Beyond Lift and Shift: Getting Real Value from Microsoft Azure

Why moving to Azure is not the same as getting value

A lot of organisations move to Azure for sensible reasons. They want to get away from ageing servers, reduce the risk of hardware failure, and make it easier to scale when the business changes. That is a good start. It is not the same thing as getting value.

The gap shows up quickly. A workload can move into Azure and still leave the business with surprise bills, unclear security ownership, slow recovery, and an environment that only one person understands. The location changed. The outcome did not.

For small and medium businesses, that is the real issue. Most are trying to keep the business moving with a mix of on-prem systems, Microsoft 365, and Azure workloads, and they need Azure to support that reality rather than add another layer of complexity. Value only appears when the environment is managed with clear priorities. In practice, those priorities are security foundations, recoverability and continuity, cost control, and targeted modernisation.

That is why it matters whether you are still on-prem, mid-migration, or already in Azure. If you are still on-prem, this will help you set better migration priorities before you move. If you are mid-migration, it helps you avoid carrying old weaknesses into Azure. If you are already in Azure, it helps you check whether the environment is actually delivering the outcomes you expected.

The wider market shows why this matters.

28%

of UK SMEs that migrated to the cloud say they have yet to see any benefit

42%

of companies say they are fully achieving their expected cloud outcomes

10%

of companies say they are capturing cloud value at scale

What Real Value Means in Azure

What lift and shift is good for

Lift and shift has a place. It is often the quickest way to make progress when you need change without a full redesign. It can reduce dependence on ageing hardware, buy time, and create breathing room when there is pressure to move quickly.

For many SMBs, that first step is sensible. If the choice is between doing nothing and moving a stable workload into a properly managed cloud platform, lift and shift can be the right bridge.

It also helps when you want to de-risk legacy infrastructure in stages. Moving to Azure can reduce the burden of physical server maintenance and shift some spending from capital outlay to more predictable operating cost.

What lift and shift does not do is solve the harder problems by default. It does not automatically improve security posture. It does not guarantee reliable recovery. It does not impose cost discipline. It does not make ownership clearer. It simply moves the workload.

