Synextra

Everything You Need to Know About

Azure Cost Optimisation

THE ULTIMATE GUIDE

While the Azure cloud platform offers incredible flexibility, scalability, and powerful productivity tools, many organisations find themselves spending more than necessary on their cloud resources. In fact, **91% of organisations admit they're wasting money in the cloud.**

At Synextra, we've seen it all too often—Azure bills that make Finance Directors wince, unexpected costs that blow through quarterly budgets, and the frantic search for savings when it's already too late. But the good news is: it doesn't have to be this way.

We know cloud costs can spiral out of control, but with the right approach, tools, and a bit of forward planning, you can keep your Azure spending in check. And that's without sacrificing performance.

This guide brings together a wide range of our thinking on cloud cost optimisation. We'll go deep into Azure cost control strategies, tools, and best practices to help you maximise value from your cloud investment.

Contents

-

1. Common cloud cost challenges.

Before diving into specific optimisation strategies, let's uncover the common pitfalls that drive up cloud costs. Three major challenges consistently appear across organisations:

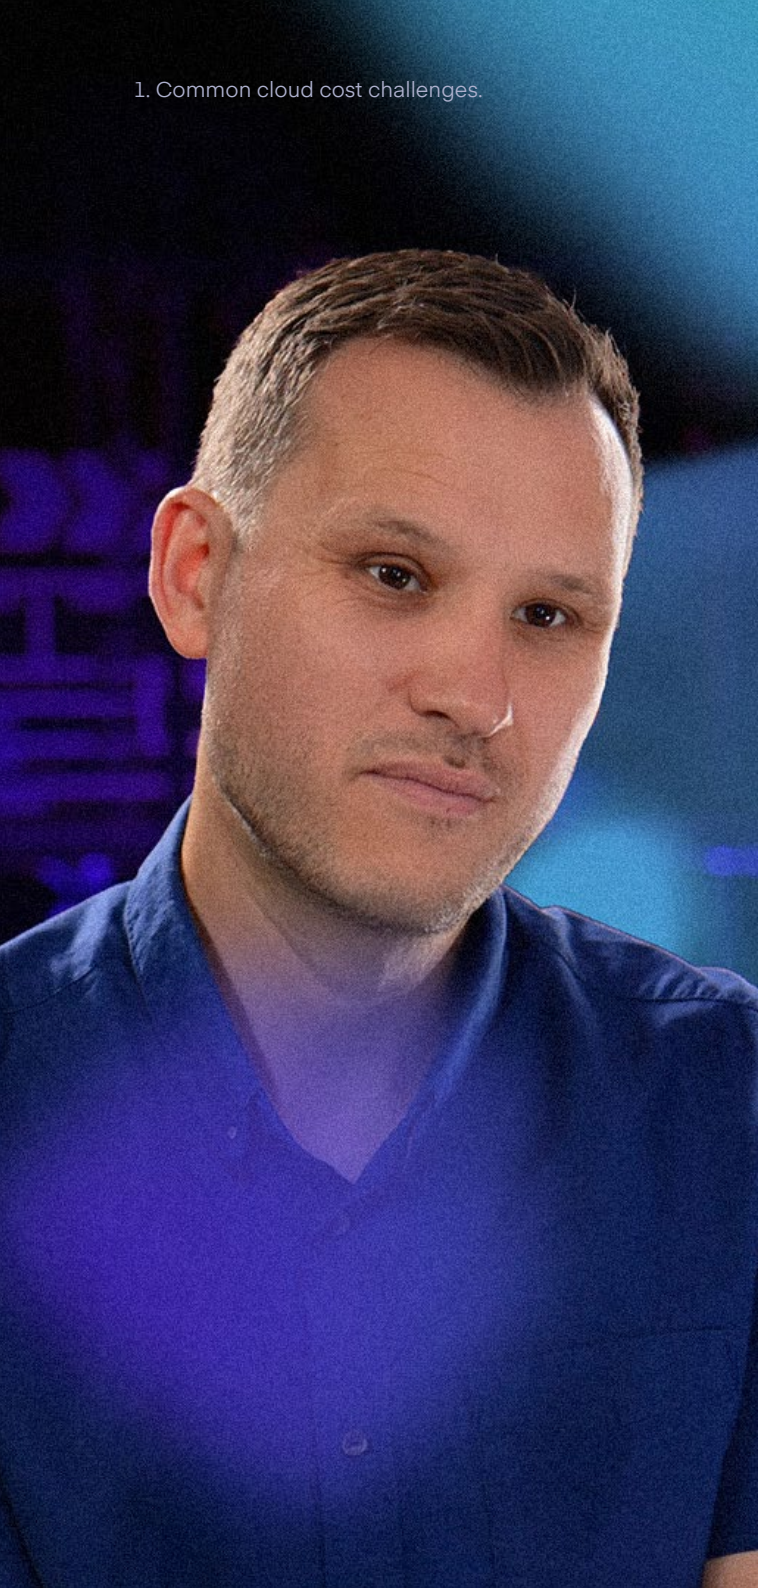

LACK OF CLOUD-SPECIFIC SKILLS

Azure's complexity means organisations without specialised expertise often overlook hidden costs or configure environments inefficiently. As our consultant Matt pointed out in a recent Experts in Polo Shirts podcast on reducing your cloud costs:

> **"Azure can be a beast. If you don't have the skills, it's easy to miss things like bandwidth or backup costs, which can spiral out of control."**

This challenge manifests in multiple ways:

- Missing hidden costs triggered by seemingly simple actions
- Not understanding Azure's cost structure
- Replicating inefficient on-premises practices in the cloud

Solution: Invest in your team's cloud-specific abilities through training, hiring specialists, or partnering with an expert MSP who can guide you through the complexities of Azure's pricing structure.

