Heavily invested in Cloud but not yet seeing the full value?

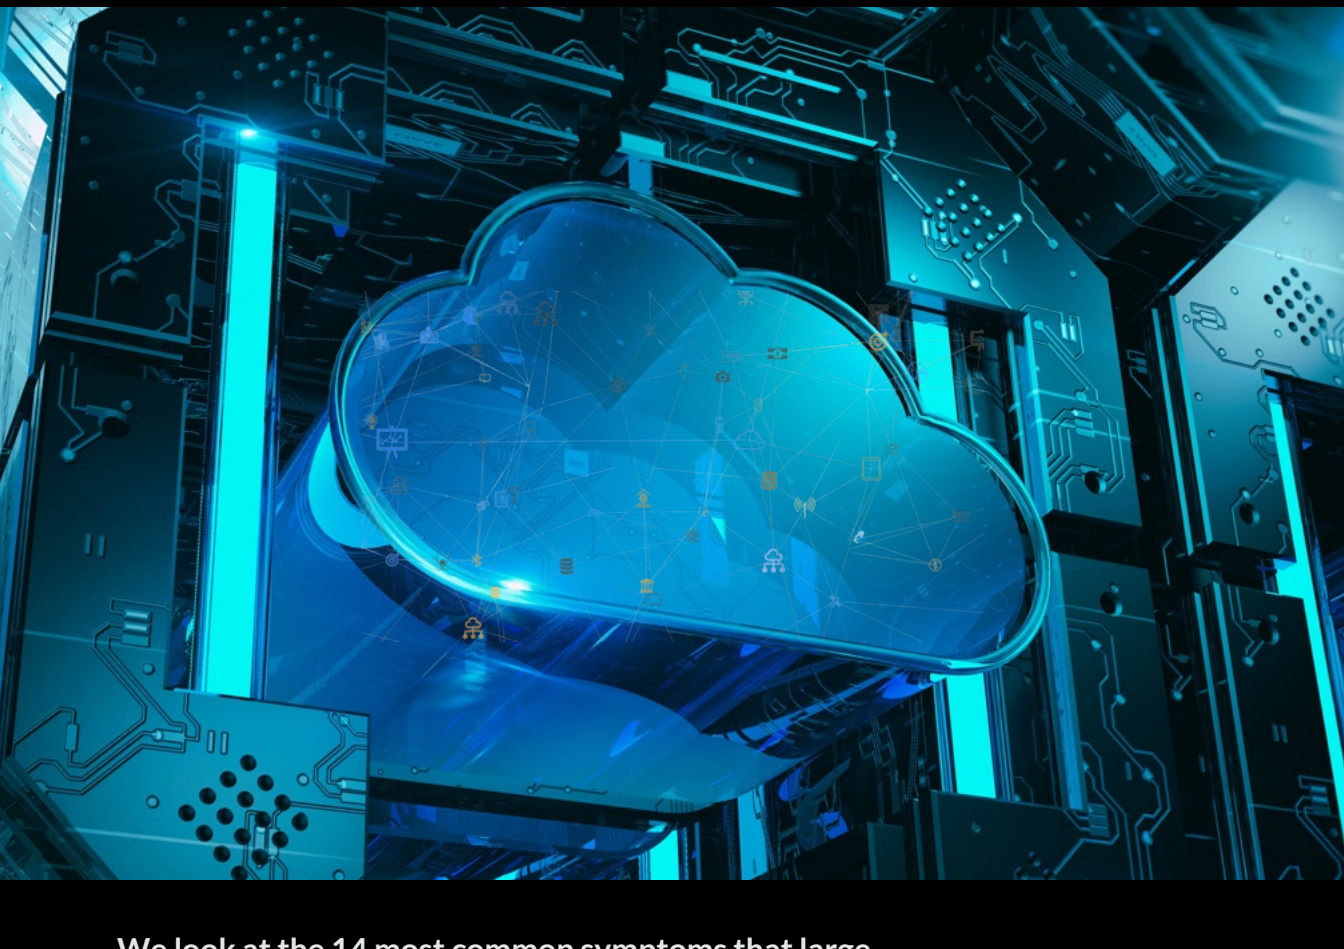

We look at the 14 most common symptoms that large organisations face when migrating to cloud, and explain how a well designed and implemented Cloud Operating Model can reset your cloud journey on a path to success.

DigiRen

Organisations are constantly looking for ways to increase revenue and reduce costs. For IT teams the demand is to become “more agile” to provide service features faster, to concentrate on valuable services and to do so with a limited budget.

For technologists, cloud computing promises to deliver all this through self-service, on-demand and pay per use features and functionality. Indeed by providing immediate access to resources that can be scaled up and down automatically, cloud computing can solve many of the legacy IT problems that have long bedevilled IT projects, notably long lead times to install hardware, poor service delivery due to infrastructure lock-in, and high upfront costs for compute resources and licences.

This argument has been made and won, and most organisations have committed significant resources to migrating their workloads to cloud. Yet having made the investment, few are realising the promised benefits and value. Why is that? Why, having built a cloud “landing zone”, is it being used only for proof of concepts? Why, having trained staff and created a “cloud guild”, are so few building cloud native applications? Why, having employed “agile coaches”, are services still expensive and taking too long to deliver?

At DigiRen, we have years of experience working in and with cloud vendors on some of the largest migration programs, and what we have learned is that when it comes to cloud, it is not about the technology, it is about people and how they use technology. For organisations, this means adopting new ways of working that enables their people to take advantage of cloud computing - a Cloud Operating Model.

A well implemented Cloud Operating Model provides businesses with the capabilities and structure to realise the value of their cloud investment. It combines many proven models in the industry today - LEAN Product Development, Human-centred Design, Agile Delivery, DevOps & CI/CD - with the core cloud computing value proposition of self-service, on-demand and pay per-use.

If your company has made significant investments in cloud but you are not seeing the value, the symptoms identified in this book will be familiar to you. They are the problems commonly seen in organisations with cloud programs that are stalled, falling short of delivering the expected business outcomes. Read on to understand how a well designed and implemented Cloud Operating Model can resolve these problems and enable you to adopt cloud at scale, with increased velocity and success.

DigiRen

Table of Contents

Symptom #1:
You've built a cloud landing zone...
but nobody is using it!

Are you wondering why no one in your organisation is using your newly built cloud landing zone or taking you up on your offer to be more 'agile'? If this is you, then chances are you are not realising the full value of your investment in cloud computing. In this chapter, we will explain why this is a problem and what you can do about it.

There are a number of potential reasons why people aren't engaging with your landing zone, but we will concentrate on those that you can directly influence; that is, how your service works and how you expect your potential customers to use the service.