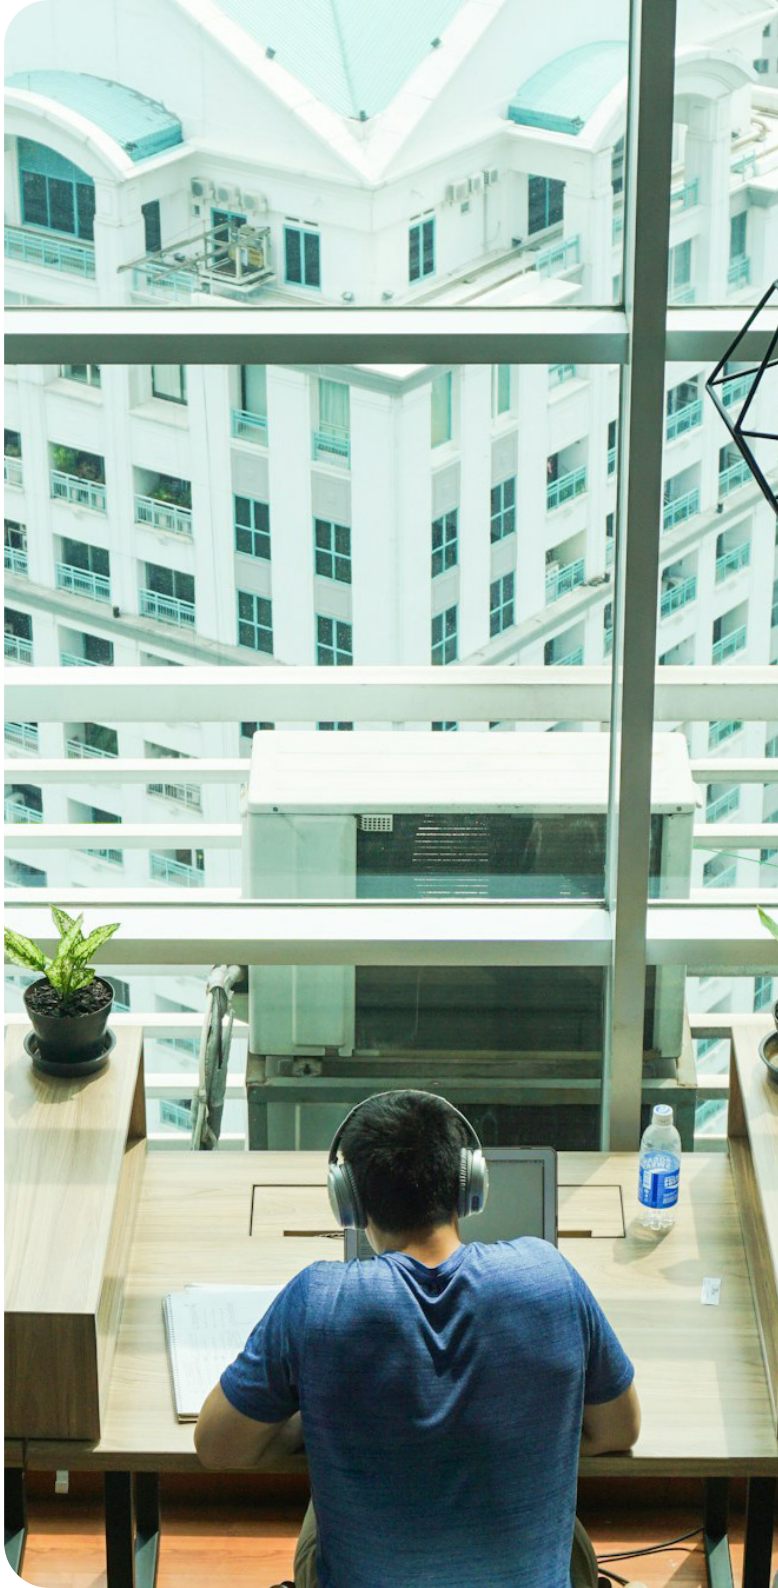

The main migration approaches

Lift and shift is only one way to move to Azure. In practice, organisations usually choose from a small set of approaches.

A **rehost** approach is the closest to lift and shift. You move the workload with limited redesign so you can get off old infrastructure quickly.

A **replatform** approach makes targeted changes along the way. That might mean improving backup, changing hosting, or adjusting the way a service is run without rebuilding everything.

A **refactor** approach goes further. The application or service is reworked so it fits the cloud better, usually when the old design is too limiting or costly to keep.

None of these is automatically right or wrong. The choice depends on the age of the system, the risk it creates, the skills available, and how much change the business can absorb. The important thing is to avoid treating every workload the same way. Some systems need speed. Some need stability. Some need a proper redesign. Good Azure decisions start with that distinction.

Where the value gap appears

This is where many organisations get caught out. They assume the cloud itself delivers the outcome. In reality, value depends on how Azure is configured and how it is run day to day.

That means the same platform can produce very different results.

One business moves to Azure and ends up with stronger identity controls, tested backups, clear alerting, and costs leadership can understand. Another moves the same way but keeps inherited admin sprawl, poor visibility, and unmanaged resource growth. Both are in Azure. Only one is getting value.

The common failure points are practical: unclear ownership and decision-making, permissions that grow over time, weak monitoring and alerting standards, cost controls that only appear after finance notices a spike, and recovery plans that have never been tested properly.

Cloud spend waste remains a major issue across organisations, with hybrid estates now the norm and waste still showing up in a meaningful way. That is a reminder that adoption alone does not create value. Governance does.

The same pattern shows up in cloud outcome studies. Accenture found that more organisations are increasing cloud activity, but many still fail to realise the expected benefits. More work does not automatically mean better results.

A simple definition of value

For an SMB, Azure value is not a technical score. It is a business result: reduced risk, faster recovery, predictable spend, and smoother operations. That definition matters because it changes the conversation. The right question is not, “Are we on Azure?” The right question is, “Is Azure helping the business run better?”

A practical test is to ask:

1. Do we know who owns cloud security and access decisions?
2. Can we restore critical systems in a timeframe the business can live with?
3. Can we explain next month’s cloud spend before the invoice arrives?
4. Do we know which workloads matter most and why?

5. Are we spending time improving the environment, or just reacting to it?
5. If the main Azure administrator was away, could someone else safely run the environment?
5. Do we have enough visibility to spot problems early?
5. Have we decided what should stay, move, improve, or be retired?

If the answer to several of those questions is “not really”, the organisation probably has cloud location but not cloud value.

