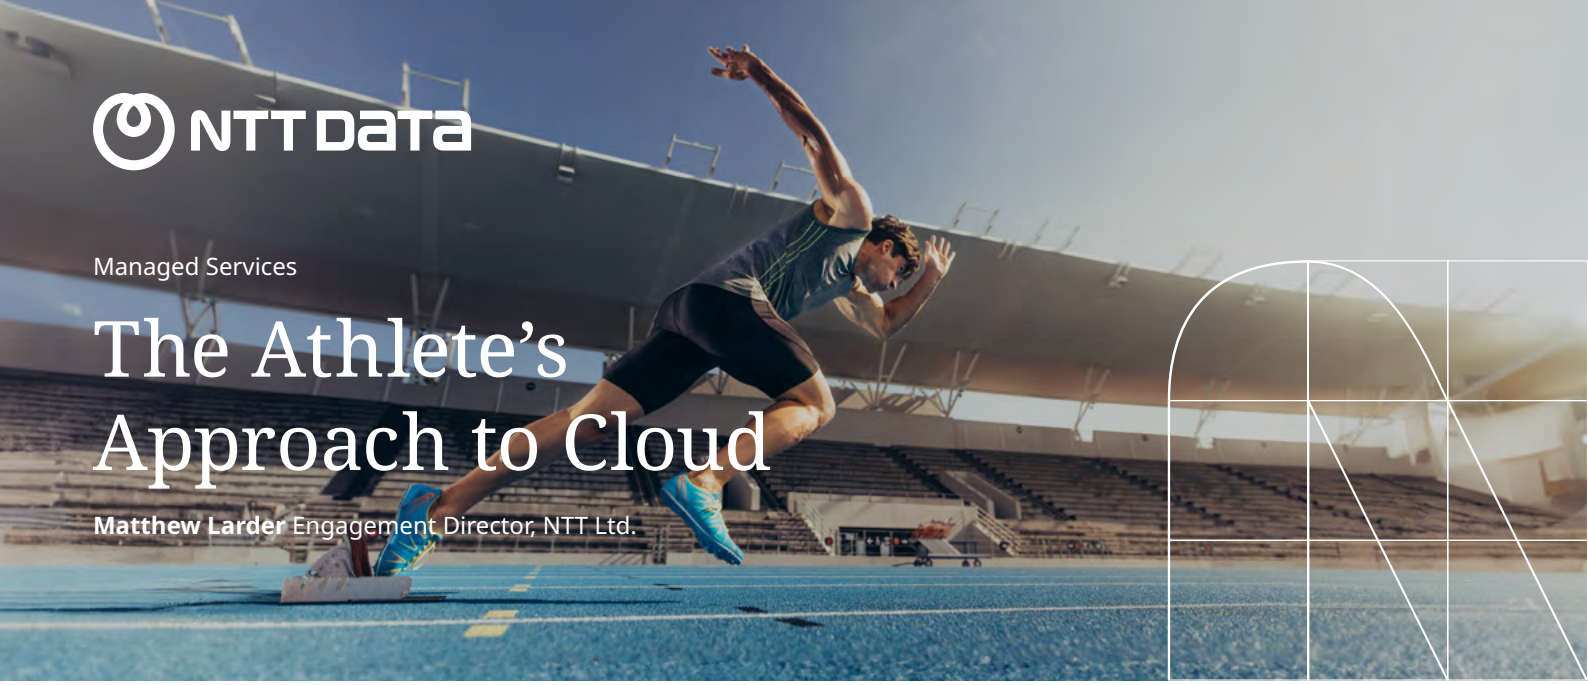

Managed Services

The Athlete's Approach to Cloud

Matthew Larder Engagement Director, NTT Ltd.

In my work at NTT Ltd. helping clients with cloud strategy, I frequently hear questions such as:

- Is public cloud cheaper or more expensive than private cloud?
- What are the hidden costs of public cloud?
- Is there a reliable way to control costs in public cloud?
- How should I change my applications to optimize them for the cloud?
- Can I make them faster, more reliable and cheaper to maintain?
- What are the triggers for migrating to public cloud?

In my opinion, to answer these questions we first need to understand the past experiences that have driven individual and business behavior in designing and operating IT systems.

Almost every decision we make in present-day life is based on habits. David T. Neal, Wendy Wood and Jeffrey M. Quinn, who collectively authored 'Habits — Repeat Performance' surmised that:

> "Habits are response dispositions that are activated automatically by the context cues that co-occurred with responses during past performance. Experience sampling diary studies indicate that **much of everyday action is characterized by habitual repetition.**

What happens if we make new decisions using experience, where the context for making the decision in the present is different to that of the past? We risk setting ourselves on a path towards a flawed future. In this thought leadership, I want to explore some key areas that I believe help to unlock an alternative approach to deciding how we design, operate and evolve IT in the present and in the future.

