WHITEPAPER

Cloudscaler

The true cost of cloud transformation

TIPS ON HOW TO AVOID THE PITFALLS AND GO FASTER

www.cloudscaler.com

Introduction

Cloud infrastructure is undoubtedly the future. It reduces organisational cost and complexity, unlocks agility, flexibility and scale as well as providing improved security. As a result, most organisations have started using cloud in some capacity – migrating their existing applications or switching to cloud native solutions.

However, many large organisations are becoming frustrated at the slow speed of business transformation underpinned by cloud, as well as their inability to realise the promised benefits of cloud including cost savings. A few early adopters are even choosing to move back to on-premise.

In this whitepaper, we shine a light on the hidden costs of cloud transformation and give transformation leaders top tips on how to avoid those costs, accelerate their cloud transformation and realise the value of cloud.

Cloud transformation budget is the tip of the iceberg

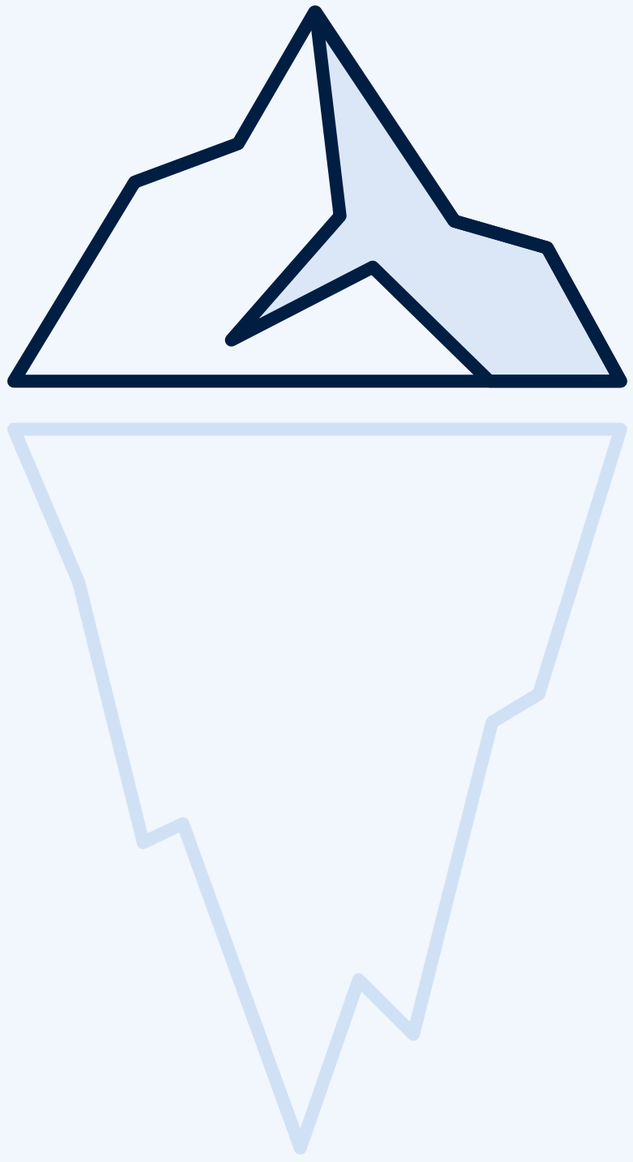

The reality for many enterprises on a journey of cloud transformation is a limited number of applications in live service and a lack of genuine value delivered. It's simply taking too long to move to the cloud, timelines are uncertain and budgets seem to extend year-on-year.

As a result, companies are failing to realise the top-line benefits of cloud, whether that's improved customer experience, improved reliability, increased cost efficiency or business growth. Instead, they face spiralling transformation and cloud spend.

The bad news is that the true cost of slow cloud transformation is far higher than most realise.

Cloud transformation budget is just the (visible) tip of the iceberg. Delays carry hidden costs for the organisation which are far, far greater. These costs are extensive and, what's more worrying is that most organisations don't realise the scale of the problem.

You can’t manage what you can’t measure

Every organisation has overarching business objectives associated with cloud transformation such as improvements in customer service, system stability, to name a few. Achievement of these business objectives is pretty straight forward to track.

It’s also relatively simple for organisations to compare the cost of running applications in the cloud to the cost of running them on-premise.

The question is, how do you measure cloud transformation progress?