Value also changes over time. A migration can be the right decision in year one and still leave major gaps in year two if nobody keeps checking the environment. Azure needs ongoing attention. Someone has to review access, test recovery, watch spend, and decide when a workload needs attention. If those habits are missing, value will drift even if the original migration was well executed.

That leads to the next question. How do you tell whether Azure is actually working as well as it should?

The Azure Value Scorecard: How to Measure Progress

The outcomes that matter

Before you measure anything in Azure, define the outcomes you care about.

Most SMBs do not need a perfect cloud governance framework. They need a practical view of what matters most: how long the business can tolerate downtime before it starts to affect revenue, delivery, compliance, or client trust; what level of risk is acceptable given the sector, client base, and data handled; what kind of costs leadership can plan for; and which systems are critical enough to deserve active management rather than passive hope. These questions are easier to answer when they are linked to real business scenarios rather than technical jargon.

A finance lead may care less about the technical difference between backup and replication and more about whether payroll can run if a key platform fails. An operations director may care less about the name of a security control and more about whether a compromise would stop the business for a day, a week, or longer.

The Azure value scorecard

A scorecard is useful because it turns a vague question into a manageable conversation. It does not need to be complex. It needs to tell you whether the environment is moving in the right direction.

Most organisations only need a few views of the environment. That might mean how quickly they can recover key systems, how many people have high-level access, whether critical workloads are being tested, and whether spend is drifting away from plan. The exact measures can vary, but the principle stays the same. Use a small set of signals that leadership and IT both understand.

If the measures are too detailed, people stop looking at them. If they are too generic, they do not tell you anything useful. A good middle ground is usually enough.

The point of the scorecard is not complexity. It is to answer one question: are we managing the environment, or just hosting it?

For most SMBs, a monthly or quarterly review is enough if the right people are in the room. Keep the discussion consistent. Ask what changed, why it changed, and what needs to happen next. That is enough to keep the environment under review without turning it into a reporting exercise.

Area	What to look for	Why it matters
Recovery	Agreed recovery targets for key systems, plus evidence they can be met	Shows how quickly the business needs to get back up and running
Restore testing	Proof that critical systems have been restored successfully	Confirms backups are real, not just assumed
Access control	Who has admin or privileged access, and how often it is reviewed	Reduces risk from excess access and forgotten permissions
Security basics	MFA, alerting, and logging on the systems that matter most	Gives early warning and reduces avoidable exposure
Cost visibility	Budget, actual spend, and the main drivers of variance	Helps leadership see whether spend is controlled
Ownership	Named owner for key systems, decisions, and follow-up actions	Stops important work becoming nobody's job

Ownership and decision-making

A scorecard only helps if someone is clearly responsible for acting on what it shows. Leadership usually sets the tolerance for risk, cost, and downtime. Internal IT or an external partner then translates that into settings, reviews, and fixes. In smaller teams, the same person may cover more than one of those jobs. That is fine as long as the responsibility is visible.

What matters most is that the important areas are not left vague. Someone needs to keep an eye on access, recovery, cost, monitoring, and changes to the environment. If those areas are shared, the handover needs to be clear. If one person owns them all, the business should still know who that is.

If no one is clearly accountable, the environment will drift.

Security Foundations That Unlock Confidence

Identity and access first

In Azure, identity is the control plane. In plain English, that means access management matters more than almost anything else because it decides who can do what.

Multifactor authentication is one of the simplest ways to reduce account compromise risk. It should be in place for admin access at minimum, and ideally for every user. That alone will not solve security, but it removes a common route in.

At minimum, the access model should be straightforward. People should only have the access they need, admin accounts should be separate from everyday accounts, and temporary permissions should not be left in place longer than necessary. Shared admin accounts, broad privileges, and informal workarounds are where problems usually start.

If the environment still relies on “we all know the password” for privileged access, that is not control. It is exposure.

Permission sprawl and unknown admins

Permission sprawl happens gradually. A supplier needs temporary access. A project needs elevated rights. Someone is on holiday. Someone else is troubleshooting. Nobody goes back and cleans up the old access.

Six months later, the environment has more people who can make high-risk changes than anyone thought.

A practical access review can be done with four questions:

1. Who has access?
2. What can they do?
3. Why do they need it?
3. Should it be removed or reduced?

That is enough to start.

You do not need a perfect identity project to make progress. You need a disciplined habit. Review privileged accounts, confirm the business reason for each one, and remove anything that no longer has a clear purpose.