The challenge

There's always one...

The gap between past and future creates a challenge in understanding how you shape the present. The traditional approach to building IT includes attempting a prediction of what the future may hold.

Many organizations have deep roots in buying, building, leasing, managing data centers and hardware. With this experience comes a level of certainty and predictability. You've done it many times and so, you rely on this to convince you of the outcome.

However, there are fundamental flaws in this.

You may take space for 10 years and build hardware platforms with a 3-5 year refresh cycle in mind. This is on the basis that your business will grow by x percent per year and with this the requirement for IT will as well.

However, all too often we hear stories where a storage system which cost millions was only ever 40% utilized over its life, or on average a production server is only 35% utilized on any given day.

In many cases, organizations learn from this and try to take a balanced view of buying the right capacity for the job.

But in a world where your business is built on IT, there is only so much risk you can take before hardware lead times, predicted growth or requirement contingencies scare you into buying substantially more than you will ever need before the hardware becomes obsolete.

Also, the cost of hardware is being driven down all the time and this breeds a mindset of getting 'more bang for your buck' each time you provision IT services.

While in some cases this is due to individual organizations and their ability to negotiate based on volume, increasingly it is due to the competitive nature of public cloud services.

The 2019 Gartner Cloud IaaS Magic Quadrant compounds these points, where eight cloud service providers were dropped; leaving just six (AWS, Microsoft, Google, Oracle, Alibaba and IBM), the first two commanding the lion's share of the IaaS market.

The point here is to understand who is making the decisions, based on what experience and for what desired outcome? This is always driven by the different priorities of individuals within an organization and the organization as a whole.

The views I commonly see include:

The business person

At the 'sharp end' of the organization, identifies customers and their requirements. This person is tasked with establishing a short-medium/long-term plan with the customer/user, including the demand over a period of time.

The IT person

Understands the requirements of their customers in the business and how these can be translated into IT requirements. This person must ensure IT is deployed to meet demand with a layer of 'headroom' to ensure it is fit for purpose.

The finance person

Juggles the cost to deliver business requirements with the business revenue or efficiency optimizations and associated profit expectations. This person performs calculations to understand how long it will take to depreciate assets followed by the cost to continue operating these assets before the next investment is required.

The analogy of 'I do'

To help put this in context and for those of you who have experienced the wonderful moment of 'I do', let's compare this decision-making to wedding planning:

- How many evening guests will be attending?
- Will food be provided for each person?
- What kind of food your guests would like: bacon butties, gourmet burgers or beluga caviar?

The challenge is traditional evening reception venues do not cater on-demand for a large volume of guests. Prior planning is required to ensure there is enough food available on the day. To address this dilemma, your wedding planner will suggest you cater food for each person or as a percentage of the total, for example, 75%.

- How many evening guests will turn up?
- Will each person consume their equivalent 'head' allowance for food?
- Should you cater food for each person?
- Will there be any uninvited guests making an appearance?

All of these impact the volume of people for which you need to cater.

I can hear you thinking, are we still discussing IT or weddings?

Well, this challenge is a balance of risk and cost in arranging enough resource (food) for all applications (evening guests), while ensuring you minimize commercial risk of having too much (food waste) or too little, causing application downtime or performance impact (unhappy guests).

Reviewing these points and swapping the wedding analogy for IT resources, it becomes apparent that many organizations who build and operate their own IT services are balancing requirement accuracy, risk and predictability.

When working with public cloud, much of this responsibility can be offset to the provider who has built an IT service for the mass market, not for the individual needs of each organization.

However, in my experience, IT provisioned on public cloud is often overprovisioned just like on private cloud and this is due to the illusion of predictability. Predicting the future is a mixture of:

- **Science:** based on historical data — but is this right for the present or future?
- **Art:** the well-documented 'finger in the air' which many IT architects have employed at some point in time.

> "We must be open to **adapt current processes and accept a new procedure.**

- **Negotiation:** the power of the organization and its employees to agree on a model which works for them all.

Therefore, to reach the most efficient outcome (economics, reliability and scale) we must be open to adapt current processes and accept a new procedure which may require slightly less control than we are used to, with varying degrees of risk.

Applying the athlete's mindset to cloud

To some people, the uncertainty around public cloud use is driven by the absence of the 'see it, touch it' element of traditionally-hosted IT- and its inherent flexibility can result in cost-creep and infrastructure sprawl. We sometimes perceive these risks as the removal of predictability.