We believe there are two critical measures to track the success of cloud transformation

For a successful cloud transformation, everyone from the C-suite through to the SREs must focus on these metrics.

They are an indicator to inefficiency and poor decision making that will impact the bottom line, while simultaneously damaging morale across engineering and business teams who urgently want to realise the benefits associated with cloud.

Read on to understand what these two measures are, and how you can manage them.

1 “Time To Cloud Value” (TTCV)

The elapsed time from the date you start building a cloud platform (sometimes referred to as a landing zone[1]) to the point at which your first business service or application is live and being used by customers. It’s the time it takes to build your minimal viable product (MVP) cloud platform and then successfully migrate your first workload. It’s the point where you begin to realise a return on investment (ROI).

TTCV represents your cloud transformation burn rate. The whole time you’re building a cloud platform and preparing to migrate applications, you’re monopolising key resources and paying for on-prem infrastructure. You’re not realising the savings outlined in your business case so it’s vital TTCV is minimised.

Minimum viable product (MVP) cloud platform

Landing zone[1]: MVP cloud platform	Security Controls · Data Controls · Operational Controls · Financial Controls
Infrastructure	IaaS · Hybrid Cloud · PaaS
Data centre & edge provider	aws

[1] There is a full explanation of what a cloud landing zone is on page 21.

Cloud platform evolution

The initial MVP cloud platform has minimal common services available.

Early cloud users must therefore implement their own services and controls to different degrees of completeness, often duplicating work.

With every platform evolution, more common services become available but cloud-based applications will need to be reworked to remain compliant (impacting delivery pace).

2 “Amount of Cloud Friction” (ACF)

Putting the first cloud workload live is a major achievement but it’s only the first step, next you’ll need a production line for migrating workloads onto cloud quickly and effectively. In parallel you’ll also need to support and manage the live cloud platform.

At this stage there will only be a baseline set of common cloud services, controls and tools. To achieve the full benefit of cloud, the cloud team must iteratively release additional controls, tools and services for users.

As the number of common services increases, friction reduces, the cloud becomes easier for business teams to manage, delivery pace increases and efficiencies in cloud spend can be realised.

Cloud platform evolution reduces the Amount of Cloud Friction (ACF)

As the cloud platform matures, friction reduces which exponentially improves agility, delivery speed, cost efficiency and security.