Cloud computing services are different to traditional IT services. In a traditional IT project, an application team will request infrastructure resources, such as servers, databases and storage from an infrastructure team. In response, the infrastructure team will provide a bill of materials (BoM). The BoM has a price tag attached, which will be included in the business case, so the full cost of the infrastructure is accounted for at the start of the project. During delivery, the infrastructure team will purchase new equipment, if required, build or connect the resources as specified and then hand it over to the application team once complete.

By contrast, in a cloud computing delivery model, the resources are delivered on-demand and provided on a self-serve basis. There is no BoM, no hand-over, and no infrastructure team to engage. Further, the cloud resources are priced per use, so there is no upfront cost.

One key mistake many organisations make when building a cloud landing zone is to hide the self-serve, on-demand, pay-per-use nature of cloud computing behind layers of traditional IT process. Ask yourself…

If you have hidden the value of cloud computing behind these traditional IT tasks, then you can be pretty sure that your potential customers are looking elsewhere to deliver their applications. The cloud native developers are going directly to the public cloud providers—AWS, Azure, GCP — while the traditional developers will stay with their on-premise infrastructure that they at least understand.

So, what can you do? First, acknowledge the problem! Don't hide your landing zone behind layers of IT process, a cloud team, ITSM tool, or access restrictions. You cannot offer your cloud landing zone like any old traditional IT service and expect your organisation to use it.

You need to start by embracing the value proposition of cloud computing — self-serve, on-demand, and pay-per-use. This is what makes cloud different to traditional IT services. It will attract early adopters within your organisation who would otherwise go directly to the public cloud providers, and over time, as your organisation matures, you will be able to attract the on-premise application teams as well.

Offering a self-serve, on-demand, pay-per-use cloud landing zone is not easy. It will impact how your organisation builds business cases, the skillset of your workforce as well as their roles and responsibilities, and it will impact how you deliver on your commitments. You will have disputes with your security team, network team, financial team, and your own management.

However, if successful, you will have built a full cloud operating model — the processes, skills and tools required to take advantage of cloud computing. More importantly, your organisation will be much more agile, with an IT group that can respond quickly to changing business circumstances.

In the next chapter, we will look at how injecting new talent and capabilities into your team. will help you re-imagine your cloud landing zone by exposing the cloud computing value proposition to the wider organisation.

Symptom #2:
Your cloud service is hidden behind layers of IT processes

So, your team of cloud experts have built a beautifully automated CI/CD pipeline that can deploy well-architected infrastructure to the cloud in minutes… but why is no one is using it?

In this chapter we will explore how the make-up of your cloud team may be hindering your attempts to expand usage of your landing zone.

No doubt your new cloud landing zone – or new iPaaS middleware or SaaS based Identity service – was built by your IT organisation. Your IT organisation has a well-defined project delivery model for introducing new services into the business, including operational integration points and a standard operating model based on existing ITSM principles.

To complete the landing zone, or iPaaS or Identity Service, the IT organisation will have employed a team of cloud practitioners; expert architects, developers, and operation staff to build and operate the solution.

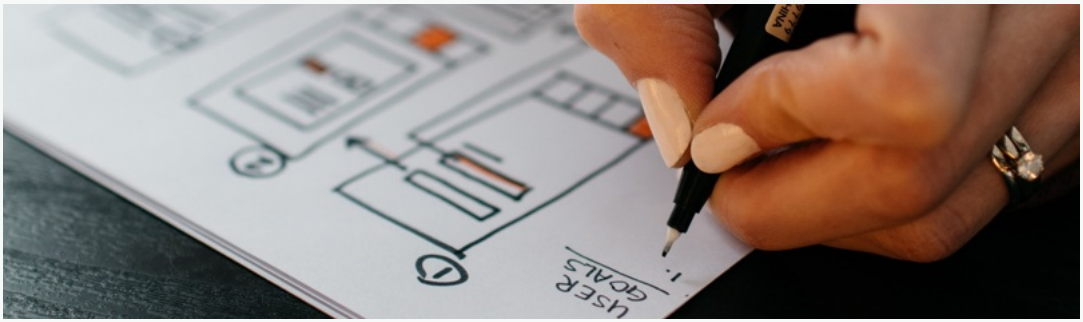

Using cloud services to build and operate solutions is what these experts know and love. They can talk for hours explaining how they architected a particular capability using a cutting-edge cloud feature, or how its implementation was guided by the need for redundancy, or for speed, or some other factor. But there is a drawback. For the most part, these experts won't be able to explain:

1) the problem their solution attempts to solve
2) why they built a particular capability
3) how to use their solution (at least not in all cases).

This is not their job. They built the solution according to the requirements. And because these requirements probably reflect how IT projects have always been delivered, chances are your new cloud service is hiding behind layers of IT process, such as an ITSM tool, SecOps-enforced feature restrictions and, possibly, the cloud team itself.

Consequently, no one is using your service.

Can this situation be salvaged? Yes, it can, but it will require new talent and a different delivery model.

What you need is talent with an incentive to challenge existing IT processes and delivery models in order to expose the self-serve, on-demand and pay-per-use value propositions of cloud computing to your potential customers.

They should be able to explain how to use the service and why different features are valuable, and also take feedback on problems with the service. Expanding just a little, you should be looking for product management, technical product management and financial management talent, responsible for, respectively, engaging customers, designing features, and ensuring feature development is financially viable.

As they are joining the cloud team, this product management talent should have a clear preference for delivering self-serve, on-demand, pay-per-use products and features.

With this new product management talent embedded, you will have a team who will want to replace your IT project delivery model with a product delivery model, incorporating industry standards such as agile, DevOps and human centred design.

You will also have a team ready to take on the next set of challenges:

- training traditional IT teams to deliver cloud-native applications,
- re-designing business cases to take advantage of the pay-per-use nature of cloud, and
- reducing stakeholder overhead by simplifying responsibility matrices.

We will explore each of these challenges and how to approach them in the following chapters.

The essential take-out here is that introducing product management talent will help your Cloud team re-imagine your Cloud Landing Zone and make it more appealing to your customers. This is the first step to expanding cloud uptake and usage.

Symptom #3:
Everyone has their cloud certs, … but no-one is building cloud apps!

Have you spent a great deal of time and effort training your staff to be cloud practitioners only to find no one is building cloud native applications? In this chapter, we look at why this is happening and what you can do about it.

Cloud computing requires a whole new skillset, so training is a logical first step. Ensuring your staff have the skills, mindset and confidence to take advantage of the self-serve, on-demand, pay-per-use nature of cloud computing takes time, and requires leadership and positive re-enforcement.