This matters even more for SMBs because small teams often assume familiarity equals control. It does not. A small environment can still have a large amount of hidden risk if access has grown informally over time.

Visibility and alerting

You cannot manage what you cannot see.

At a high level, the things worth logging and alerting on are the events that matter to business continuity and security. That includes:

- **Admin sign-ins**
- **Failed login spikes**
- **Privilege changes**
- **Changes to security settings**
- **Service outages that affect critical systems**
- **Unusual resource creation or deletion**
- **Backup failures**

A good alert is not just noisy. It should be actionable.

That means it should answer questions such as: What happened? Why does it matter? Who needs to do what next?

If alerts are vague, ignored, or sent to no one in particular, they will not help. The goal is not more alerts. The goal is fewer surprises.

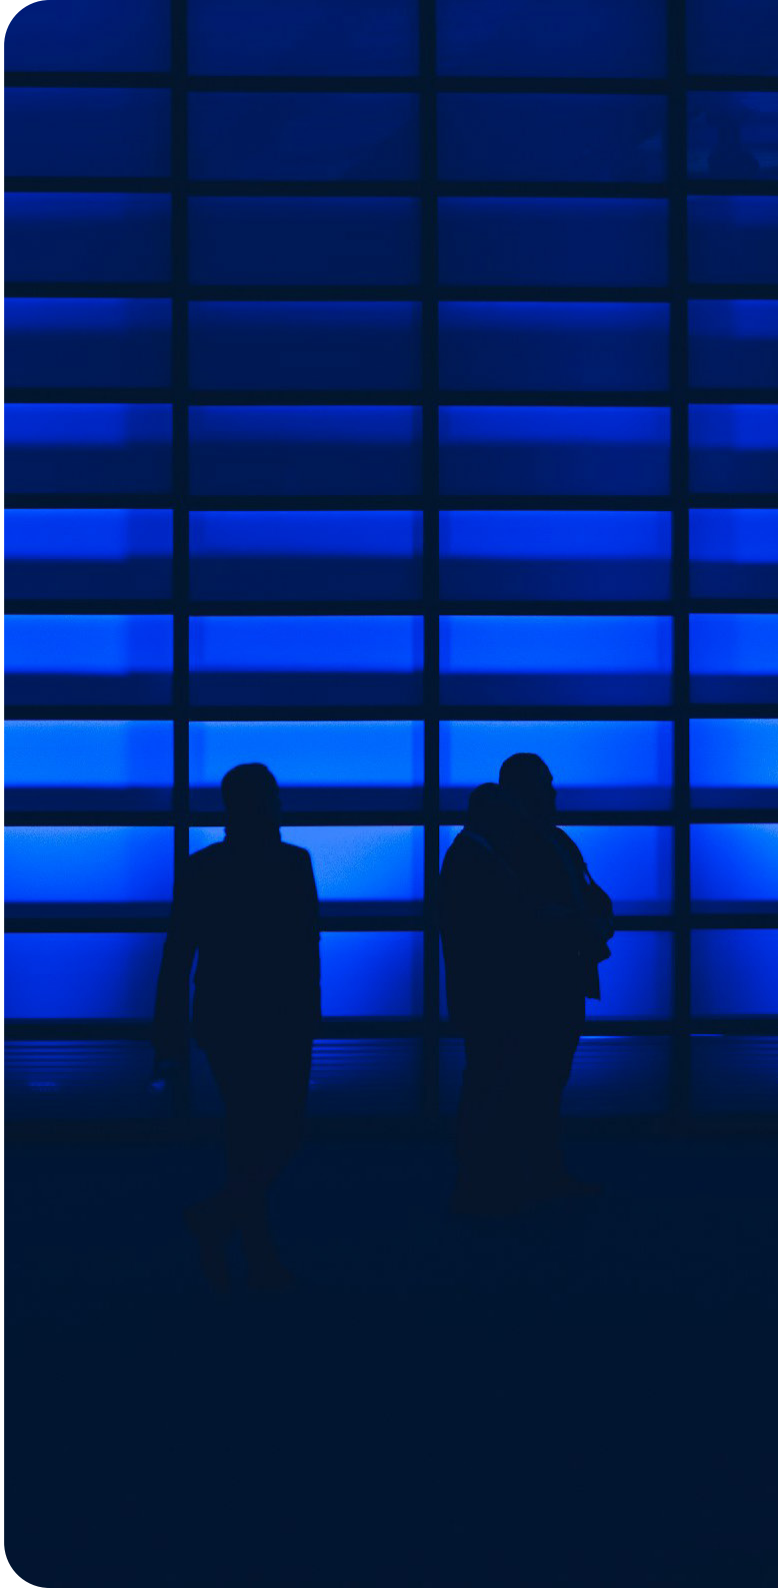

Secure-by-default habits

Security in Azure is not only about tools. It is about habits.