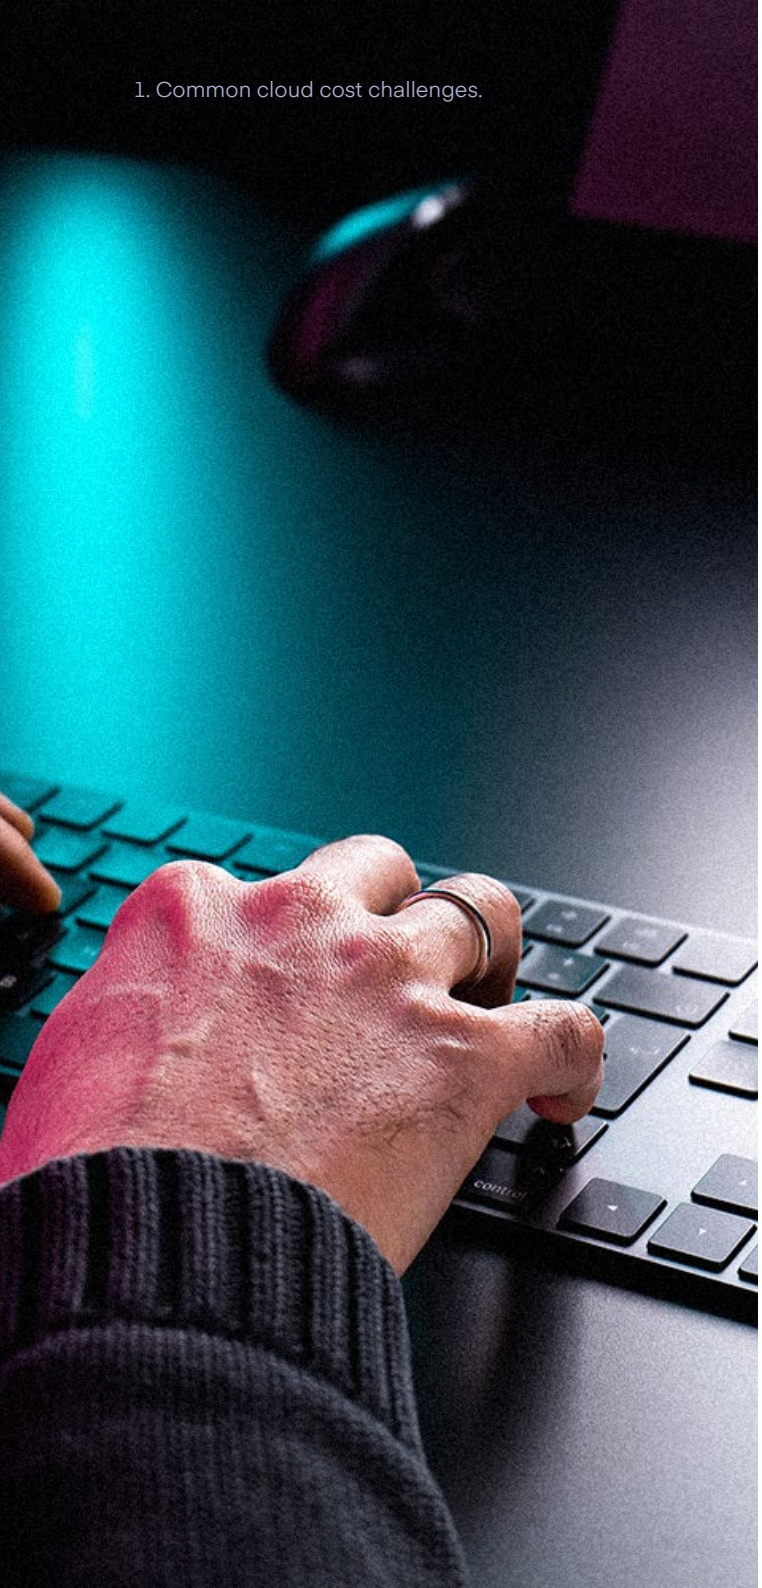

OVERPROVISIONING RESOURCES

Overprovisioning is a common issue, especially during cloud migrations. Many businesses simply copy their on-premises infrastructure to the cloud without optimisation, replicating inefficiencies and driving up costs.

Migration is one of the biggest causes. Someone might simply copy exactly what they had on-prem without thinking about how Azure's resources compare to their 10-year-old servers.

Modern cloud resources are often more powerful than legacy hardware, meaning you could achieve the same performance with smaller, more cost-effective resources. Overprovisioning appears in:

- Unnecessarily large VMs
- Excessive disk sizes
- Inappropriate storage types
- Unused services and features

Solution: Don't blindly copy your existing setup. Assess what's truly needed and use Azure's monitoring tools to identify opportunities for scaling down.

IDLE AND UNDERUSED RESOURCES

Idle resources (including virtual machines, disks, and other services that are no longer actively used but still run) silently drain your budget.

It's easy for things to get away from you—maybe a dev clones a VM for testing but forgets to delete it, or you end up with dozens of disks from old projects just sitting there, adding up month by month.

This challenge becomes even greater in larger organisations where multiple teams access Azure without clear ownership or accountability.

Solution: Implement proper governance with tagging, naming conventions, and automated policies to identify and clean up unused resources.

2. Azure cost optimisation tools.

Microsoft provides several powerful tools to help you understand, manage, and optimise your Azure costs.

Each one serves a different purpose in your cost optimisation journey—let's clarify which ones you should use, and when.

AZURE COST MANAGEMENT

The cornerstone of your cost optimisation strategy, Azure Cost Management gives you real-time visibility into what you're spending across your Azure estate. You can monitor spending patterns, identify anomalies, and drill into where costs are coming from—whether that's specific departments, projects, or resources tagged in a particular way.

It also lets you set up budgets and configure alerts, so you'll know when you're approaching thresholds, or something starts to go off the rails. Want detailed reporting? You've got it. The platform makes it easy to slice and dice consumption data in a way that suits your internal reporting needs.