A traditional mindset to provisioning IT resources is finite, you buy so much at a fixed cost with minimal if any variation in the cost to operate. This implies a degree of control.

Moving to at a minimum public cloud IaaS should involve relinquishing some of the responsibility and control present in a traditional world.

In my experience, this must be embraced.

Picture any professional athlete. The best golfers do not hold a tight, rigid grip on their club. Instead, a golfer maintains the right balance of control and flexibility as the key to momentum.

Athletes share common ingredients that an adopter of public cloud would be wise to adopt. As Ben Hogan said:

> "Reverse every natural instinct and **do the opposite of what you are inclined to do,** and you will probably come very close to having a perfect golf swing.

If we recognize that what we deploy into public cloud will likely need to evolve on day two, then we will be prepared to support such a change. But preparation is key, as otherwise you will provision IT services on public cloud the same way as you provision it in a traditional environment and this may even lead to costs rising rather than falling.

Think of the present

The nature of strategy implies long term, but this doesn't have to mean designing IT systems to be the same in five years' time. A key value of public cloud is that it is designed to support change through flexibility.

Public cloud is designed to support change through flexibility.

Based on a 2018 survey by the Project Management Institute:

> "Almost three-quarters (71%) of organizations **report using Agile approaches sometimes, often, or always.**

This promotes flexibility and ultimately does not change the purpose of the application but helps to evolve how that purpose is achieved.

Many organizations looking to consolidate systems and adopt public cloud, may be familiar with the five Rs to application transformation — replatform, retire, refresh, re-architect and retain.

Re-architecting (often referred to as refactoring) is the process of changing an application's architecture to take advantage of native public cloud features as part of its move — a nice medium compared to refreshing an application entirely. Re-architecting still requires valuable time, effort and resources.

The alternative is to re-platform (often referred to as a re-host or 'lift and shift') the application as-is to the cloud.

This requires less time and effort while enabling you to take advantage of on-demand resources to re-architect the application, with no long-term fixed commitment to those resources. if you make a mistake along the way. In either scenario, there are some key principles which in my opinion help to maintain or improve flexibility, increase reliability and control or reduce cost.

Four key principles

Establish a set of cloud architecture guidelines

The beauty of public cloud is that it's there, ready and waiting. However, before you jump in at the deep end, take some time to think about how your organization will integrate with public cloud.

Explore how you will manage user identity to make the experience seamless.

- What operating model including roles and responsibilities works for your business?
- What security controls will you need in place and how will you name, then tag resources to ensure governance and avoid sprawl?
- Lastly, what service levels does the business require of IT?
- How will you achieve this in a public cloud world where infrastructure and applications can be architected to take advanced of native functionality, reducing traditional machinery and complexity?

Plan for a race, sprint to the finish

Many organizations have experienced multi-year contracts, where over time (short, medium and long term), requirements change, meaning the IT service no longer serves its purpose; but the nature of the contract makes it prohibitive to evolve.

If you are exploring a large-scale adoption, then take some time upfront to plan aspects such as a business case to understand what the total cost of ownership might be, but expect this will change as requirements do.

To manage this change, consider planning cloud adoption as a relay, with multiple sprints as part of an overall race to the finish.

Don't be afraid to start low

Recognize that your traditional on premises IT is most probably over-provisioned.

The same is often true for many public cloud systems.

- For low-criticality systems, start with provisioning a small, low-cost specification on public cloud.
- For high-criticality systems, explore the actual utilization before migration. Use any number of native tools at your disposal, then select an equivalent specification on public cloud.

Monitor, resize and right-size over the short term as needed.

> "Public cloud is designed to **support change through flexibility.**

Reserve away

Once you are confident the provisioned resources are right-sized, work with the native tools and also support technology to take advantage of committed use discounts (sometimes referred to as 'reservations'), available from most providers, to further reduce cost.

Conclusion

Following these suggestions will enable you to focus on the present, while keeping an eye on the future. This will start to alter your habits and, therefore, decision-making. This will help lay a foundation for cloud adoption, allowing an evolution of the organizational culture to plan for the present.

You will be able to:

Right-size IT for **the immediate business requirements,** while supporting evolution for future change. This will minimize the cost of IT, great for the finance person and the business person helping reduce cost to the customer.

Ensure IT is fit for purpose, simpler to maintain and capable of change, great for the IT person.

Accurately **track cloud charges** and ensure public cloud provisioned IT is right-sized, great for the finance person.

Next steps

NTT Ltd. offers free cloud strategy workshops to organizations looking to start or continue their strategic cloud transformation. If you would like to take advantage of this, please contact Matthew Larder (matthew.larder@global.ntt), or visit service.global.ntt


NTT DATA