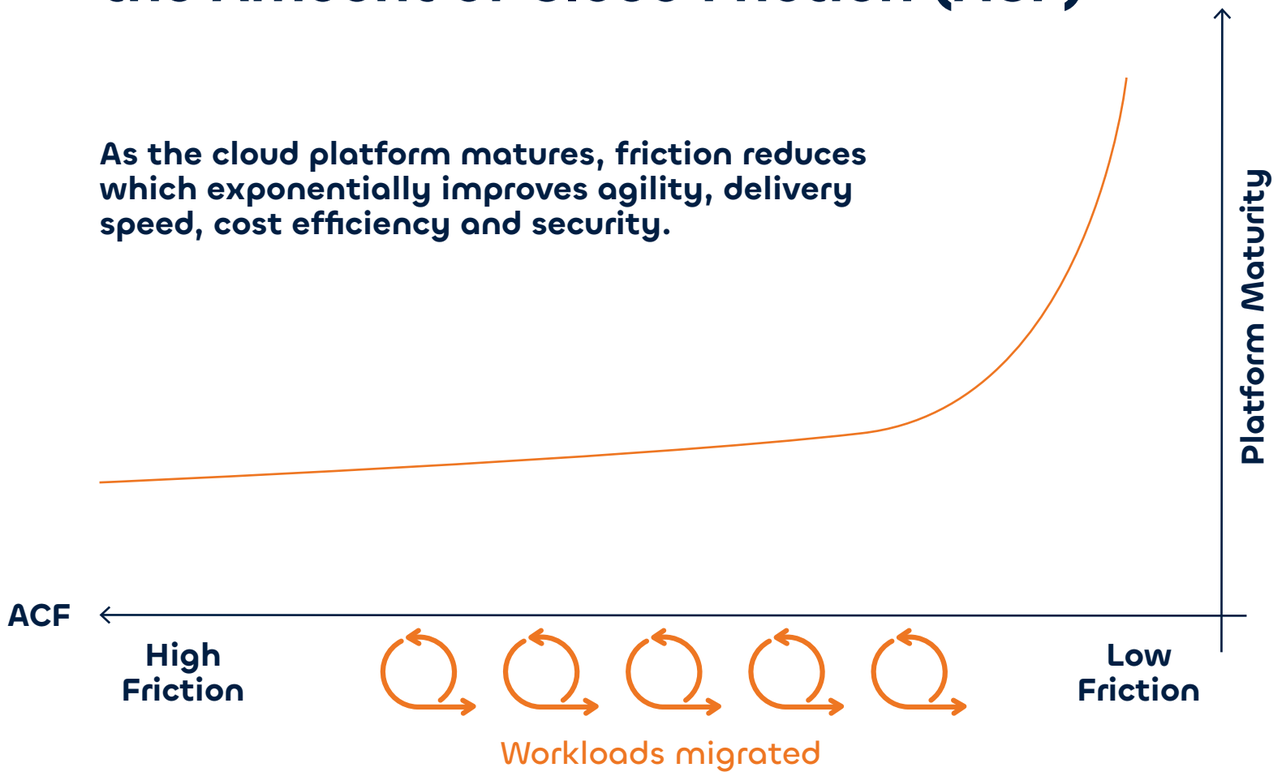

Indicators of High ACF:
- Business teams must implement their own services and controls
- Business teams need highly-skilled cloud resources
- Duplication of core functions across teams (backups, logging)
- Tooling sprawl – each team chooses different solutions

Indicators of Low ACF:
- Centralised services are easy to consume (backups, logging, monitoring) and enable business teams to go faster
- ‘Out of the box’ implementations of controls which reflect and standardise the organisations risk policy (compliance as code).
- Business teams no longer need deep cloud skills, agile methods are fully enabled

Drivers leading to high cloud transformation costs

We know that organisations are challenged by escalating transformation and cloud spend. We also know that delays incur extensive hidden costs.

Let’s consider what is influencing the TTCV and ACF. What is causing delays and ultimately leading to escalating costs?

Competing cloud platforms

How enterprises end up with competing cloud platforms

It's surprisingly common for organisations to have multiple landing zones built on top of the same hyperscaler, instead of following best practice, which is to have one single central platform. Typically, this is because:

// **The central cloud platform rollout is too slow.**
Urgent business priorities, critical strategic projects and frustrated delivery teams, force individual teams to build their own cloud environments while waiting for the centralised platform.

// **No mandated cloud strategy or cloud governance framework.**
If there is nothing to prevent teams from building their own platform, why wouldn't they seek to take control? It may seem obvious, but few large organisations have a unified, clearly-communicated cloud strategy and associated governance framework.

"The landing zone is where your workloads land, or where your projects land, and it enables you to hook in a lot of automation to just drive consistency and efficiency."

Martin Bishop, Chief Technologist, UK Public Sector, AWS

Landing Zones explained...

In this **video** of Martin Bishop, Chief Technologist, UK Public Sector, AWS

The cost impact of competing cloud platforms

Cloud platform teams within the same organisation rarely collaborate effectively unless they're encouraged to do so through rigidly enforced governance - which they rarely are. If they're not working together, teams will be trying to solve similar problems in parallel, re-inventing the wheel and likely implementing different solutions – this has significant cost implications:

// Decreased productivity

With cloud platform teams duplicating resources and effort, overall productivity is impacted. Merging those teams means fewer people are needed to work on a specific problem or new feature so you can either divert resource to a different task to accelerate delivery or, even better, you can reduce cloud resources without impacting productivity.

// Technology sprawl

Parallel teams are likely to solve the same problem by implementing different solutions. For example, they might select different monitoring tools, which means:

- // Different tooling licences and fewer opportunities to leverage bulk discounts.
- // Additional procurement time to negotiate, buy and process renewals for those licences.
- // Skilled resources required to install, configure, maintain, and support each tool.
- // Knowledge share is more difficult, experts in specific tools become a critical point of failure and too valuable to lose.

// Demand for resources and skills

In organisations with multiple cloud platforms, we've observed inter-departmental competition for cloud resources. Teams are all advertising similar roles, or worse, those same teams start poaching resources from one another because an internal transfer is faster and less risky than hiring externally. Resource churn and internal competition become a distraction. Each cloud platform team has open vacancies and you're spending more on recruitment, induction, IT equipment, ongoing HR and management.

// Increased security risk

Multiple platforms create multiple access points which means a wider attack surface to manage. The CISO will have a larger, more complex estate to manage which will inevitably drive up costs. Your organisation will be exposed to increased risk and will need to work harder and invest more resources to mitigate those risks.