Introducing the cloud guild

A cloud guild is one way to improve cloud knowledge, generate excitement and momentum, and produce cloud leaders who will evangelise the power of cloud computing throughout the organisation. They do this by incorporating activities such as hackathons, game days and workshops, as well as more traditional learning experiences, such as training courses that lead to certifications.

Some companies have gone all in on cloud guilds, asking their entire staff to become cloud practitioners and obtain certifications. In fact, some companies proudly promote the number of certified employees as part of their pitch during hiring. Certifications, particularly cloud certs, are valuable in the broader industry and prospective employees look favourably on organisations that provide the training and resources to achieve them.

Despite the initial cloud guild success, for many companies the momentum has waned, the energy dissipated, and the cloud skills learned in the early days are starting to atrophy. Why? What has happened?

A fading cloud guild is a symptom of a larger problem: the lack of a cloud operating model that connects the value of cloud computing—self-serve, on-demand, pay-per-use—with the skills, processes and tools of your organisation.

We previously explained why, for many companies, so few teams are using their organisation's cloud landing zone - that being they are not really designed with customers in mind because they were delivered using a traditional IT project delivery model. To solve this problem, we recommended the cloud team replace their IT project delivery model with a product development model.

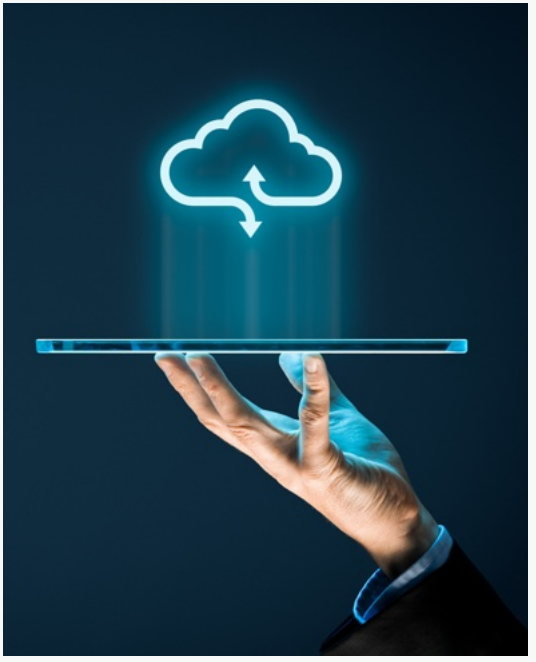

With a product development model in place, you can introduce the IT product team. Their role is to build cloud-native applications and services to solve business problems.

IT product teams are skilled to take advantage of the self-serve, on-demand, pay-per-use nature of cloud computing. Just like the cloud team, IT product teams are focussed on delivering features that customers need.

With your cloud team and IT product teams using the product development model to deliver features, your organisation will have the two fundamental teams in place to implement a fully functioning cloud operating model.

These teams will have the incentive to make further changes that we will discuss in later chapters, including simplifying your organisation's RACI, introducing a cloud business case, and expanding the self-serve, on-demand cloud model to other services, such as networking and security. They'll also form the core of your cloud guild.

Essentially, a waning cloud guild is a symptom of a larger problem in your organisation: the lack of a functioning cloud operating model. By introducing a product development model and IT product teams you will have the fundamentals in place to implement a cloud operating model that will enable the adoption of cloud at scale across your organisation.

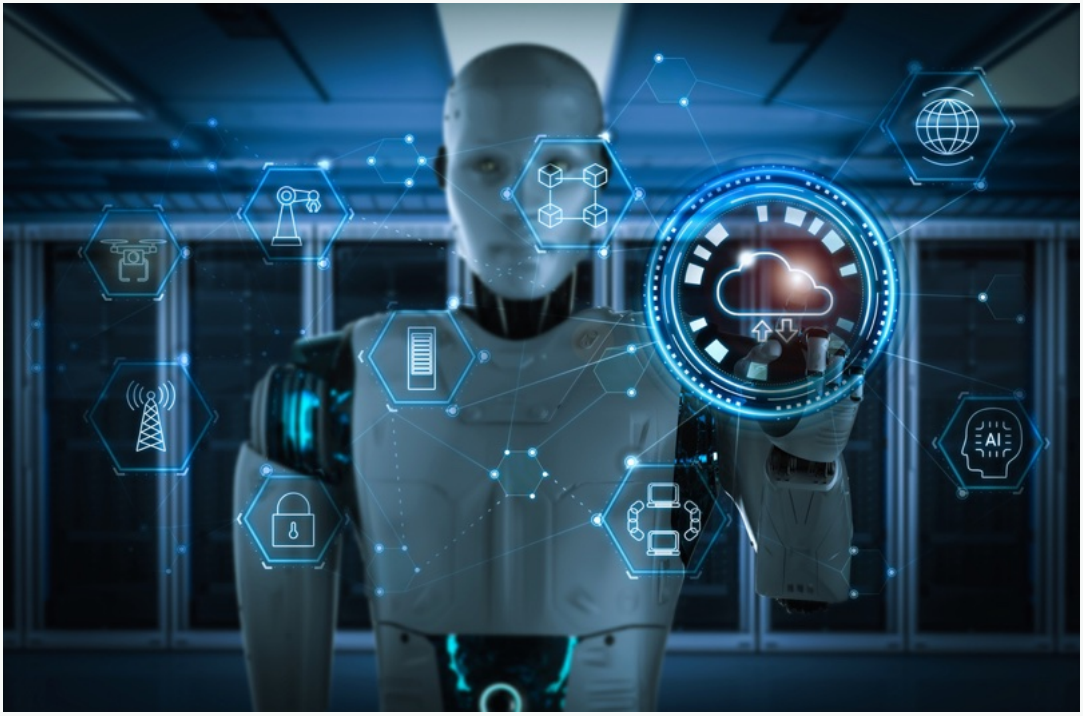

Symptom #4:
Your RACI matrix is overcomplicating operations and causing friction

If you're in the process of making the move to the cloud but are still dealing with the friction that comes with complex responsibility matrixes in your IT department, then it's time to introduce a new model of working. In this chapter, we will discuss why your existing RACI is holding you back and what you can do about it.

The problem with RACIs

A RACI matrix sets out who is Responsible, Accountable, Consulted, or Informed for completing tasks or deliverables in a project. They're a key deliverable that all IT teams must produce before putting their service into production.

Modern IT departments are flush with complex RACIs, usually reflecting the different technology teams and stakeholders involved in delivering and operating IT services.

One key problem with complex RACIs are the demarcation disputes, resulting in delivery delays, cost overruns and in some cases, cancellation of initiatives altogether. The problem can be so bad that some delivery teams have looked to bypass stakeholders altogether, going directly to public cloud providers, and contributing to the problem of shadow IT with its own cost, operational and security risks.