The main ones are simple:

- **Make change control normal**
- **Document who approved what and why**
- **Define who owns vulnerability management**
- **Patch in line with business risk**
- **Separate routine changes from emergency changes**

The UK Government Cyber Security Breaches Survey found that 43% of UK businesses identified a cyber attack in the previous 12 months, with the majority of those businesses reported phishing attempts, and only 25% had a written incident management plan. That combination matters because it shows how often attacks start with human compromise, and how often recovery planning is still informal.

The lesson for Azure is straightforward. Secure the account layer first. Then make sure there is a clear plan for what happens when something goes wrong.

Recoverability and Business Continuity

The real question: how fast can we be working again?

When leaders talk about resilience, the useful question is not whether there is a backup somewhere. The useful question is how quickly the business can get back to work.

That means you need to think about realistic scenarios such as ransomware, accidental deletion, misconfiguration, supplier outage, a regional or platform issue, or user error that breaks access or data. Each scenario affects the business differently.

A short outage might be irritating. A longer outage might stop sales, block client service, delay payroll, or create compliance issues. The right recovery plan takes that into account.

To translate impact into priorities, think in three layers:

People
Who is unable to work?

Process
What business activity stops?

Systems
Which apps and data are needed to keep going?

The recovery target should match the business, not the technology.

Backup and restore, not backup and hope

Backup is not the same as recoverability. Backup copies data so it can be recovered later. Replication keeps another copy of a system or data set in sync for faster recovery. Retention defines how long data is kept. Archive is for data that must be preserved but is not needed for daily operations. Those are different purposes, and confusing them creates gaps.

A good SMB recovery design usually includes backups for critical systems and data, tested restores rather than just successful backup jobs, clear retention settings, documented recovery steps, and ownership for who runs the restore when needed. If a restore has never been tried, you do not know whether it works in practice.

A minimum restore testing cadence should be at least quarterly for critical systems, and more often for anything that changes frequently or supports core operations. The test does not need to be elaborate to be useful. In many SMBs, the most valuable exercise is a simple one: pick a critical system, restore a small but real set of data, and confirm that the business can use it again. If that sounds basic, that is the point.

It also helps to write down what success looks like before the test begins. If the team knows what must be restored, where it will be restored to, who signs off the result, and how long it is allowed to take, the test becomes informative rather than stressful. Over time, those tests build confidence. They also expose weak spots before an incident does.

This is where the ransomware statistics matter.

46%

of ransomware-hit SMBs paid the ransom

20%

of small businesses fail within a year of a cyber attack

Designing for resilience

Resilience improves when you reduce single points of failure.

That does not always mean doubling everything. It often means understanding dependencies well enough to remove avoidable risk.

The main dependencies to document are:

- Identity
- Connectivity
- Critical applications
- Data flows
- Third-party suppliers
- DNS

If users cannot log in, the application may as well be offline. If DNS fails, the service may still be running but effectively unreachable. If the backup is healthy but the restore path depends on a person who is away, that is a weak design.

Documenting dependencies helps you prioritise sensibly. It shows where resilience is about technical design and where it is about process discipline.

Cost Control Without Slowing the Business Down

Why Azure bills surprise people

Azure spend tends to drift for predictable reasons. The usual causes are always-on resources that never get turned off, storage growth that nobody reviews, data transfer costs that were not expected, backup policies that are broader than necessary, forgotten test environments, and duplicated services after a migration. That is why Azure cost and Azure value are not the same thing. Cheap is not always good, and expensive is not always bad. What matters is whether spend matches business need.

Cost management continues to be one of the biggest cloud challenges, and a meaningful share of cloud spend is wasted. That is not a reason to avoid Azure. It is a reason to govern it properly.

The other reason bills surprise people is that cloud cost is easy to defer. A small amount of waste on one server or one test environment may not feel important. Add a few more services, then a few more users, then a migration project, and the total starts to matter. By the time finance spots the drift, it is usually spread across several resources, not one obvious mistake.