Lack of centralised cloud services and tooling

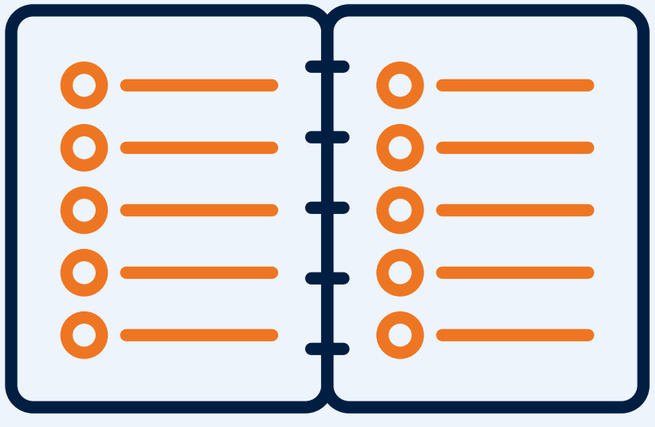

The list of pre-configured services, tools and processes available to tenants is often referred to as the **Service Catalogue**.

Hyperscalers provide a range of ‘ready to use’ cloud services and tooling designed to help organisations build, deploy, and manage applications in the cloud including Compute, Storage, Database, Networking and Automation tools.

It’s not uncommon to see organisations establish cloud platform teams who simply provide access to the hyperscaler account within which users can do whatever they want. Inevitably some users implement robust controls and use approved corporate tooling while others are less thorough. As a result the security posture of the organisation is weakened and other risks are introduced.

In order to mitigate this problem, the central cloud team should:

- **Pre-configure and “wrap” hyperscaler services** to reflect your security posture and embed cost efficient practices.
- **Create a catalogue of additional services, tools and processes** for example account vending, network vending and DevOps pipelines.

If your cloud platform only provides a few common services, Amount of Cloud Friction (ACF) is high, the cloud is complex and time consuming to use. Tenants must build a large proportion of their own services and duplication across teams is common. As the number of services provided by the platform increases, friction reduces, migration speed and delivery pace increase and efficiencies in cloud spend can be realised.

This will accelerate cloud adoption and better protect the organisation from risk.

Balancing demand and platform maturity

Internal demand for cloud tends to be high. Choosing which applications should be prioritised for cloud and when those teams can be given access to the cloud platform has significant cost implications, particularly as full landing zone maturity can take up to two years.

Teams know they will go faster, be more productive and have fewer dependencies on other teams if they can leverage cloud infrastructure. When they're under constant pressure to reduce budgets while increasing output, using cloud is an obvious way to do more with less.

It's tempting to approach the problem of demand and prioritisation by adopting one of two extreme approaches:

1. **Giving tenants unrestricted access to cloud before a basic landing zone is available.**
2. **Keeping cloud access closed until the landing zone has reached full maturity.**

Those teams given early access will gain some benefit however, they will need highly-skilled cloud resources to plug the gaps associated with using an immature landing zone. The organisation must be aware, and accept, that only the most significant risks associated with running production workloads on the cloud will have been mitigated in these early stages. Early adopters will also need to refactor their applications every time the landing zone introduces new controls and features. This can be disruptive and will undoubtedly impact productivity.

Waiting for the landing zone to reach full maturity before any workloads are migrated carries an opportunity cost. Similar to waterfall projects, it's likely to mean the cloud platform is delivered late because creation of the landing zone is too big and complex to manage effectively. Delivery teams in the queue will become increasingly frustrated and staff attrition can often increase.

Five steps to reduce costs while accelerating TTCV

01

Start with a well-defined cloud operating model and governance framework

Definitions

A cloud operating model is made up of the processes, tools and organisational structures used to manage the cloud environment. It includes the policies, procedures and best practices that optimise cloud usage, ensure security and compliance, and manage costs effectively.

A cloud governance framework is the set of policies, procedures and controls that ensure the risks associated with cloud adoption are managed and that the use of cloud is aligned with business objectives.

Implementing an effective **cloud operating model** and **governance framework** will reduce costs and accelerate TTCV by providing:

// **Clarity on strategy, priorities and direction:** Aligning your teams, focusing on the most important objectives and setting expectations with key stakeholders.

// **Effective risk management:** Cloud adoption introduces significant risks, with the right governance you can mitigate and monitor those risks.