Introducing the cloud shared responsibility model

So where does your cloud investment fit into this discussion? Cloud services come with their own type of RACI—the cloud shared responsibility model. It defines the responsibilities between just two actors: the cloud service provider and the cloud service user.

You'd be familiar with the model if you have used a public cloud provider like AWS, Microsoft Azure or GCP. These providers are responsible for providing resources, such as compute, storage and networking using the core cloud characteristics: self-serve, on-demand, pay-per-use.

Emulating this relationship between cloud providers and cloud users in your organisation can address the demarcation disputes and reduce the appeal of shadow IT by exposing what makes your cloud investment so valuable: the ability for users to access resources quickly with no high upfront costs.

Introducing a fully functioning cloud shared responsibility model is a long-term undertaking, but it is possible to start small with just two teams: your cloud team, representing the cloud provider, and an IT product team, representing the cloud users.

These two teams must forge a cloud shared responsibility agreement where your cloud team is responsible for providing a self-serve, on-demand, pay-per-use service, and the IT product team is responsible for using the service to deliver value to the business.

Putting the new model to work

With the introduction of a cloud shared responsibility model, and your cloud team now accountable to more IT product teams, you'll find the demarcation disputes causing cost overruns, delays and cancellations will be greatly reduced. It will also reduce the appeal of shadow IT and mitigate the accompanying cost, operational and security risks.

This first agreement is only the start of the discussion as it will impact other parts of the business. Two functions in particular—security and networking—will need to be folded into the cloud shared responsibility model before IT product teams can realise the full value of the model. This will be the subject for a future chapter.

Separately, finance will also be impacted as the business moves away from funding large capital outlays required to purchase equipment and licenses to operational costs to fund feature development. Again, this will be the subject of a future chapter.

Introducing the cloud shared responsibility model is a significant undertaking and has quite far-reaching implications for team members and their roles. For the introduction to be successful, it needs to be accompanied by a change program that places team members and their needs at the centre of the discussion. We will cover how to do this a little later in the book.

Symptom #5:
Your business is not willing to fund your cloud service

Despite offering a pay-per-use cloud service to your business, are you finding it difficult to compete with your existing on-premise solutions?

Nothing comes for 'free'

In many organisations, the IT department is responsible for a large budget. Each year, the department is allocated budget to cover headcount, software, services and a certain amount of infrastructure. The scope of the budget is hard fought but, once settled, the department is then funded (in theory) to service the needs of the business. This is BAU (business as usual).

Every now and then, the business will request a service not offered by IT. To add this service, a project must be spun up. IT will estimate the cost of the service, it will be included in a business case and, if the business case is successful, then IT will go about standing up the service as part of an IT delivery project. Once the project is complete, the service is ordinarily included in BAU and cost for upkeep is added to subsequent budgets.

Unfortunately, this funding model can create perverse incentives, causing problems for the business such as reduced agility and increased costs. In future chapters we will explore two serious issues in depth (1) encouraging empire building and (2) fostering monopoly behaviour around a "strategic" asset or service. For the purposes of this chapter, we will deal with a third problem that arises from the funding model: The impression that IT services are' free'.

Of course, IT services are not free; their true cost is hidden deep within the overall IT budget.

What does this mean for your cloud service? First and foremost, it turns one of the key value propositions of your cloud service, pay-per-use, into a liability. Assuming your cloud team plan to pass on the cost, why would the rest of the business move to cloud when they can use the 'free' services offered by IT?

How to solve this problem? First, acknowledge its source. In traditional IT departments, IT costs are not accurately allocated to the business. Many organisations have acknowledged this problem and have attempted to solve it via a charge back mechanism. In fact, many traditional IT solutions include chargeback as a feature of the service.

Introducing a chargeback mechanism to correctly allocate IT costs is a good start—even without cloud services—as it will make it possible to accurately estimate the real value of traditional IT solutions

However, this solution shines a light on a second problem; the fact that for any new IT service, the first IT project to use the service will pay for most of it, whilst subsequent projects will only pay a small proportion to cover the delta. As a result, every project team has an incentive to wait until another project team pays for the IT service.

This often leads to horse-trading between project managers as each attempt to shift costs to others, wasting everyone's time. To solve this second problem, your business will need to re-imagine how to deliver IT services to apportion the costs fairly across projects.

Fortunately, this is where the pay-per-use nature of cloud-based services becomes important. By providing a pay-per-use cost model for your cloud service, you will remove the incentive for project managers to horse-trade and provide a fair way to allocate costs.

Of course, delivering your cloud service will require funding and ordinarily it would be provided via an IT project – so we are back to the horse-trading problem. However, in a previous chapter, we explained why the IT project delivery model is a problem for cloud services and recommended your cloud team move to a product development model instead.

It's time to change your funding model

The product development model is based on the cloud team delivering features to meet the needs of your users. A cloud business case is a funding model designed to complement the product development model, by funding new features on an incremental basis. Compared to the traditional business case that is designed to fund a time-boxed project, there is no end date for a cloud business case. Instead, funding is on-going, like BAU activities.

Tip: You will need to work with your finance team to build a financial model to support your cloud business case.

Of course, there is nothing magical about a cloud business case: You must still demonstrate that the benefits outweigh the costs. So, in addition to charging your customers, you will want to implement new methods for tracking progress such as feature utilisation, feedback from users and even bug fixes.

Introducing a cloud business case and charge back will enable your organisation to compare costs between cloud services and traditional IT solutions.

Symptom #6:
Wasting time on building IT business cases

A well-executed cloud strategy should have eliminated large IT business cases. If your business is still building them, then you're probably not realising full value from your cloud investment. Let's look at the likely causes and what you can do to address them.

The Company Toll

Building large IT business cases for funding new technology services will drain agility for most organisations. It takes too long, the costs are eye-watering, and the results rarely match the promised benefits. It is, in a very real sense, a Company Toll.

But what choice do you have?

In previous chapters, we mapped out what an alternative might look like. You can start by replacing your IT project delivery model with a product development model; hire talent with experience creating cloud products; introduce a cloud shared-responsibility model and fund incremental feature development with cloud business cases. These are key components of a well-functioning Cloud Operating Model that enable you to optimise your cloud investment, providing the agility your stakeholders expect of you.

If you have started to implement your Cloud Operating Model and you're still building large IT business cases, then chances are the problem lies with other parts of your IT department still working under the old operating model. Even if your Cloud team is no longer contributing to the Company Toll, you cannot rid yourself of it until all IT teams have made the transition.