Guardrails that stop waste early

Cost control works best when it is built in, not bolted on later. The most useful guardrails are budgets and alerts, tagging and accountability, environment hygiene, and reporting that leadership can understand. Budgets and alerts should tell you when spend is moving outside the expected range. Tags should show who owns what and why it exists. Environment hygiene should remove leftover resources from old projects, non-production systems, and temporary work.

A practical question to ask is what should stay on, what can be turned off, and what must never be turned off without warning. That matters because not every resource can simply be stopped at the end of the day. Some workloads are business critical. Some can be paused. Some should be redesigned entirely. The point is to make the decision consciously rather than by accident.

The easiest way to keep control is to define three simple rules: every resource needs an owner, every meaningful cost needs a reason, and every exception needs a review date. Those rules sound basic, but they stop a lot of waste. They also make conversations with leadership much easier because there is a visible link between spend and business purpose.

Leadership reporting should avoid excessive technical detail. A good report should answer what we spent, what changed, why it changed, whether it was expected, and what we are doing about it. If the answer to those questions is not clear, the report is not helping decision-making.

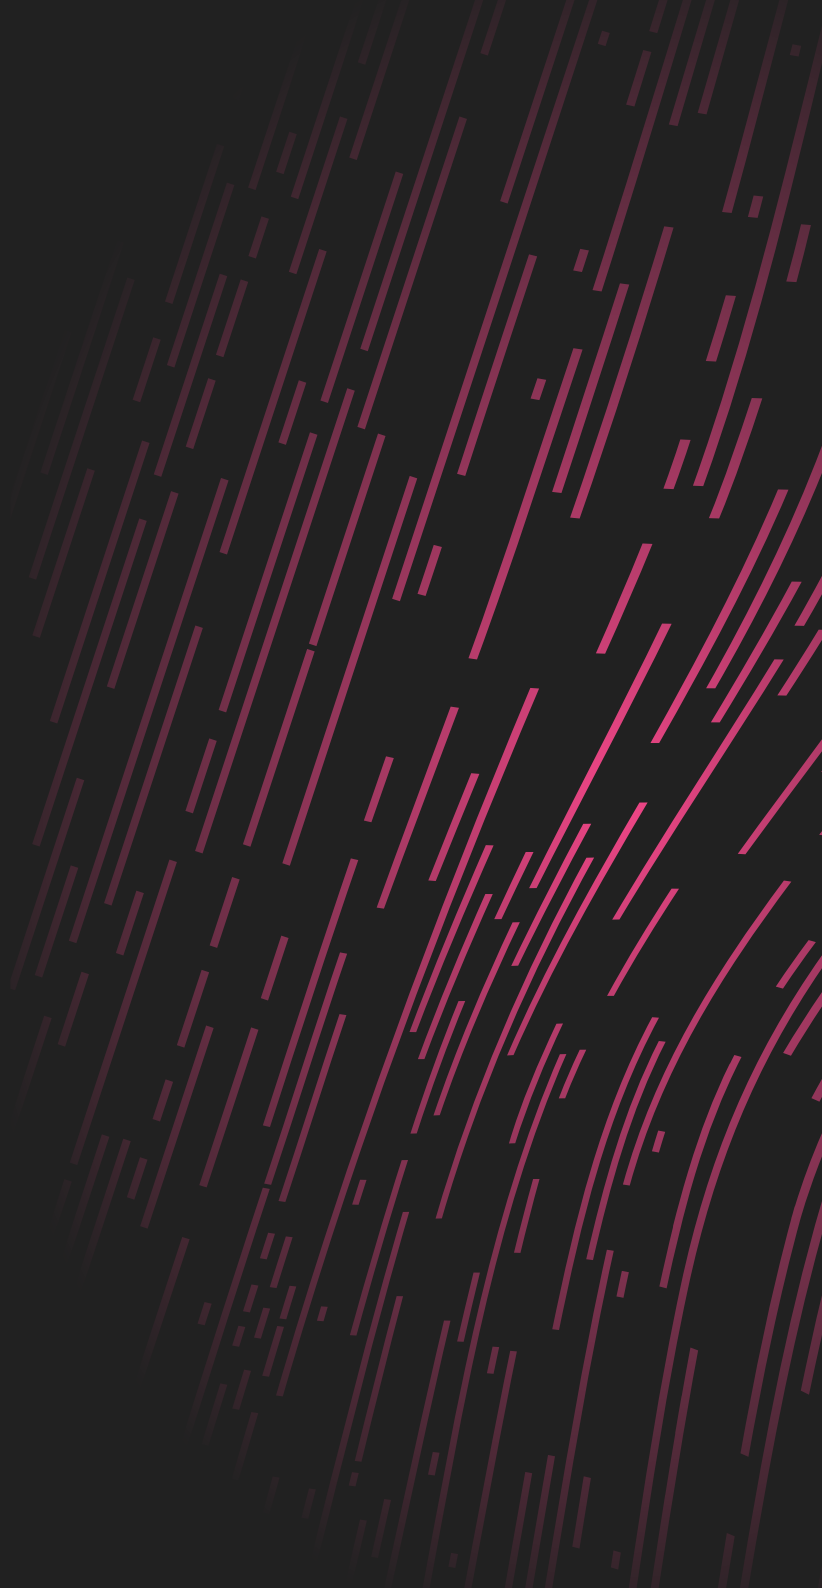

Environment hygiene

A lot of cloud waste hides in plain sight. The usual culprits are test systems left running after a project ends, oversized virtual machines, storage that has grown but is never reviewed, old snapshots or backups that are kept longer than needed, duplicate services created during a migration, and resources attached to systems nobody actively uses anymore. This is where hygiene matters more than heroics. You do not need to redesign every workload to get better cost control. Often the first gain comes from removing obvious waste and standardising how environments are managed.

A simple monthly hygiene review can focus on what is running and why, what is no longer needed, what has grown, what can be resized safely, and what should be shut down outside working hours. The review should be small enough to complete. If it becomes a large project, people will stop doing it.

Tagging and accountability

Tags are not just an admin detail. They are how you keep ownership visible. A useful tagging approach can include owner, department, environment such as production or test, business purpose, and review date. The exact labels matter less than the habit. If no one can tell who owns a resource, it tends to linger. If no one knows why it exists, it tends to be left alone. That is how cloud estates become messy.

Tagging also helps when more than one team touches the environment. Leadership needs enough visibility to see which costs belong to which part of the business. IT needs enough context to work out what can be changed. A partner needs enough information to make sensible recommendations without guessing.

