monitoring and patching are performed structurally.

For most organisations, the financial business case is positive. But the value of peace of mind and operational focus is at least equally real, even if it does not fit in a spreadsheet.

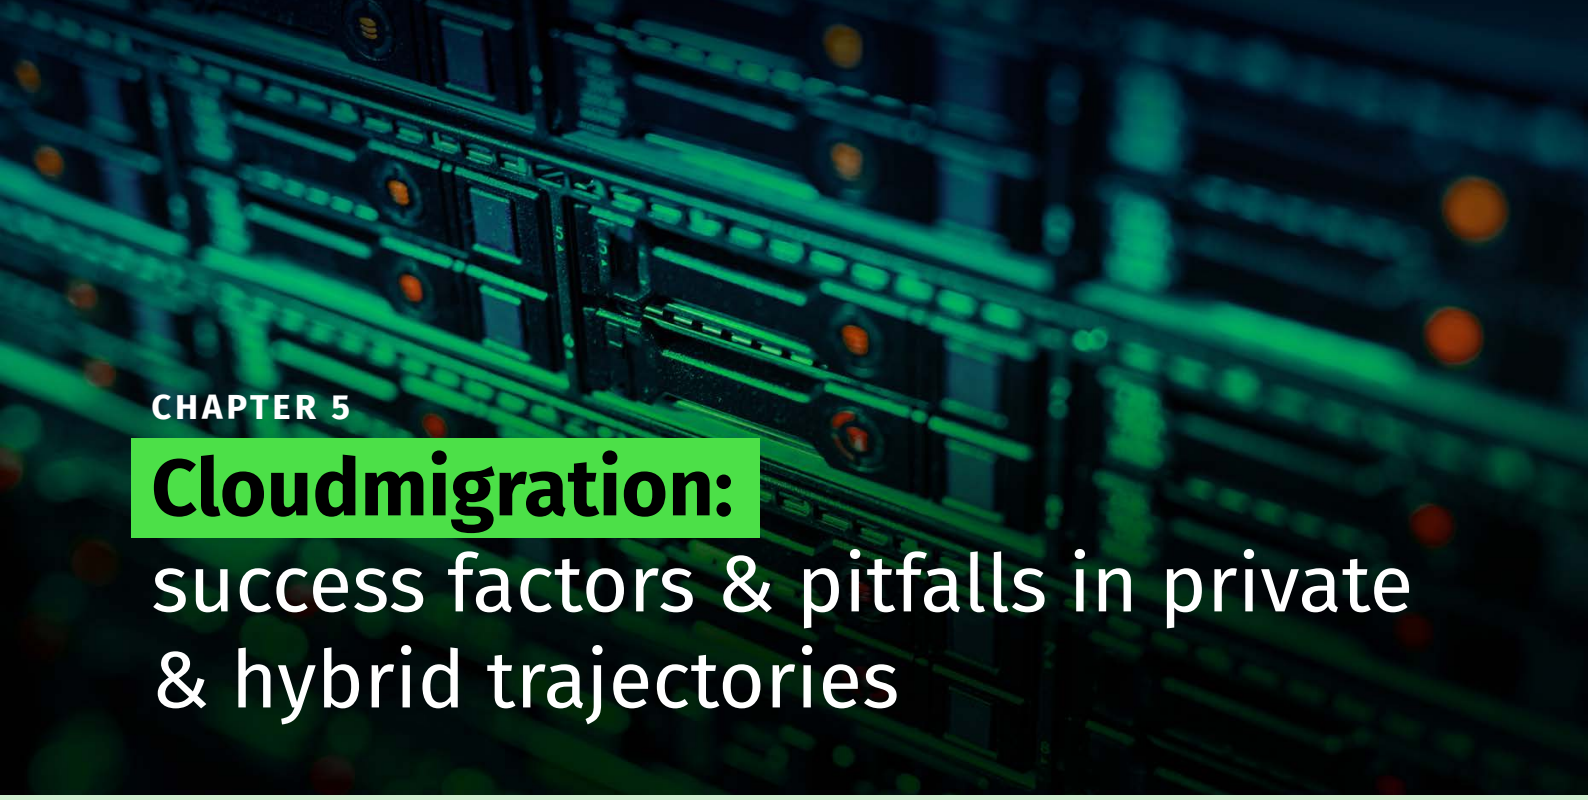

CHAPTER 5

Cloudmigration: success factors & pitfalls in private & hybrid trajectories

A cloud migration project succeeds or fails largely before the first server is moved. The technical execution is relatively manageable. What makes migrations complex are the dependencies that nobody had fully documented, the applications that behave differently than expected in a new environment, and the alignment between IT and the rest of the organisation that was insufficiently prepared.

A structured approach is no guarantee of a perfect migration, but it is the most reliable means of managing the risks.