Operating cloud responsibly and intelligently

At this point, let's take a moment to discuss how large IT business cases create a Company Toll and how moving to a Cloud Operating Model eliminates it.

The issue starts with stakeholder engagement. During the development of an IT business case, you will need to get quotes from impacted teams. For many teams, this is a perfect opportunity to correct past mistakes. For example, a team may have underquoted a previous engagement and now want to square their ledger.

Other teams may take the opportunity to purchase a fancy IT gizmo or platform they "need". Some teams will simply provide go-away quotes because they just don't want to do it.

This issue is compounded by the fact you'll need to build in contingency accounting for every possible thing that may go wrong, as there will be limited opportunity to revisit the quote. You may even include a provision to account for two or three change requests – extra costs your technology partners will invariably present whenever you need to deviate from the original requirements.

As an aside, it is these issues that contribute to the eye-watering quotes that makes Shadow IT so attractive.

Creating enterprise-wide buy in

Extending your well-functioning Cloud Operating Model to other IT teams can address the Company Toll. By introducing product management talent, these IT teams will want to offer features rather than technology gizmos. Introducing cloud business cases funds these teams to deliver features, rather than use initiatives to square ledgers. By introducing a cloud shared responsibility model, these teams will be service providers attracting service users rather than turning them away.

We're talking primarily about your network and security teams, followed by your managed service providers and on-premises teams. Shifting these teams will be a challenge for different reasons.

Your network and security teams will likely push back, explaining their domain is different to "cloud" and cannot possibly be provided on a self-serve, on-demand, pay-per-use basis. Your managed service providers, used to making most of their profit via change requests will see pay-per-use as a challenge to their business model. Finally, some existing teams may see the Cloud Operating Model as a challenge to their empire, particularly the change in roles and responsibilities. We'll address each of these issues in future chapters.

Extending your Cloud Operating Model to other teams to eliminate the Company Toll is a significant undertaking, but is achievable if you know how to address the challenges.

Symptom #7:
Your IT-project pipeline is progressing slowly

Is your IT team waiting for a network firewall burn or to spin up a server, which is impacting your project – costing you time and money as your engineers wait around? If so, you are suffering from a serious condition: the Company Queue. Let's explore what it is and what you can do about it.

The costly waiting game

No matter where you are on your cloud journey, company queues indicate your organisation is not making the most of its cloud investment.

Company queues are a natural outcome of the complex RACIs we discussed in a previous chapter, *shared responsibilities*. These RACIs determine who is responsible for *what*, as well as *how*, teams are engaged to fulfil a request.

As noted in that discussion, project teams are generally required to engage an expert who will perform the firewall burn, the server request, or the service-integration build on behalf of the project. Since most organisations only have a limited number of experts that can perform these functions, project teams must wait their turn to access the service, thereby waiting in the Company Queue.

Waiting in a queue is bad enough, but if the queue is not first-come, first-serve, then the problem compounds. This often happens when projects are bumped by a more “strategic” initiative. This leads to frustration, escalations, and attempts to circumvent the process as project teams leverage connections or look to position themselves as more “strategic” than others.

Making matters worse, organisations may have a Gate Keeper who, due to their unique position, can determine who has priority to their IT service. We will discuss this problem in a future chapter, suffice to say, priority disputes simply lead to more delays and increased costs.

The obvious solution is to break up the sequential nature of the Company Queue, making it possible for project teams to access the service in parallel.

The self-serve, on-demand nature of cloud services makes this possible with the Cloud Shared Responsibility Model, that has the cloud provider responsible for operating the service and the cloud user using the service either via the console, CLI or API.

Understanding the user journey

In this chapter, we’ll take another step and introduce the cloud user journey. This is the journey the cloud user takes to make use of, or “consume”, the service offered by the cloud provider.

We explored cloud users, your IT product teams, in a previous chapter, but to quickly recap: Your IT product team is a mix of product and technical staff who deliver valuable features to the business.

The traditional IT journey for project teams is well-known:

- Kick off new project
- Request quotes from the IT (network, server, security ….) Teams
- Business case approval
- ***Engage the IT experts to fulfill the requests***
- Deliver the project
- Hand over to the operations team

In contrast, the cloud user journey for IT product teams is based on the product development model:

1. Kick off new product line
2. Request seed funding for initial feature development
3. Register with required cloud services
4. ***Product consumes cloud services via console, CLI or API***
5. Deliver the product features

The key difference is step 4; in the traditional journey, multiple project teams engaging a limited number of IT experts leads to the *Company Queue*. In the cloud user journey, each IT product team can consume the cloud service at the same time and, as a result, experience no waiting.

If your teams are waiting in a *Company Queue*, you're simply not getting the best value or performance from your cloud investment.

Symptom #8:
You have Gate Keepers who hinder progress

Does your team need buy-in from the IT department to get things done? Perhaps a manager responsible for a "strategic" service or someone who prioritises access to IT services? These individuals are Gate Keepers who often hinder organisations trying to realise the full value of cloud investments. Let's explore why this is a problem and what you can do to fix it.

The Gate Keeper

Gate Keepers are a legacy of traditional IT project delivery processes. As we discussed earlier, whenever a business requests a service not offered by IT, a new project is spun up to deliver it. From time to time, IT managers will see this as an opportunity to introduce technology that becomes "strategic" for the company. Familiar examples include:

- the "data lake", which places all company data in one spot,
- the "identity solution", whereby customers can use one login for all company services, or
- the "message broker", that allows data to be transmitted over a secure channel and single platform.

Of course, each of these examples have merit. The issue is not the solution; it is how the solution is delivered to the rest of the business. If the solution is delivered through traditional IT delivery models, there's a strong risk that the IT manager will build an empire around the technology, becoming the "technology owner". For some technology owners, responsibility for a "strategic" piece of technology provides some form of job security as they hold the keys to a service used by everyone in the business.

This isn't necessarily a problem, especially if the rest of the business can access the service on fair and equitable terms. Unfortunately, as discussed in our previous chapter on the Company Queue, if the business needs to engage experts to use the service, then the different business initiatives must be sequenced and prioritised.

This opens the door for disputes, as teams vie for priority. As with any large organisation, teams will make use of personal relationships, past agreements, and promises of future help to get ahead of the queue. This type of environment offers strong incentive for technology owners to weigh in on the priorities, becoming, in effect, Gate Keepers.

Working in these conditions reduces agility, leading to delays and cost increases. The solution, as we discussed in our previous chapter, is to introduce a cloud-user journey, so your IT Product teams can access the service in parallel, retiring the queue altogether.

Rethinking Roles and Responsibilities

The cloud-user journey is a core component of a well-functioning Cloud Operating Model. Gate Keepers are an impediment to the introduction of your cloud user journey, to your Cloud Operating Model, and to ultimately improving your organisation's agility.