Limitations to consider	Best practices
Data latency: Some metrics have refresh cycles of 24-72 hours	Implement a comprehensive tagging strategy from the start
Complex pricing integration: New pricing benefits from Reserved Instances or Savings Plans aren't always immediately reflected	Create custom dashboards that highlight the metrics most relevant to your organisation
Resource-specific blind spots: Some services have limited cost visibility	Connect Cost Management with automation workflows to trigger actions based on cost events

AZURE PRICING CALCULATOR

This is your up-front Azure budgeting tool.

Before deploying resources, the Azure Pricing Calculator helps you estimate costs and find the most cost-effective configuration.

How it works

The Azure Pricing Calculator is your go-to tool for planning cloud spend before anything goes live. It lets you model your exact setup—from the type of resources you'll use to where they're hosted—and gives you a detailed cost breakdown for every component.

It's particularly handy when you want to estimate the cost of a specific workload before deployment. You can experiment with different configurations to see what delivers the best value, and get a clearer view of how discounts like Reserved Instances or Savings Plans will impact your bottom line.

Planning to scale? The calculator lets you model that too, so you can understand how costs will behave as demand goes up—or down. It's a bit like test-driving your cloud bill before committing to the real thing.

Cost-saving opportunities

The calculator can help you discover potential savings through:

1. **Reserved instances:** Lock in discounted pricing for one or three years
2. **Spot VMs:** Take advantage of unused Azure capacity at heavily discounted rates
3. **Scaling schedules:** Model scenarios with scaled-down resources during off-peak times
4. **Hybrid benefits:** Factor in savings from existing on-premises licences
5. **Configuration optimisation:** Experiment with different VM sizes, storage tiers, or regions

How to use the calculator

1. Access the Azure Pricing Calculator
2. Select your currency
3. Choose your services from the menu
4. Add usage details for each service
5. Look at commitment options like Reserved Instances
6. Save and export your estimates for further analysis

Common pitfalls to avoid

It's easy to miss the mark with Azure cost estimates if you're not careful. Hidden costs like bandwidth, backup storage, and transaction fees can sneak up on you if they're not included from the start.

Don't assume your workloads will stay flat, either. Usage spikes and scaling can significantly affect costs, so build in that variability.

Be mindful of regions—Azure pricing varies by location, and the wrong choice can skew your numbers. Also, don't overlook cost-saving options like Reserved Instances, Savings Plans, or spot pricing—they can make a big difference.

And finally, plan for growth. If your usage is likely to increase, factor it in now—not when your next invoice surprises you.

For more on how to get the most out of the calculator, check out our guide linked below.

AZURE ADVISOR

Think of Azure Advisor as Microsoft's built-in consultant that analyses your Azure usage and proactively recommends optimisation opportunities.

How it helps with cost optimisation

Azure Advisor will identify:

- Idle and underutilised resources
- VMs that could benefit from right-sizing
- Reserved Instance purchase opportunities
- Redundant resources that could be consolidated

Best practices

- Schedule regular reviews of Advisor recommendations
- Prioritise recommendations based on potential savings
- Automate implementation of low-risk recommendations where possible

3. Payment and pricing models.

Choosing the right payment model is a smart way to optimise your Azure costs. Different workloads benefit from different pricing options. Here's how they work.

PAY-AS-YOU-GO VS RESERVED INSTANCES

The two main pricing models are quite different approaches to cloud resource management.

Pay-As-You-Go

This model does exactly what it says on the tin—you're billed only for the resources you use, charged by the second for most services. There's no upfront commitment, which gives you maximum flexibility.

It's ideal for dev and test environments, short-term projects, or any workload with unpredictable usage. Startups often lean on this model while they're still figuring things out.

The upside? You can scale resources up or down on a whim, try out new services without risk, and avoid being locked into anything long term. The downside? You'll pay the highest per-unit cost, and your monthly bill might look like a rollercoaster ride.

Reserved Instances

Reserved Instances are the opposite. You commit to using specific Azure resources for one or three years, and in return, you get hefty discounts—up to 72% compared to Pay-As-You-Go.

This model is made for stable, always-on workloads. Think production environments, core infrastructure, or any project that hums along 24/7 with predictable usage.

The big benefit is cost savings and predictable pricing, which makes budgeting a lot easier. You can even choose to pay upfront or in monthly instalments, and Microsoft offers some flexibility to exchange reservations if things change. Still, it's a commitment—less suited to workloads that fluctuate or projects with an uncertain future.

Making your decision

There are a few factors to think about when choosing between Pay-As-You-Go and Reserved Instances.

- **Financial impact:** Use Azure's pricing calculator to compare scenarios
- **Future needs certainty:** How confident are you in predicting requirements?
- **Management overhead:** Reserved Instances require more active management
- **Available expertise:** Do you have the skills to optimise your chosen model?