// **A blueprint for new governance:** Most large organisations are constrained by waterfall governance which has become stale and bureaucratic. Moving to cloud is an opportunity to implement new, highly-automated governance which in turn can unlock the true power of cloud and enable more agile ways of working.

// **Improved collaboration:** Common tools, processes and best practices will underpin collaboration across departments and geographies. Teams will be encouraged to share knowledge and assets as well as re-use architectural patterns which will avoid technology sprawl.

Your cloud operating model and cloud governance framework provide the foundation for monitoring, reporting and managing TTCV and ACF.

02

Don’t build your own landing zone

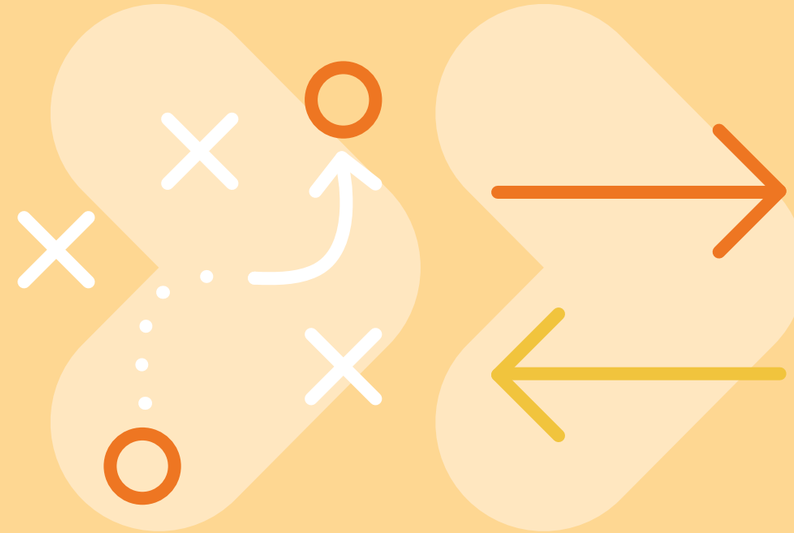

If you choose to build your own landing zone you will be choosing to embark on a journey of exploration filled with trial and error – it is likely to take several years and incur significant costs.

Even though cloud best practice continues to evolve, there are a number of mature landing zone solutions available on the market. Very few organisations are truly unique enough to justify building their own so taking advantage of these ‘off the shelf’ solutions is a far more effective route to cloud transformation success.

Landing zones embed common controls and guardrails to mitigate risk exposure – essentially, these are the “rules of cloud use” for your organisation and all tenants will need to comply. This in-turn enables centralised compliance reporting and monitoring which is useful when dealing with internal auditors or external bodies like the UK financial regulator.

Deploying a proven landing zone product will mitigate many of the risks associated with running production workloads in the cloud. If your organisation needs additional controls, they can be easily added. It’s a fast-start solution which could save you months or even years of effort.

Whether you decide to build your own landing zone, use one provided by a hyperscaler or deploy a landing zone product from a specialist provider, you will need to ensure that your cloud operating model defines who is accountable for prioritising new landing zone features as well as building out, maintaining and supporting the landing zone.

03

Simplify cloud adoption for delivery teams

Organisations must determine what aspects of the end-to-end cloud service the central cloud team will be accountable for versus those elements the users will need to provision themselves. This is commonly referred to as the ‘shared responsibility model’.

Some tenants will argue they need total flexibility with a minimum number of services pre-configured or centrally controlled. This can be a mistake and lead to several problems that will increase your TTCV.

- **Inconsistent use of services:** Hyperscalers release new services and update existing features every day which can lead to inconsistent use across your organisation and also cause costs to rise unexpectedly.
- **Technology sprawl:** The more flexibility and choice you provide to tenants, the greater the risk of each tenant choosing a different service option or configuration. By adding popular software products to a Service Catalogue, tenants will be forced to seek permission to use non-standard alternatives.
- **Wasted effort and duplication:** Without access to pre-configured tools and services, teams will waste time building their own versions, often not realising that other teams have already created something similar.
- **Inconsistent or insufficient risk mitigation:** If tenants are expected to apply their own risk mitigation controls and guardrails, there will be inconsistent (and possibly insufficient) risk mitigation across your cloud estate.

By creating an opinionated (wrapped) version of services (Service Catalogue) available to tenants via an enhanced landing zone, you force users to adhere to best practice which in-turn ensures consistency.