1 inventory and preparation

The first and most underestimated step of any migration trajectory is a thorough inventory of the existing environment. That sounds obvious, but practice shows that most IT environments are less well-documented than assumed.

Which servers run which workloads? Which applications communicate with which other systems? What dependencies exist with external services, on-premise hardware or specific network configurations? Are there applications with hardcoded IP addresses or server names that will immediately break during a migration?

A point of attention often missing in the inventory phase: software licences and hypervisor dependencies. Some software — in particular Microsoft products such as Windows Server and SQL Server — has licence terms that depend on the underlying virtualisation platform. A migration from VMware to an alternative hypervisor can therefore have licensing consequences that only become visible when the invoice arrives or the software audit takes place. Map this out before architectural choices are made, not after.

A good inventory also maps how critical each workload is. Which systems are business-critical and cannot have downtime? Which can be temporarily taken offline? That classification determines the order and approach of the migration.

2

architecture choices and planning alignment

Based on the inventory, the target architecture is determined. Which workloads go to the private cloud? Which are set up in a hybrid model? Are there systems that are better candidates for rationalisation or replacement rather than migration?

Planning alignment transcends the IT department. Business stakeholders need to know when certain systems will be temporarily less available. External vendors or partners who depend on certain systems must be informed. Migration windows must be aligned with the business cycle: planning a migration during the busiest month of the year is a risk that is easy to avoid.

3

integration with existing systems

Migration to private or hybrid cloud rarely means a complete break with the existing environment. There are legacy systems that continue to run, on-premise hardware that is not being replaced, and external connections that must continue to work.

The integration of the new cloud environment with what remains requires attention to network configuration, authentication and access management, and the consistency of monitoring and logging across the entire environment. A hybrid environment that is not coherently monitored creates blind spots that only become visible during an incident.

4

testing, downtime management and rollback

No migration should be put into production without thorough testing. Performance tests validate whether the new environment can handle the expected load. Functional tests confirm that applications behave as expected. And for critical systems, there must be a documented rollback scenario: a plan to return to the original environment if the migration causes unexpected problems.

Downtime management is a separate discipline. For systems that cannot go offline, techniques such as live migration or parallel operation of old and new system during a transition period are available. The choice of approach depends on how critical a system is and the technical capabilities of the applications involved.

5

documentation, knowledge transfer and training

A migration is only fully complete when the internal team understands the new environment and can manage it independently — at least at the agreed level. Documentation of the architecture, configuration choices and management processes is not a luxury, but a requirement for operational continuity.

Training and knowledge transfer deserve explicit time in the project planning. Not a half-day course, but a structured trajectory in which the internal team is actively involved in configuration and setup, can ask questions, and builds familiarity with the new environment before it goes into production.

The common thread: communication and aligning expectations

The pitfalls in cloud migrations are rarely purely technical. They result from unclear expectations, insufficient communication between IT and business, and a planning that was too optimistic about the complexity of the dependencies. A partner who not only handles the technical execution, but also assists with preparation, alignment and knowledge transfer, makes the difference between a migration that goes unnoticed and one that is remembered for a long time.

CONCLUSION

A cloud migration is not a project, it is a change trajectory. The technology is manageable. The complexity lies in the dependencies, the communication and the expectations you have or have not aligned in advance.

As CIOs or IT managers, it is better not to follow this trajectory alone. Choose a partner who masters the technical execution, but also understands how an organisation works: how to create buy-in, how to inform stakeholders, and how to choose a migration window that impacts the business as little as possible.

The best migrations are the ones nobody notices. Not because nothing has changed, but because everything was so well prepared that the transition felt natural. That is the goal. And it is achievable, with the right approach.

CHAPTER 6

Business continuity & compliance in cloud hosting: what you need to know as a CIO or IT manager

Cloud hosting is not just a technological choice. It is a risk decision. Which risks do you accept? Which do you transfer to a partner? And which are so critical that they must be structurally mitigated, regardless of cost?

The questions that matter are not how many IOPS the storage achieves, but what happens if a data centre fails, who is liable in the event of a data breach, and how quickly critical processes can be restored after an incident. CIOs and IT managers who assess cloud hosting through that risk lens consistently ask themselves and their partners better questions and consistently receive better answers.

Redundancy and failover: architect for the worst case

High availability begins with architecture, not with the SLA. Redundancy at every level — power supply, network connectivity, storage, compute capacity — ensures that the failure of one component does not lead to a total outage of the environment.

For critical workloads, that goes beyond internal redundancy. Failover to a second location or environment, activated automatically or manually, ensures continuity even in the event of a serious incident at the primary location. RTO and RPO are the parameters that set the tone here. But it is the business context that determines what is acceptable, not the IT department alone.