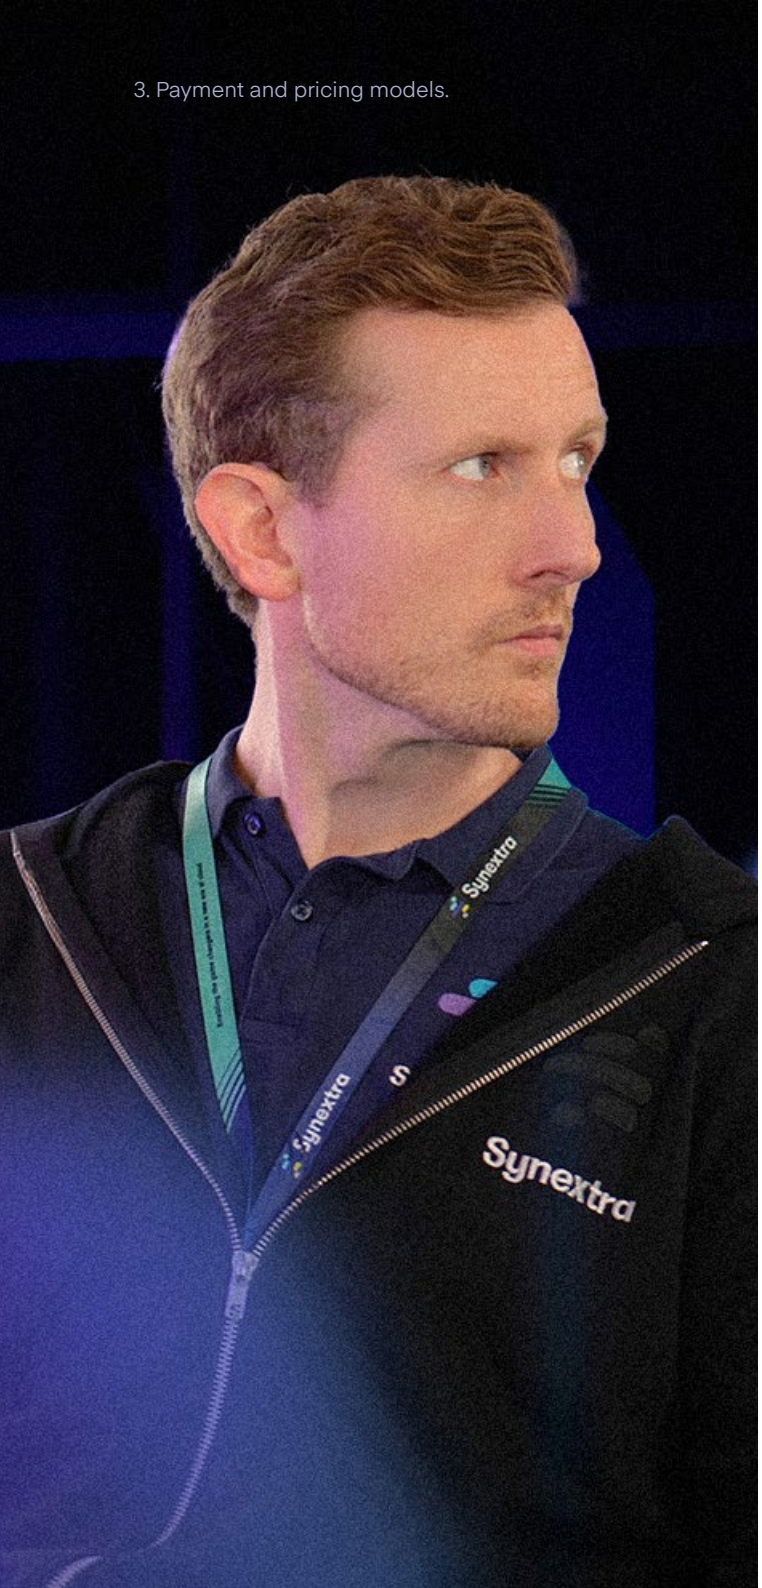

AZURE SAVINGS PLANS VS RESERVED INSTANCES

For organisations that want a middle ground between flexibility and cost savings, Azure Savings Plans are a decent alternative to Reserved Instances.

Azure Savings Plans

Instead of locking into specific resources, you commit to spending a fixed hourly amount for one or three years. That commitment is then automatically applied across a wide range of eligible services—almost like a prepaid deal that flexes with your needs.

There are two types to choose from:

- **Compute Savings Plans** cover VMs, Container Instances, Functions, App Service, and Dedicated Hosts.
- **General Purpose Savings** Plans go broader, including databases, analytics, storage, and integration services.

The real draw here is flexibility. You're not tied to a particular VM size or region, which makes these plans ideal for organisations with mixed or evolving workloads. You still get up to 65% off Pay-As-You-Go pricing—without needing to micromanage every resource.

Reserved Instances

Still, if you've got predictable, stable workloads and want the biggest savings possible, Reserved Instances are hard to beat. By committing to specific resources over a one- or three-year period, you can slash costs by up to 72%.

They also give you more control—perfect if you've got specialised needs or tightly defined capacity plans. But that control comes with a trade-off: more admin overhead and a greater risk of under-utilisation if your plans change.

Choosing between them

- **Choose Savings Plans** if you need flexibility in resource selection and have evolving workload patterns
- **Choose Reserved Instances** if you have stable, predictable workloads and want maximum possible savings
- **Consider using both** for different parts of your infrastructure to maximise both savings and flexibility

These are the basics you'll think about when making the decision, but there is some nuance to consider.

AZURE HYBRID BENEFIT

Azure Hybrid Benefit is a licensing benefit that helps you significantly reduce the costs of running workloads in the cloud. It works by allowing you to use your existing on-premises Windows Server and SQL Server licences with Software Assurance on Azure.

How it works

By reusing your existing licences, you can cut costs by up to 40% on Windows virtual machines and up to 55% on SQL Database services. And if you combine this with Reserved Instances, the savings can climb even higher—potentially up to 80%. Not bad for using what you already own.

Eligibility requirements

- You must have active Software Assurance or subscription licences
- Specific rules apply depending on the licence type and deployment scenario
- Different conversion ratios apply for different licence types

Best practices

- Conduct a licence inventory to identify eligible licences
- Use the Azure pricing calculator to estimate potential savings
- Document your licence usage for compliance purposes
- Consider the total cost of ownership, including licence management

4. Optimisation strategies.

Beyond using the right tools and pricing models, several smart approaches can help you further optimise your Azure costs. They're not quite quick fixes, but they're worth looking into—the savings could really add up.

MANAGING ORPHANED RESOURCES

Orphaned resources—assets left running after the projects they supported have ended—are a common source of unnecessary cloud spend. Here's how to avoid them.

Identification strategies

Finding orphaned resources starts with implementing consistent tagging practices across your Azure environment.

If you tag each resource with its owner, associated project, and expected expiration date, you create a clear trail of accountability. Azure Resource Graph can be particularly helpful here, allowing you to run queries that identify resources without recent activity or those missing critical tags.