Planning spend for growth and change

Azure spend should also be reviewed whenever the business changes materially. That includes new sites, new applications, acquisitions, seasonal demand spikes, major user growth, and security or resilience improvements that change infrastructure needs.

Do not treat Azure cost as a static annual budget. Review it when the business changes, because the cloud changes with the business. It is also worth reforecasting when a new system goes live or when the business decides to keep a legacy application running for longer than planned. Both decisions affect cost. Both should be visible before they hit the invoice.

What good cost control feels like

When cost control is working, leadership should not need to chase the IT team for explanations. The pattern should be visible before it becomes a problem. Good cost control usually looks like this: there is a budget and someone watches it, major changes are reviewed before they create spend, test and non-production resources are not left running without reason, each major workload has a clear owner, monthly reporting is short, clear, and consistent, and surprises are rare and explained quickly when they happen.

That does not mean spend stays flat forever. A growing business should expect some movement. The goal is not to freeze Azure. The goal is to make sure growth is deliberate, understandable, and connected to business value.

Modernising
What Matters

The 80/20 approach to modernisation

Modernisation should not become a sprawling programme that tries to fix everything at once. The smarter approach is 80/20. Focus on the small number of workloads that create most of the risk, cost, or friction.

A workload usually needs attention when it is hard to support, expensive to maintain, risky from a security or resilience point of view, dependent on skills that only one person has, or blocking productivity or growth. That does not mean every old system should be replaced. It means the business should know where the pain is highest and start there.

Common SMB modernisation paths in Azure

There is no single right modernisation path. For SMBs, the outcome matters more than the architecture label. Some modernisation work improves stability so systems fail less often. Some reduces admin time by removing manual steps. Some strengthens security with better identity and access control. Some improves recovery so the business can resume faster. Some standardises deployments so changes are repeatable. Some reduces support burden by retiring awkward legacy tools.

The goal is not to modernise for the sake of it. The goal is to remove business friction.
A modernisation project should also stay small enough to finish. Pick one workload, one outcome, and one measurable improvement. For example, reduce outage risk for the finance system, cut manual admin time for file access, improve recovery for a critical line-of-business app, or replace a brittle legacy server with a supported service. That is a much more manageable way to work.