Backup strategy: beyond a daily copy

A back-up that has never been tested is not a back-up. In practice, it regularly turns out that back-ups are made, but that the recovery process has not been validated. A back-up is only valuable if the recovery works, is fast enough, and restores the right data point. A robust backup strategy rests on three pillars: the 3-2-1 rule (three copies, two media, one offsite), immutability (back-ups that cannot be modified or deleted, even by an attacker with access to the production environment), and regular restore tests, documented and monitored.

The immutability of back-ups is not a nice-to-have in the context of increasing ransomware attacks — it is a hard requirement. NIS2 also makes the expectation around immutable back-ups explicit for organisations in essential and critical sectors. But even outside those sectors, it has become an architectural minimum.

GDPR, NIS2 and data sovereignty: beyond the legal checkbox

1 GDPR

GDPR has been in force for years and yet organisations still approach cloud hosting without checking exactly where their data ends up, how long it is retained and who has access to it.

That negligence is a risk — not only because fines are real, but because a data breach also causes reputational damage that is difficult to repair.

2 NIS2

NIS2 adds an additional layer for organisations in essential sectors. The directive sets concrete requirements for risk management, incident reporting and supply chain security — including cloud service providers.

That last point is crucial: as CIO, you are not only responsible for your own environment, but also for the choices of the partners that form part of your IT chain. You must be able to demonstrate that your hosting provider meets the requirements — and that requires more than a contractual clause.

Ask for certifications (ISO 27001 is a widely used reference point), audit rights, and the concrete procedures in the event of a security incident.

3 Data sovereignty

Data sovereignty is the third dimension, and one that is growing in importance. With a Belgian or European cloud hosting partner, control over data location is structurally guaranteed: architecturally, not merely contractually.

That distinction has become relevant in light of the Schrems II ruling and the ongoing tension around the extraterritorial reach of American legislation such as the CLOUD Act. American cloud services physically running in Europe do not offer absolute protection against access by American authorities.

For organisations with sensitive data or strict compliance requirements, this is not a theoretical risk but a real legal issue that must be explicitly addressed when choosing a cloud partner.

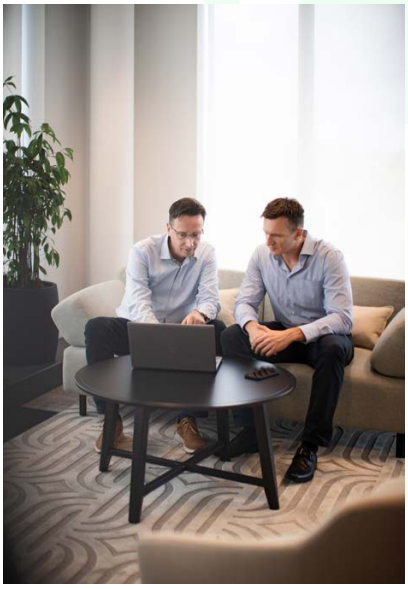

CONCLUSION

The most valuable mindset shift you can make as a CIO around cloud hosting is the transition from a technical to a risk-based approach. The question is not only “does it work?”, but “what if it doesn’t work?”, “how quickly can we be operational again?” and “can we demonstrate that we acted correctly?”

Organisations that have answered those questions and translated the answers into an architecture, a management model and a set of agreements with a partner, are not only better protected technically. They are also better prepared for the strategic, legal and reputational dimensions of a world in which IT continuity is increasingly a boardroom topic.

CLOSING REMARKS

Private and hybrid cloud are more relevant than ever before

This ebook is not an argument for one specific technology or one particular model. It is an invitation to ask the right questions before decisions are made.

Which workloads belong in a private environment? Where does hybrid cloud add value? What are the real costs and risks of your current setup? And which partner helps you not only with the technical execution, but also with the thinking?

At Epact, we guide organisations through all these questions: from intake and sizing to implementation, management and compliance.

With our eCloud environment, built on the open source Proxmox platform, we offer a fully-fledged, scalable and affordable alternative for those who want to build their cloud strategy without depending on a single large vendor. Managed from Belgium, by a team that knows you and that you can call. Not as an off-the-shelf product. But as a tailored solution, with the long-term relationship that belongs to infrastructure choices that truly matter.

Curious what this concretely means for your environment? Get in touch! We are happy to think along with you.

Want to know whether a private or hybrid cloud fits your cloud strategy?

The question is no longer “public or private cloud?” but “which workloads belong where?” A well-considered private cloud environment as the core of your strategy, supplemented with public cloud capacity where it makes sense, gives you the best of both worlds. Not as a compromise, but as a deliberate architectural choice.

Get in touch

Contact us for a no-obligation assessment or technical deep dive. That way, you’re not taking a leap in the dark, but making an informed decision.

Christof Ugau
Managing Partner

christof.ugau@epact.be

Leo Chang
Business Developer

leo.chang@epact.be

EPACT