Making orphaned resource audits a regular part of your cloud management routine is a must here—schedule these reviews at least quarterly to catch abandoned assets before they drain your budget for too long.

Cleanup approaches

Once you've identified potential orphaned resources, having a systematic cleanup approach saves both time and money.

Consider creating automated processes that identify suspicious resources and alert the relevant teams. Building governance through policies that require proper tagging for all resources helps prevent the problem before it starts. For resources that slip through, Azure Policy can be configured to automatically shut down untagged resources after a suitable warning period, providing a safety net for your environment.

Before deleting any identified orphaned resources, always archive any associated data first—this prevents accidental data loss and provides a recovery path if a supposedly abandoned resource turns out to be important after all.

Back to Contents

IMPLEMENTING BUDGET ALERTS

Budget alerts provide early warnings when spending approaches defined thresholds, allowing you to take action before costs escalate. They're like financial guardrails for your Azure environment, giving you time to respond before a minor overspend becomes a major budget issue.

Setting up effective alerts

To do this, you'll want to go to the Azure Portal and open "Cost Management + Billing". Here, you can:

1. Create budget thresholds at multiple levels (e.g., 70%, 85%, 100%)
2. Configure alerts for both actual and forecasted spending
3. Set up different alert recipients based on severity and scope
4. Define actionable responses for each alert level

Advanced monitoring

- Create custom alert logic based on anomaly detection
- Implement programmatic responses to alerts (e.g., scaling down non-critical resources)
- Integrate alerts with your existing operations management tools

Back to Contents

OPTIMISING STORAGE TIERS

Azure offers multiple storage tiers with different performance characteristics and price points. Matching your data access patterns to the right tier can yield significant savings.

Storage tier overview

Premium tier:
For high-performance needs, highest cost

Hot tier:
For frequently accessed data

Cool tier:
For data accessed less than once a month

Archive tier:
For rarely accessed data, lowest cost but with retrieval fees

Optimisation strategies

To really maximise your storage cost efficiency, you'll want to use lifecycle management policies that automatically move data between tiers based on usage patterns. (You can do this in the Azure Portal, under "Blob service" then "Lifecycle management".)

As data ages, these policies can shift it from Hot to Cool to Archive without manual intervention. Blob index tags help you track metadata and create more granular policies for specific data types.

When setting up new storage accounts, consider how frequently the data will be accessed and choose the appropriate tier from the start. Remember to balance performance needs with budget constraints—sometimes paying a bit more for faster access is worth it for business-critical data.

Back to Contents

AUTO-SCALING RESOURCES

Auto-scaling dynamically adjusts your resource capacity based on actual demand, making sure you only pay for what you need when you need it.

In Azure, you can configure auto-scaling for various services like Virtual Machine Scale Sets, App Service plans, Azure Kubernetes Service, and more through the Azure Portal, Azure CLI, or infrastructure-as-code tools like Terraform or Azure Resource Manager templates.

While powerful for cost optimisation, setting up effective auto-scaling isn't always straightforward and typically requires some technical expertise.

Implementation approaches

- **Schedule-based scaling:** Adjust your capacity based on known usage patterns (e.g., business hours vs. nights/weekends)
- **Metric-based scaling:** Automatically scale based on performance metrics like CPU utilisation or request queue length
- **Predictive scaling:** Use AI/ML to predict future loads and scale proactively

Best practices

- Define appropriate scaling metrics that truly reflect user experience
- Set appropriate minimum and maximum instance counts
- Implement gradual scaling to avoid performance issues during rapid changes
- Regularly review and adjust scaling rules based on actual performance data

If you're new to auto-scaling, we recommend starting with a simple schedule-based approach for predictable workloads before moving to more complex metric-based rules. Microsoft provides detailed documentation for setting up auto-scaling for specific services like Virtual Machine Scale Sets, App Services, and Azure Kubernetes Service.

Back to Contents

HOW PRICING REGIONS WORK

Azure services are priced differently across regions, and strategic region selection can significantly impact your costs.

Regional pricing factors

Pricing varies considerably between Azure regions, with some locations like UK South typically costing more than regions such as East US. These differences are driven by factors including local infrastructure costs, market conditions, and regional demand.

When planning your Azure deployment, you'll also need to consider data transfer costs between regions, which can add up quickly for data-intensive apps. Service availability is another thing to think about, as not all Azure services are available in every region. Finally, data residency and compliance requirements may dictate where certain workloads must be hosted, regardless of cost implications.

Optimisation strategies

To optimise costs across regions, you might want to think about deploying non-latency-sensitive workloads in lower-cost regions while keeping performance-critical services closer to your users.

Where possible, consolidate resources in fewer regions to minimise expensive inter-region data transfers. For customer-facing services, calculate the trade-off between cost savings and performance—proximity to users often justifies a slightly higher regional cost by delivering a better user experience.

The key is finding that balance between regional pricing advantages and the potential latency impacts that could affect your application performance or user satisfaction.

Back to Contents

5. Service-specific optimisation.

Different Azure services have unique cost drivers and optimisation opportunities.