However, there is a solution – transforming the Gate Keeper from technology owner to Product Manager, converting their "strategic" solution into a cloud service.

Why would Gate Keepers make this move? Change can be difficult, especially if those asked to change see no benefit in doing so. Furthermore, the Gate Keeper's team of experts may see the shift towards self-serve, on-demand service as a threat to their jobs. After all, other teams have no choice but to use them to get things done, so what will they do if Product IT teams can self-serve? We'll discuss this challenge in a future chapter.

Organisational change like this is easier said than done. It requires the implementation of a robust change management program, able to explain:

1. Why the business needs to change. In this case, the improved agility that comes with implementing a Cloud Operating Model.
2. How individuals can take advantage of the change. For example, taking on more responsibility with better remuneration, or increasing capabilities and marketable skills.
3. The consequences if change does not occur. For example, loss of market share, reduced profits, or job losses.

The Gate Keeper is an example of a legacy IT department role that does not fit in an organisation moving to a Cloud Operating Model.

Symptom #9:
The DevOps Divide

Are your IT teams still divided between engineering and operations? Alternatively, has your organisation tried to implement a DevOps team, but not seen the responsiveness you were expecting? DevOps is critical for successfully implementing your cloud strategy, without it, you won't realise the agility that should come with your cloud investment. In this chapter, we'll explore the value of DevOps and the steps you can take to make it work for your organisation.

Competing internal incentives

Traditional IT project delivery splits engineering (more recently, "development") from operations. Engineering builds the service while operations deliver the service. This has a certain logic, as most traditional IT initiatives require long lead times to account for procuring and configuring infrastructure, up-skilling teams on new technology, and working through the roles and responsibilities. In this situation, you do not want your operations team wasting time and energy better spent operating systems that are already in production.

However, splitting engineering from operations creates a big drawback – different incentives. Since the engineering team is part of a larger project team, their work is time-boxed to finish before the project ends and funds run out. On the other hand, the operations team only want to take responsibility for a new service after all operational issues have been addressed.

These competing incentives create tension, particularly as the project deadline nears. Except for unicorn projects that run on time and on budget, engineering teams will start to produce sub-optimal implementations, or postpone less important features, leaving what is sometimes referred to as "tech debt". Operations teams will, of course, push back as they don't want to be responsible for the debt. This back and forth means the project will be subject to escalations, time wasting, and finger pointing.

Aligning development and operations

To address competing incentives, most organisations create an Operational Readiness Checklist (ORC) when an IT service is handed from engineering to operations. ORC is a contract between engineering and operations that defines responsibilities. In some organisations, this checklist can grow extraordinarily long after numerous battles between engineering and operations. Getting sign-off on the ORC is a vital goal for the engineering team, because once they've thrown the service over the fence, it becomes the operations team's problem.

Caught in the crossfire between engineering and operations are the customers of the service – your organisation! Battles between engineering and operations can severely impact business agility, making it difficult to remain competitive in your market.

One obvious solution is to combine engineering and operations into a single DevOps team. Some organisations have gone down this route. Unfortunately, combining teams without changing their incentives will not provide the agility you need. For DevOps teams to succeed, the incentives of all team members – engineering and operations – must be aligned.

The first step to align incentives is addressing the fundamental issues that divide your teams: the need to wait for infrastructure to be built, training, and sorting out roles and responsibilities. We have discussed solutions to these issues in previous chapters. These solutions form part of a well-implemented Cloud Operating Model that exposes the self-serve, on-demand value proposition of your cloud investment, includes a modern cloud user journey, and defines a Cloud Shared Responsibility Model.

Once the fundamental issues are addressed, the next step is to replace the time-boxed IT project delivery model with the evergreen product development model. We discussed the product development model in a previous chapter. In short, the model is built around continually funding feature development – not dissimilar to business-as-usual (BAU) funding for traditional operations teams. Once the project delivery time-box is removed, incentives can be aligned, and your organisation's DevOps model can succeed.

If you are not seeing the expected agility from your DevOps teams, or your teams are still split between engineering and operations, then you're not realising the full value of your cloud investment. Aligning incentives to the product development model is critical to a successful cloud operating model.

Symptom #10:
Your teams are still jumping through approval hoops

Are your IT teams still required to jump through hoops to get approvals - from an architecture review board, financial controllers, to security and operations – before continuing their work? Do these approvals take time, cause disruption and leave your teams frustrated? If your teams are still jumping through Approval Hoops to get their work done, then this is a sure sign you are not realising the full value from your cloud investment. In this chapter we explain why and what you can do about it.

Let's be honest: Nobody wants to be responsible for the company website going down for six hours during the busy season; or be responsible for the company logo splashed across the news following a customer data leak; or be the one to explain to management why the budget for the new-fangled IT solution needs to be doubled before it can be deployed to production.

To reduce the likelihood of such events, IT organisations have added approval steps to handle change including:

- architectural approvals to ensure that the designs meet company standards,
- security approvals to ensure that customer or sensitive data is not exposed to hackers,
- financial approvals to ensure that funds can be released.

We discussed one of the more significant approval steps, the operational readiness checklist, in our previous chapter. In that chapter our solution to that Approval Hoop was to align incentives across engineering and operations teams. In this chapter, we will look at how we can codify and automate approvals to reduce the frustration and disruption that comes with most approval hoops.

Approval Hoops in most companies generally take the form of a gating process, where stakeholders are given the opportunity to review a particular proposal. In the usual run of events, project teams are required to present documentation, estimates & designs for their proposals to company principals who then examine the information and provide feedback. This feedback is then integrated into the proposals before final approval is given, and the project team can continue.

While the process sounds innocuous, Approval Hoops can be the source of serious disruption due to time delays. For example, simply preparing the documentation, ensuring the principals have read it, and then organising the meeting to tick off the proposal can sometimes take weeks. If someone hasn't read the documentation, can't make the meeting or has serious concerns with the proposal, then the project may be delayed even further.

While disruption is a serious problem for Approval Hoops, there is a second, less obvious problem. Since Approval Hoops take so long, they generally take place only once - during the initial phase of a project. Subsequent changes do not receive the same level of scrutiny, and as a result bugs can creep in, exposing the organisation to operational, financial and security risks.

So, Approval Hoops are time consuming and generally happen only once. How does your cloud investment fit into this discussion? In a previous chapter, we examined the cloud user journey and how it differed from the traditional IT user journey. In the traditional IT journey, teams had to engage an expert to get something done, whereas in the cloud user journey, Product IT teams can consume the service themselves via a GUI, console or API.