04

Adopt a phased migration approach

As with agile software development, the true measure of successful cloud migration is working software running reliably in production. The best way to achieve that is by taking a phased approach - the deployment of an MVP landing zone for low-risk pilot workloads followed by rapid uplifts of the landing zone enabling the migration of workloads with increasingly higher risk profiles.

As the cloud platform continues to evolve, minor landing zone features can become blockers on the critical path for a specific migration. You will need the right governance and a collaborative planning process to resolve such dependencies.

The priority order of application migrations is influenced by number of factors:

// **Risk:**
You cannot migrate higher risk and mission critical workloads to the cloud until your landing zone is mature enough to manage them effectively.

// **Competitive pressure:**
All industries face increased competition, you may need to prioritise specific applications in order to stay competitive.

// **Potential cost savings and contractual pressure:**
The organisation will benefit from moving expensive on-prem workloads quickly. Data centre contract renewals have significant cost implications and might force organisations to transition out of the data centre when subjected to contractual deadlines and sales pressure.

// **Scalability and resilience:**
Bottlenecks and failings in the current on-prem solution may necessitate a move to cloud but that transition is likely to require re-architecture and re-engineering which will need to be factored into your roadmap.

05

Leadership and people are key

While technology and processes play vital roles in driving transformation, it is people who play a critical role in achieving sustainable and successful outcomes. When it comes to the people aspect of transformation, key considerations include:

// **Change management and communication**
Wherever you are on your cloud journey, it is likely you'll need to change what you're doing and how. Effectively managing and communicating change will be critical. Explaining what you're doing and why is key. Listen to and address any concerns, provide clear communication channels, and create a culture of transparency and trust. This will ensure new processes are adopted and governance is observed.

// **Skill development or acquisition**
You may need new skills and capabilities to define and implement a new cloud operating model and governance framework. Identify skill gaps. You may need to design or acquire training, recruit or get third-party support to address those gaps. Upskilling will not only improve output but will foster a learning culture that will see teams embrace and drive change.

// **Leadership and organisational alignment**
Executive endorsement will be needed to secure resources and support. Leaders must clearly explain any changes. They must also lead by example, show commitment and actively participate in the process. You may need to recruit individuals with cloud transformation experience to avoid common pitfalls. This will also enable you to leverage their network to build a team with the right skills and experience to keep you on track. Your team needs to be aligned on the transformation goals. This will enable collaboration, reduce resistance, and create the supportive environment necessary for success.

Conclusion

Starting a cloud transformation without a strategy or well-conceived plan is likely to result in considerable overspend and a constantly extending timeline – achieving your goals will become a complex journey.

The key to minimising TTCV is to build an MVP landing zone that can be opened to low-risk application teams as quickly as possible. This should be followed by rapid iteration or development of the landing zone - measuring and continuously improving the ACF.

Key mistakes to avoid are:

// **Lack of funding** for the cloud platform team. Moving to cloud at scale isn't cheap, pace is largely driven by the capacity and capability of the cloud platform team.

// **Poor cloud leadership** will create a lack of clarity over the target end-state and steps needed to reach it.

// **An inability to modernise approval governance** will mean the benefits of cloud are constrained by bureaucratic processes that stifle agile delivery.

// **A poorly defined cloud operating model** leads to in-fighting and conflicting priorities rather than a consolidated, aligned effort that prioritises the urgent removal of roadblocks.

// **A lack of accountability** in the cloud platform team. Let's never forget the cloud platform team operate a service for delivery teams and they should always listen to feedback from their customer and improve their service accordingly. Cloud transformation is a team effort and will only be as successful as it's component parts.

Free Cloud Acceleration Workshop

Identify key opportunities to improve your cloud operations with a Cloudscaler Cloud Acceleration Workshop.

This two-hour session with an experienced Senior AWS Architect will explain how you can accelerate your cloud programme, whilst also reducing risk and spend.

To book your workshop, simply complete the online request form at:

www.cloudscaler.com/workshop

Appendix

The foundation: cloud platform or landing zone

What is a cloud platform?

A cloud platform (sometimes referred to as a landing zone) is the infrastructure layer between a hyperscaler and the business applications you're building and operating. It provides common cloud standards, tooling and services to mitigate the risks of running cloud-based production services. It simplifies cloud use across your organisation – it's the most important component for realising organisation-wide cloud-related cost savings.

The landing zone layer