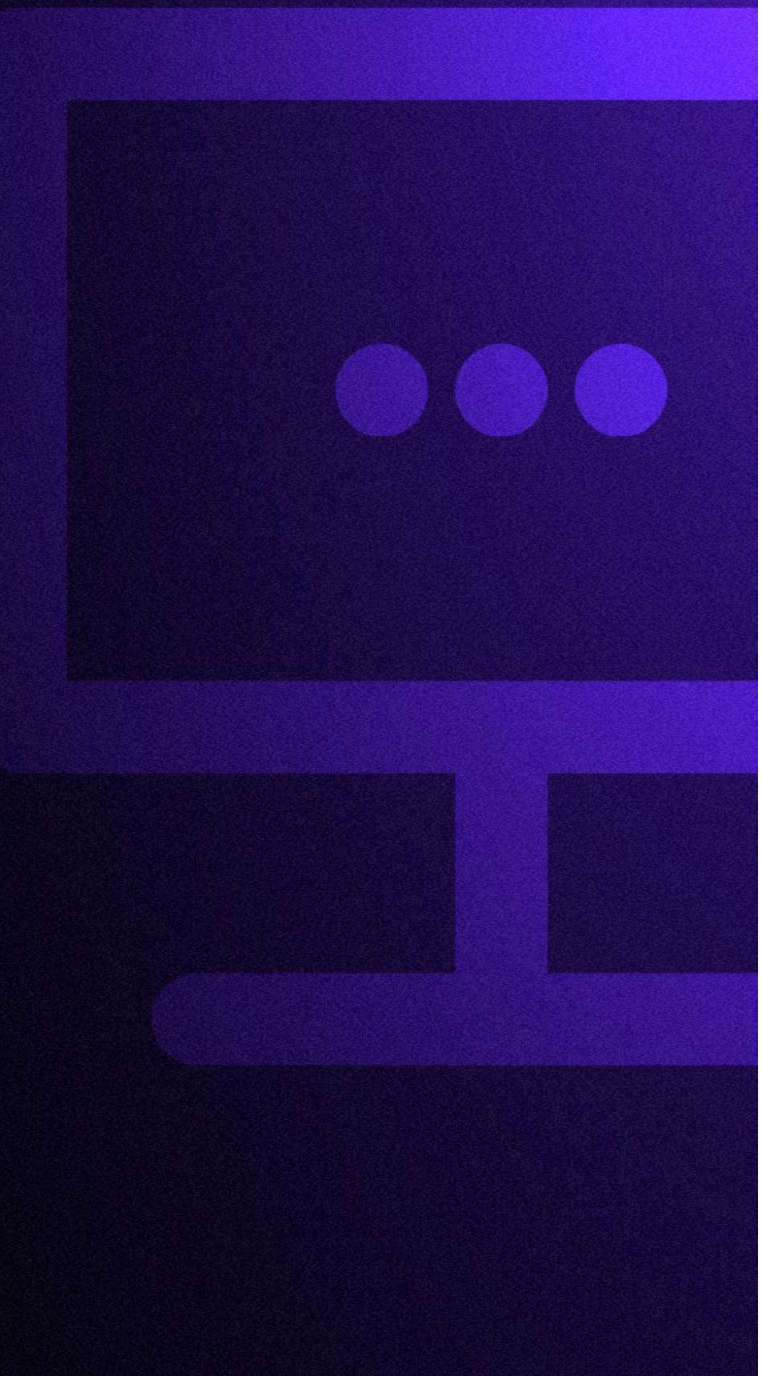

VIRTUAL MACHINE OPTIMISATION

Virtual machines often represent the largest portion of Azure spending for many organisations. We recommend taking a look at our Azure Masterclass on choosing the right size Azure VM—something that'll definitely impact your VM spending.

Right-sizing strategies

Start by using Azure Monitor to pinpoint underutilised VMs—you might be surprised how many resources are sitting idle. For workloads that don't need constant performance, B-series VMs offer a cost-effective, burstable option. If your workloads are more variable, VM Scale Sets let you scale up and down automatically. And for jobs that can handle interruptions (like batch processing), Spot instances are a great way to save even more.

Commitment options

If your VM usage is steady and predictable, Reserved Instances can cut costs significantly. For more flexibility, Savings Plans offer discounts across multiple VM types. And when you combine either of these with the Azure Hybrid Benefit, you're stacking savings on top of savings—up to 80% in some cases.

Operational optimisation

Set up start/stop schedules for non-production VMs, especially dev and test environments. Use auto-shutdown to avoid racking up costs overnight. You can also lower costs by switching to more affordable storage options for VM disks and making sure your OS disk sizes aren't oversized for what they're actually doing.

Back to Contents

AZURE VIRTUAL DESKTOP COST CONTROL

Azure Virtual Desktop (AVD) offers powerful remote work capabilities, but costs can spiral without proper management.

The five biggest AVD cost drivers tend to be:

1. **Oversized compute resources:** Session host VMs represent the largest expense in AVD deployments, and are often provisioned larger than necessary
2. **Idle session hosts:** Without proper scaling plans, session hosts run 24/7 even when idle
3. **User profile storage:** FSLogix profiles stored on Azure Files can quickly accumulate
4. **AVD image bloat:** Outdated images accumulate in the Azure Compute Gallery
5. **Management overhead:** Ongoing administration requires specialised skills

Optimisation strategies

To bring these costs down, there are a few different things you can try.

For compute costs:

Start by analysing what different user groups actually need in terms of performance. Resource-heavy users can be placed into dedicated session pools to prevent unnecessary overprovisioning. Also, make sure the number of session hosts adjusts with staffing changes—there's no point running full capacity when half the team's on holiday.

For scaling plans:

Enable automatic scaling to align capacity with real-world usage. It helps avoid idle resources racking up costs. For session hosts that run 24/7, Reserved Instances can provide a cost-effective alternative. Just be sure to fine-tune your scale-in and scale-out thresholds—defaults often miss the mark.

For profile storage:

FSLogix profiles can quietly balloon in size, so exclude any files users don't really need. Only provision the storage that's necessary, and regularly shrink profiles to keep them lean. And clear out profiles belonging to former employees.

For image management:

Set up a routine to clean out old VM images. Keeping the last three months' worth is usually enough, and anything older can go after each update. A tidy image library not only saves money, but makes management easier too.