For a Product IT team to take full advantage of these automated services, they must eventually implement a CI/CD ("Continuous Integration / Continuous Delivery") pipeline. CI/CD pipelines have been used by DevOps teams to manage code changes and code deployment for years. There are many excellent articles explaining the benefits of these pipelines and how to use them.

For our purposes, it is enough to say that sophisticated pipelines can now include operational, financial and security controls in order to codify and automate many of the Approval Hoops that previously used a gating process. As a result, Approval Hoops no longer need to take days or weeks to complete, but minutes or hours and can be performed every time there is a change rather than just once at the start of the project.

In other words, codifying and automating Approval Hoops to be used in CI/CD pipelines can reduce disruption to feature delivery as well as the operational, financial & security risks to your organisation.

If your teams are still expected to jump through hoops for approvals, this is a clear symptom that you're not realising the full value of your cloud investment

Symptom #11:
Cloud metrics from cloud services

Did your cloud operations team smash through another 50 service tickets last sprint? Is the team now turning around these tickets in hours and days rather than weeks and months? In the traditional IT world, closing more and more service tickets is a cause for celebration, but if your organisation is making a move to cloud, it is a cause for concern. In this article we will discuss how the measure of success in the old world is, in many instances, no longer applicable and may in fact be driving bad behaviour as you implement your cloud strategy.

Service tickets are the bane of operations teams. Resetting a password, adding a new user, provisioning a virtual machine; they are endless. But if these tickets are not closed in a timely fashion, then users start to complain. In many companies, operation team members have found ways to improve efficiency when it comes to managing tickets, such as writing scripts to automate some of the process or to batch up similar jobs and complete them all at once.

At the same time, companies are introducing project management techniques from newer frameworks like Agile; so, teams are working within sprints, tasks are pushed to the done column by sprint-end and scrum masters are carefully watching burn down rates and completed story points.

These new techniques fit nicely with the service ticketing approach; tickets are just tasks, and they can be measured as they move from the backlog to in progress until they finally land in the done column as closed tickets. Managers can set KPIs around number of tickets closed per sprint and easily measure results – holding individuals to account. Thus, operational staff have an incentive to find ways to automate the tasks so they can meet and beat their KPIs.

So, what's the problem? The problem is not that the operations team want to get better at closing tickets, the problem is the existence of the ticketing system itself. The system is a classic example of the Company Queue, a serious issue which we discussed in a previous chapter. Company Queues form when teams must engage an expert to get something done. Since companies have so few experts, teams must wait their turn resulting in delays and increased costs. In that chapter, we argued companies should introduce a cloud user journey, so that users engage an automated service via a GUI, CLI or API rather than an expert. With a cloud user journey, all teams can reset a password, add a new user or provision their own virtual server in parallel with other teams and as a result experience no wait time.

But if the KPIs for the cloud operations team is set around the number of tickets closed, what incentive does the team have to introduce a cloud user journey? Further, if the team manager is effectively a Gate Keeper, then it will be doubly difficult to make the mental shift required. In short, measuring service ticket velocity is a good example of a metric from the old IT world that helps drive bad behaviour in the cloud world.

So far we have looked at the different elements of a successful Cloud Operating Model, including:

- Dividing your organisation into service providers and users mediated by a Cloud Shared Responsibility Model,
- Introducing a product development model along with product management talent,
- Introducing a cloud business model to fund feature development,
- Introducing DevOps with the correct incentives and
- Codifying approvals in a sophisticated CI/CD pipeline.

In order to successfully implement your new Cloud Operating model, new metrics, *cloud* metrics, are required in order to measure success (or otherwise) and to help drive good behaviour.

For example, in the case of the cloud user journey, you're likely to focus on user experience metrics such as time on task, error rate as well as adoption and retention rates. For feature development, you will look at utilisation, feature requests & bug reports. For DevOps, you will need to consider DORA metrics to measure team performance.

If your cloud teams are still using metrics from the old IT world, then you're missing an opportunity to realise the full value of your cloud investment.

Symptom #12:
The Cloud Migration Gap

Are your cloud migration program costs mounting, is the promised "agility" failing to materialise, are you struggling to convince application owners to make the move? If so, then you may be experiencing the Cloud Migration Gap. In this chapter we will explain the Gap, why it is a problem, and what you can do about it.

Migration to the cloud is all the rage today. There is no shortage of articles explaining how to execute on your data centre exit strategy or why you should concentrate on value add and leave the undifferentiated heavy lifting to a cloud provider. There is also an abundance of material explaining how to build your business case, from how to calculate the benefits, to the innovation dividend and the ability to reduce costs by turning off unused servers and databases.

All these benefits are real, if they can be captured. So why is it taking so long and therefore costing so much? There may be several reasons, but by far the most likely, is your organisation is focused on migrating *applications* rather than *application teams*. This is a critical mistake, since it is your application teams who need to make the transition from traditional IT to Cloud if your organisation is to realise the agility you have been promised.

Let's look deeper. In the traditional IT world, how do your teams get things done? Typically they engage an expert – to build a server, to perform a firewall burn, to add storage, etc. In the cloud world, we expect teams to engage with a system – a console, a CLI, an API - and get things done for themselves. We discussed this in a previous chapter where we introduced the cloud user journey.

If your organisation is concentrating on migrating applications, then, without even realising it, you are likely replacing the expertise your teams rely on with a system they do not know how to use, creating the Cloud Migration Gap. Is it any wonder your application owners are pushing back?

The problem is particularly acute if you have engaged a cloud partner to handle the migrations. The partner, with deep cloud expertise, will have been tasked with re-engineering your on-premise applications to run on a cloud platform before throwing them over the fence to the application teams. This is the engineering/operations problem we discussed in a previous chapter – only compressed, as your cloud partner will be time-boxed on delivering the entire migration, not just one application.

Suffice to say, application teams will not readily accept these re-engineered applications without the expertise they have come to rely on, or an extended operational readiness checklist, or extensive training – or all three. In all cases, adding costs to your migration program and reducing the agility of your organisation.

To close the Cloud Migration Gap, you need to start by re-orientating your migration program around uplifting application teams - taking them on a journey from a world where they rely on expertise, to a world where they can act for themselves. In doing this they become the Product IT teams we discussed earlier on, and will bring about the agility your organisation has been promised.

Symptom #13:
The Cloud Change Program

Is there is a wall dividing your DevOps team from your traditional IT team? Do the teams sit in different areas, use different tools, and seemingly speak different languages? Do they even look at each other with mutual mistrust and suspicion? If so, then your cloud strategy could be at risk. It's important to resolve this quickly, otherwise it will impact the ability of your IT organisation to respond to changing business demands.