Decide what to keep, improve, replace, or retire

A practical decision framework is straightforward. Keep it if it is low risk, low effort, and still useful. Improve it if the system is basically sound but has clear gaps in security, cost, or supportability. Replace it if the current option is too costly, fragile, or hard to secure. Retire it if it no longer adds value.

The key is to be honest. Some systems should not be modernised. They should be retired. Others are fine to keep for now. The value comes from making that decision deliberately.

Turning Insight
Into an Action Plan

Choose priorities based on business impact

The best action plans start with outcomes, not technical tasks. Instead of saying, “We need to improve Azure,” say what improvement you actually want. That might be reducing the risk of compromise, reducing downtime risk, making cloud spend predictable, or removing operational bottlenecks for IT and users.

From there, identify the five business-critical systems that matter most. Do not start with everything. Start with the systems that would hurt most if they failed. For each one, agree what good enough looks like. How available does it need to be? How quickly do you need to recover it? How much data loss is acceptable? Who is responsible for it? Those answers give you a practical starting point.

Examples of Now, Next, Later

The simplest way to organise the work is to separate it into what can happen now, what needs planning, and what should wait. Now is the low-disruption work that reduces immediate risk or waste. Next is the work that needs coordination, approvals, or a change window. Later is the useful work that is not urgent yet.

That keeps progress visible without turning the plan into a heavy programme. It also helps teams stay honest about what can actually be delivered in the near term. Many organisations make progress not by doing everything, but by doing the few things that reduce the most risk.

What the first 90 days should look like

If you want momentum without creating a mammoth programme, start with a narrow review and a few visible fixes. The point is not to tidy everything. The point is to prove that the environment can be brought under control. A sensible first pass usually starts with clarifying who owns Azure security, access, backup, and cost decisions, reviewing privileged accounts and removing anything that no longer has a clear purpose, checking that MFA is enforced on every admin account and ideally for every user, identifying the five systems that matter most to the business, verifying backup coverage and completing one restore test for each critical system, setting budgets and alerts so unexpected spend is caught early, and making sure monitoring goes to an owner who can act on it rather than just to a mailbox. Those actions are unglamorous. They are also the difference between a cloud estate that looks modern and one that actually behaves well.

If the business can see progress inside 90 days, leadership is more likely to back the next round of improvements. It also gives a way to show that cloud governance is not an abstract exercise. It has direct operational value.

How We Purpose Technology
helps you get value from Azure

Most SMBs do not need more cloud services. They need help getting the right outcomes from the cloud they already have. A good MSP can help by spotting risk and waste early, translating technical findings into business decisions, owning the day-to-day operational work that keeps Azure healthy, and keeping access, monitoring, backup, and cost management under control.

Accenture's research suggests that organisations which achieve cloud outcomes are more likely to lean on managed services. That makes sense. Managed support does not replace ownership, but it can add the specialist discipline that generalist teams often do not have time to maintain.

The most useful partner review is outcome-led. Whether you are on-prem, hybrid, or already in Azure, ask for a practical review of identity and privileged access, security configuration and visibility, backup, restore, and recovery testing, cost governance, and operational runbooks. Who can do what? What can you see and respond to? Can you get back to work quickly? How is spend tracked, controlled, and justified? How are changes made and issues handled?

If a partner cannot explain the impact in plain business language, they are probably too focused on the platform and not focused enough on the result.

For many SMBs, the most useful external help is not a bigger cloud project. It is a clear-eyed review of what is already there. That review should identify the few changes that will materially reduce risk, stabilise recovery, or make spend easier to predict. Once those are done, the environment usually feels a lot more manageable. The team spends less time firefighting, leadership gets a cleaner view of cost and risk, and IT can focus on improvements that are worth the effort.

Bringing it all together

Azure should not be judged by how modern it sounds or how quickly the migration happened. It should be judged by whether it helps the business run better.

For SMBs, four things matter most:

- **Spend is visible, predictable, and controlled**
- **The right people have the right access**
- **Critical systems can be restored quickly**
- **Modernisation is focused on the workloads that matter most**

If you can improve those four areas, Azure becomes more than a hosting platform. It becomes a business asset. If you cannot, the organisation may be paying cloud bills without getting cloud value.

If you want a practical review of where you are today, whether you are on-prem, hybrid, or already in Azure, we can help you work through identity, recovery, security, and cost control in plain English.

We Purpose Technology

+44 (0)1463 572042 | **wepurpose.tech** | info@wepurpose.tech