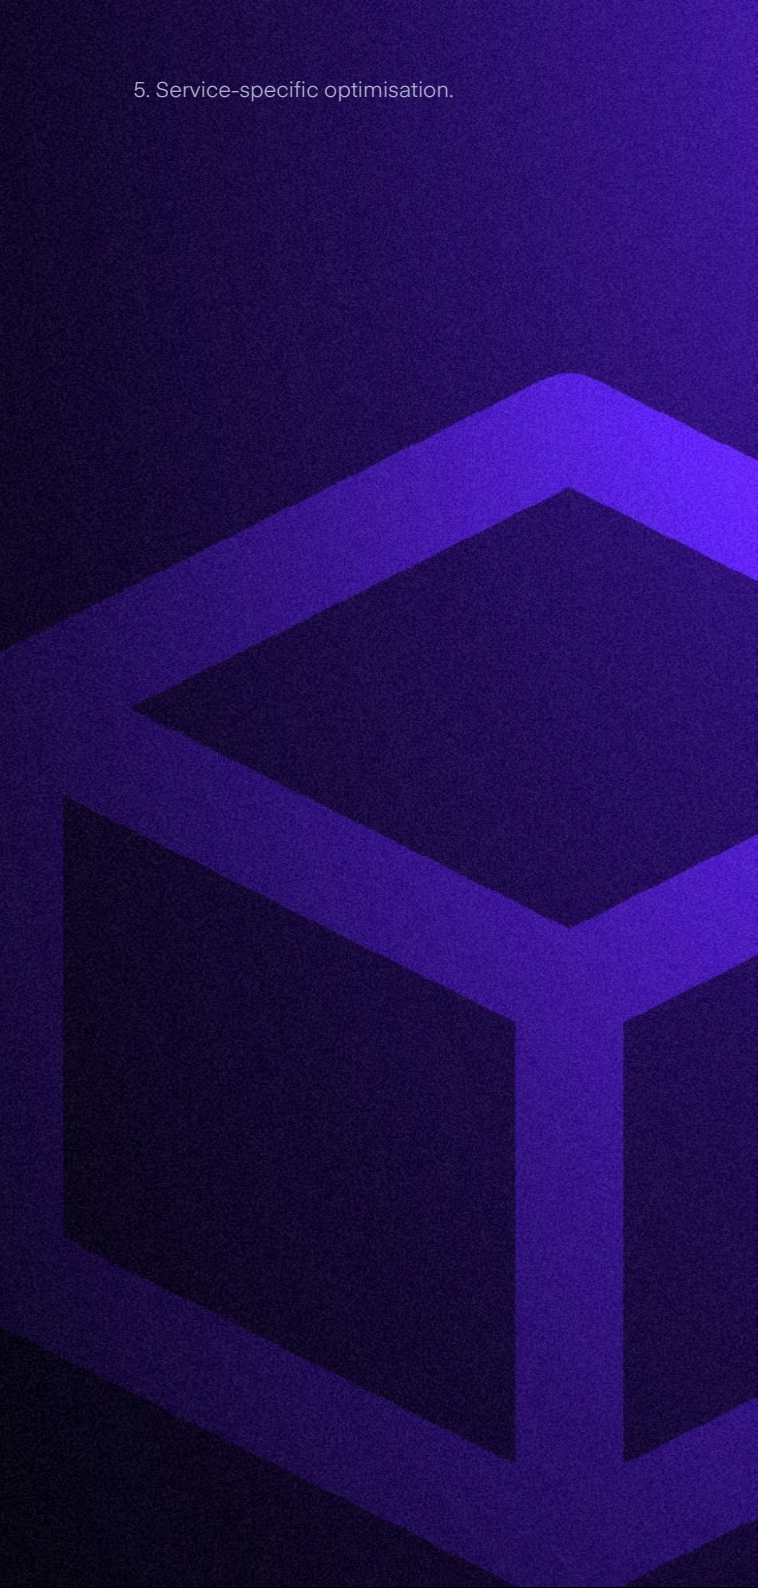

DATA AND ANALYTICS COST MANAGEMENT

Data and analytics services can consume substantial resources, with costs growing unpredictably if not properly managed.

Many firms find their analytics bills increasing faster than any other part of their Azure environment due to exponential data growth and complex processing requirements. So consider the following strategies, which focus on the three main cost drivers for analytics workloads: storage, processing efficiency, and compute resources.

Storage tier optimisation

Start by setting up automated lifecycle management policies that shift infrequently accessed data to cooler, cheaper storage tiers—this alone can save a chunk. Combine that with solid data retention rules to automatically archive or delete anything obsolete, and you're on your way to leaner, cleaner storage. Just make sure you're always balancing performance needs against cost—some data might be cheap to store, but expensive to retrieve when you need it fast.

Query and processing efficiency

Data platforms like Synapse, Databricks, or Data Factory are powerful—but they're also capable of racking up massive bills if you're not optimising your queries. Keep things efficient by partitioning large tables sensibly, applying indexing strategies that match your query patterns, and using materialised views when you're hitting the same queries over and over. The less your systems have to work, the less you have to pay.

Right-sizing analytics clusters

Wherever possible, enable auto-scaling to match the actual workload—especially for Spark. If your clusters don't need to be running 24/7, set them to shut down when they're not in use. For more consistent SQL workloads, make sure your pool sizes reflect your real-world needs, not overestimates. And if your usage is more sporadic? Consider going serverless to avoid paying for idle compute altogether.

Find out more in our deep dive into Azure Data Services.

Back to Contents

6. Governance and best practices.

Effective Azure cost management needs proper governance practices—if you want consistent cost control across your organisation, it's a must. Without some solid rules of the road, even the best cost tools won't save you from budget overruns. Here's how to get your governance house in order.

IMPLEMENTING A TAGGING STRATEGY

Many organisations struggle with cloud costs simply because they can't easily see who's spending what and why. A good resource tagging strategy solves this problem by creating transparency and accountability across your Azure environment. Implementing the right tags from the start, you'll not only understand your current spending better but also spot patterns that highlight potential savings.