Introducing new technology into an organisation will always present challenges, particularly when it concerns up-skilling, ongoing support, and cloud technology. This can completely transform team roles and responsibilities and may ultimately create hurdles for your business.

Previously, we've discussed the need for Product IT teams with a specific talent profile, as well as introducing a Cloud Shared Responsibility Model to explain the new responsibilities.

These – among other changes – are needed to ensure your cloud migration is successful.

A case for change

However, it is not enough to introduce new technology, new team structures, new ways of working and expect to immediately reap the benefits. Before all this, you must first make the case for change and build confidence and support for the move. The case for change must be consistent with your overall cloud strategy and help answer questions such as:

Why should teams adopt cloud?

How will they benefit?

Will it help individuals with their career path, will it increase their work satisfaction, and will it help them build a profile within the organisation?

Will this take us to where we need to be as an organisation?

Why is there division within your team?

If your teams are still divided between DevOps and traditional IT, then chances are you have not made the case for change.

You can test this by asking some simple questions: Do your traditional IT teams look at cloud services like AWS, Azure & GCP and see just another data centre? Do they complain that the DevOps teams just don't understand the need for approvals, for a detailed RACI, for operational support? On the other hand, do your DevOps teams complain that the traditional IT teams are just too slow, can't change and don't know how modern IT works?

An IT department split like this may start to become a burden rather than an asset to the rest of your business. You will be forced to decide between the old and the new. Moreover, your business will see different results depending on which type of team you champion.

Uniting your DevOps and IT teams

To correct this situation, it is important to revisit the reasons you adopted your cloud strategy in the first place. Remind yourself of the cloud value proposition - self-service, on-demand, pay-per-user and what this can do for your business - the ability to respond quickly to changing market circumstances as well as test new ideas without the need for a large business case or significant investment.

Implementing these changes will pose different challenges to different parts of your organisation. We have discussed the potential impact on network and security teams and reimagining them as service providers. We have looked at promoting your technology owners from gatekeepers to product owners, and we have walked through implementing a migration journey for your traditional IT teams.

A robust change program is a critical component of your cloud strategy that will allow you to:

- Provide transparency and explain to the rest of your organisation how the changes will benefit the business.
- Build the support you will need to align all parts of your IT organisation around your cloud strategy.
- Address any questions and potential pushback around the new Cloud Operating Model.

A successful cloud strategy requires strong leadership and an actionable cloud change program that encompasses your entire IT and business organisation, and that supports your people to engage, understand, adopt, and utilise new ways of working.

Symptom #14:
Not another Change Request

Change Request, contract variation, scope change – do these words leave you with a sense of dread? Do your IT partners have a glint in their eye when you start to discuss new requirements? If you're making a move to cloud but still need to pay for yet another variation invoice, then chances are you're not taking full advantage of your cloud strategy. In this chapter we discuss how to reset your relationship with your suppliers and leave change requests behind.

Engaging partners to supply services to your organisation is a standard part of the IT toolkit. In the traditional IT world, engagement generally starts with a review of the market. This quite often takes the form of an RFI – Request for Information. This is then followed by an RFP – Request for Proposal - and finishing with an RFQ – Request for Quote. The key characteristic of the RFI/RFP/RFQ processes is that, in the traditional IT world, the IT organisation will engage the partner to build and then possibly operate the service on behalf of the organisation for an agreed, upfront price.

In previous chapters, we have discussed how many IT organisations hide the value proposition of cloud services behind traditional IT processes and tools. We have learnt we can address this problem by exposing the self-serve, on-demand, pay-per-use nature of cloud to the rest of the organisation by implementing a well-formed Cloud Operating Model. This starts with the creation of Product teams and includes replacing your traditional IT project model with a product development model, adding product management talent, realigning incentives, and using cloud business cases that fund feature development.

However, if you are still paying vendors for change requests, then something is not quite right with your Operating Model. The most likely reason is that you are still engaging your vendors under the traditional IT model, that is, you have requested proposals, you have obtained quotes and you are now working with your preferred supplier who will deliver exactly to the agreed specifications – and no more. Incidentally, vendors have become good at gaming this engagement model, specifically by underquoting to win the business in the not unreasonable expectation that customers will pay for variations to account for unforeseen issues – change requests.

To end change requests and the pain that comes with requesting more funding, you need to integrate your vendor engagement model into your Cloud Operating Model; just as you expect your own Product teams to deliver on-demand, self-serve, per-use services, you should expect the same from your IT vendors.

In other words, rather than RFIs, your teams should be able to perform a google search and learn everything there is to know about vendor services. Rather than RFPs, your team should be able to sign up to potential vendor services and test them immediately. Rather than RFQs, your team should be able to estimate costs by calculating expected usage against the vendor's consumption-based pricing model.

Once engaged with the vendor, if the service does not quite match your requirements, rather than make costly change requests, you should be able to make feature requests or submit bug reports. You should expect your vendor to incorporate feature updates or bug fixes into the per-use cost of the service - not an add-on cost.

Vendor services come in different flavours, IaaS, PaaS, iPaaS, SaaS and they all share the same cloud attributes - self-serve, on-demand and pay-per-use. To this list, we can add one more, *evergreen*, where the features of the service are always up to date and the cost for any new feature or fix is incorporated into the pay-per-use price. IT vendors that offer a self-serve, on-demand, pay-per-use and *evergreen* service, will help rid your organisation of costly change requests and move your organisation even closer to

By choosing IT vendors that offer a self-serve, on-demand, pay-per-use and *evergreen* service, you can rid your organisation of costly change requests and continue to optimise your Cloud Operating Model. If adopted along with the other elements we have discussed in this book - product development, product IT teams, cloud user journey, cloud business case, cloud shared responsibility model, cloud metrics and DevOps - you will have a scalable and sustainable Cloud Operating Model that can deliver your organisations strategic imperatives.

If you recognise the symptoms discussed in this book and are looking for a way to move forward in your cloud journey, come and speak to the cloud experts at DigiRen. While every organisation and cloud journey is different, we will work with you to figure out where you are in your journey, and help design a roadmap that will take you to where you want to be.

DigiRen is focused and highly specialised in cloud strategy creation and evolution, cloud operating model design and implementation, and cloud program orchestration. Our experience working in and with hyperscale cloud vendors on some of the largest and most complex cloud migrations, has provided us unique insights on how large organisations can achieve velocity, value and scale in their cloud programs.

This is the sole focus of our business and our highly skilled cloud professionals. We are experienced at working across business and technology domains, engaging executive stakeholders and deep technical experts, and we are passionate about helping customers realise the value from their cloud program investment. To find out how we can help reset your cloud journey, reach out to us directly at solutions@digiren.com.au or through your cloud vendor partner.

DigiRen