Here are the key components of an effective tagging approach:

Essential tags

- **Owner/Department:** Who's responsible for the resource
- **Project/Application:** What business function the resource supports
- **Environment:** Production, development, testing, etc.
- **Expiration/Review Date:** When the resource should be reviewed for continued need

Implementation

- Create a tagging policy document with standardised naming conventions
- Use Azure Policy to enforce mandatory tags
- Implement automated validation of tag values
- Regularly audit and correct tag compliance

SETTING UP COST CONTROLS

Proactive cost controls help prevent budget overruns before they occur.

Think of them as guardrails for your Azure environment—they don't restrict legitimate business needs, but they do stop runaway spending before it becomes a problem. Without these controls, it's easy for teams to unintentionally exceed budgets, especially in larger orgs where multiple people have deployment privileges.

Implementing controls like these creates a safety net to protect your cloud budget while still enabling innovation:

Resource quotas

- Set subscription and resource group quotas
- Implement spending limits for development/test subscriptions
- Configure service-specific quotas (e.g., VM sizes, storage capacity)

Deployment controls

- Use Azure Blueprints to standardise compliant environments
- Implement Infrastructure as Code with built-in cost guardrails
- Require cost impact analysis for new deployments

Back to Contents

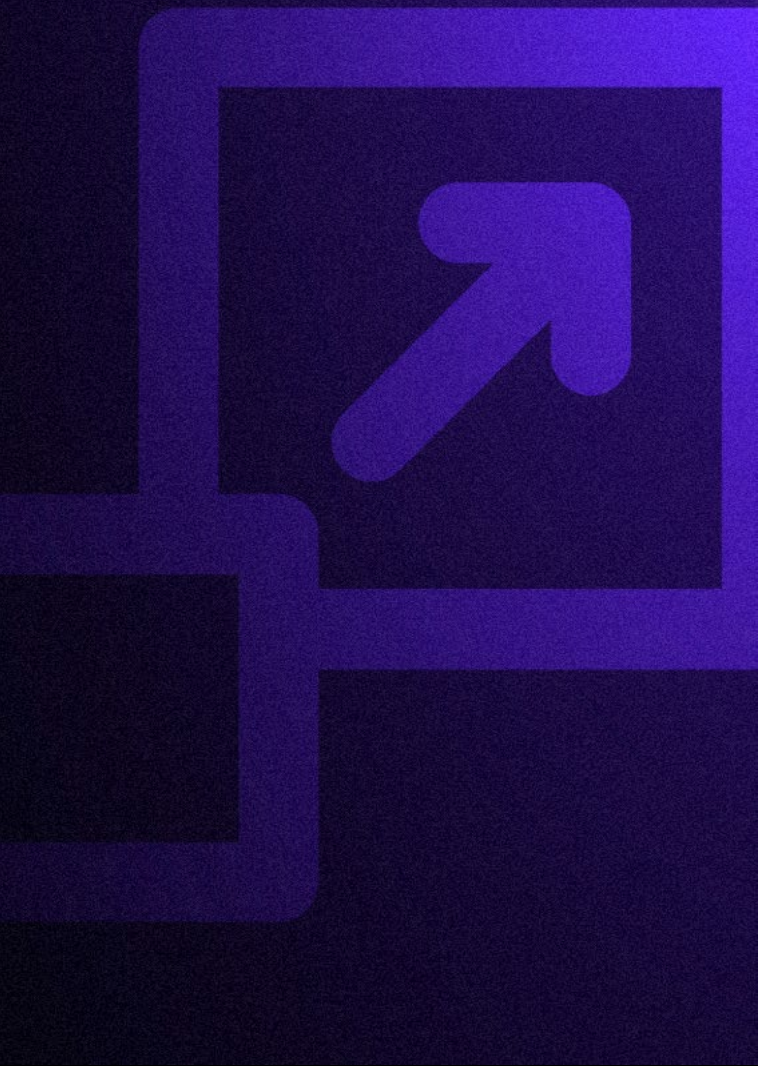

REGULAR REVIEW PROCESS

Establishing a cadence for cost reviews helps you with ongoing optimisation.

Cloud environments are dynamic, with resources and usage patterns constantly changing. A one-time optimisation effort simply won't cut it—you need regular check-ins to stay on top of your spending. Without scheduled reviews, cost optimisation becomes reactive rather than proactive, often triggered only after significant overspending has occurred.

A structured review process helps you identify trends, catch issues early, and continually refine your approach:

A potential review schedule might look like this:

- **Weekly:** Quick check for anomalies and unexpected spending
- **Monthly:** Detailed service-level analysis and trend identification
- **Quarterly:** Strategic review and major optimisation initiatives

You might consider focusing on areas like:

- Spending trends and variances from your budget
- Utilisation patterns and right-sizing opportunities
- Reserved Instance and Savings Plan coverage
- Orphaned and underutilised resources

Back to Contents

7. Getting started with Azure cost optimisation.

Cloud cost optimisation is about spending smarter so you can achieve more. From our experience helping businesses across the UK, we've seen organisations achieve remarkable savings, often while improving their cloud performance.

The beauty of Azure cost optimisation is that the gains compound over time. Small changes like implementing proper tagging might save you some pennies, while right-sizing your VMs adds more substantial savings. Then adopting Reserved Instances for your steady workloads could significantly reduce your monthly bill. Before you know it, you've dramatically reduced your cloud spending and created a more efficient environment.

Best of all, many of these optimisations can be implemented quickly. While full optimisation is an ongoing journey, those first few changes often deliver immediate and substantial returns.

Whether you're just starting your Azure journey or looking to rein in an established cloud environment, the potential for savings is significant. Every pound saved on unnecessary cloud resources is a pound that can be reinvested into your core business or new, game-changing digital projects.

Ready to optimise your Azure costs?

Ask us about our no win, no fee Azure cost optimisation service.
If we don't save you money, you don't pay us a thing. Simple.

SPEAK TO AN EXPERT

SYNEXTRA.CO.UK

0161 883 1383