White Paper

Cloud Migration for the UK Public Sector

Bridging Old and New One Step at a Time

Tim Skinner, NetApp

June 2021 | WP-7346-0621

TABLE OF CONTENTS

Foreword

The UK Government National Data Strategy consultation paper, published in 2020, demonstrates a growing appreciation for the inherent value of adopting good data practice and joined resources. This paper, coupled with the government's Cloud Guide for Public Sector, shows a strong drive from central government to realise the full digital potential of our public service.

However, such a drive does not negate the fact that, within a wider data strategy framework, cloud adoption and migration is an intimidating task. Much of the industry's discussion seems to jump from A to Z in the migration process; cloud and legacy systems are positioned as almost opposing forces, with departments even organising their resources in the same, uncoordinated way.

The reality we face is a complex array of interdependent systems that cannot simply move from A to Z; some of them may never move.

This situation is further compounded by the breadth of cloud services available — some 38,000 options to choose from, based on the 2020 G-Cloud tender. Other factors include interoperability limitations and a need to avoid technical dependencies with sole providers.

This paper explores the challenges that our public sector faces in achieving its digital and data-driven ambitions, for which it is to be commended, and in migrating to a cloud-based infrastructure. Although the goals outlined in the National Data Strategy do not implicitly depend on a move to cloud, and are only in consultation paper stage at present, decisions about storage, infrastructure, and access are intrinsically linked to such a move. Therefore this paper attempts to lay out realistic, incremental planning considerations for the next 3 to 5 years that could have significant impact on both service and cost efficiency for the treasury and the tax payer alike.

Tim Skinner, Director, Public Sector, NetApp

Definitions

Data fabric. An architecture and consistent set of data services that are independent of the underlying infrastructure, on premises and/or in multiple cloud environments.

Hybrid infrastructure. Uses a common set of tools, processes, and standards to deliver business services and resources from a combination of traditional on-premises IT and/or one or more public cloud providers.

Articulating the challenge

The UK Government National Data Strategy explicitly states that a need exists to "transform the way data is collected, managed, used and shared across government, including with the wider public sector, and create joined-up and interoperable data infrastructure."[1] Combined with its Cloud Guide, government has tasked the heads of IT across our public sector with making some of the most significant and challenging decisions of their careers. Those decisions include how to make a cost-effective, future-proofed, interoperable cloud migration strategy and enact it while minimising service interruption, technical debt, and vendor dependencies.

This challenge is further compounded by a rapidly evolving cloud market. Traditional IT vendors are venturing out of the data centre market into cloud services, while cloud providers are simultaneously moving into data centre real estate, each attempting to capture mind share. This market convergence is dominated by a small number of large hyperscalers providing infrastructure as a service (IaaS), predominately Amazon Web Services (AWS) and Microsoft Azure but also IBM and Google Cloud. Each of these companies is making eye-popping investments in data centres and establishing a global footprint in an effort to create unprecedented scale and agility.

As if that weren't enough, the breadth of available services, bolstered by an ecosystem of value-add service providers, is staggering. The 2020 G-Cloud tender had more than 38,000 different cloud services submitted to review. And while this extensive menu promotes competition, allowing small-to-medium enterprises to enter a once dominated arena of global service integrators, it doesn't make the decision-making process any easier.

None the less, headway is being made. According to a freedom of information (FOI) request made to governmental departments, agencies, and public bodies, we know that over half (53%) of respondent services are using hybrid cloud infrastructure in some form. Although this suggests notable progress, the majority of public service infrastructure is still on premises, with more than two-thirds of services currently less than 50% cloud-based. With 16% of respondents aiming to be fully cloud based in 3 years, and petabytes of data currently held on legacy systems, the inquiry suggests that a hybrid environment will be necessary for the foreseeable future.

This progress is to be expected. The adoption of relatively straightforward cloud-based solutions such as Microsoft Office 365, Teams, and even IT service management platforms like ServiceNow provided vital cloud efficiencies and quick wins. What remains, however, is the cloud migration and integration of far more complex systems, many of which have evolved over years and operate in isolation.

In a call for evidence for issues related to data foundations, cited as primary challenges were these complex legacy systems alongside varying, often incompatible systems for inputting and recording data. These issues need to be addressed within the current real estate, and also be considered as part of the decision-making process for future digital services and cloud adoption.

[1] Gov.uk https://www.gov.uk/government/publications/uk-national-data-strategy/national-data-strategy#data-2-3

Without developing a cloud migration journey that keeps data, its access, and its usability as a constant rather than an end point or goal, there is a risk that these islands of data and functionality will continue to exist. These islands result from information being stored on different platforms in different hosting environments with different service levels and data protection methods.

The challenge is multifaceted. Remaining agile and keeping options open is of paramount importance, as are realistic and achievable objectives. So if hybrid cloud is here to stay, it becomes a question of how to holistically manage data storage, access, interoperability, and mobility.

Understanding storage as a key element in achieving data access and agility

Data and resources already exist across complex ecosystems of in-house data centres, private/public clouds, and trusted delivery partners. Embracing new cloud technologies and attempting to place data at the heart of everything requires alignment of policy, governance, and risk management through software development and architectural disciplines, right down to the underlying IT infrastructure. Realising this data-led vision and shaping its future pathway ultimately comes down to how data is stored and processed at the very foundations of the IT stack.

The government recognises this need, stating that it "must expand work to treat data as a strategic asset, and create a whole-government, collectively responsible approach to investing in data foundations, so that everyone can benefit from the improved outcomes data can offer."[2] Given that data storage can make up as much as 45% of cloud consumption costs (and this figure is likely to grow as digital services mature), getting this process right is vital.

Fortunately, the message is getting through. When asked what the top priorities for cloud migration are and what public services must consider in their journey, 54% of respondents selected "improved organisational agility." Agility in this case was described as "the ability to move data around and take advantage of new commercial or technological opportunities to optimise costs and improve services." This need for smooth data portability and access is echoed by 46% of respondents, who selected "negating service interruption" as a key challenge for data strategy implementation. In essence, the big-picture objective is self-evident, but the actual how-to needs to be developed.

[2] Gov.uk https://www.gov.uk/government/publications/uk-national-data-strategy/national-data-strategy#data-4-2

Removing the “switch”

We have established the need for a cloud migration journey that keeps data, its access, and its usability as a constant rather than an end point or goal. This is in the context of hybrid, multicloud being the foreseeable reality and storage being the foundational rock. The hurdle here is avoiding technical debt and managing legacy technology build-up. The One Government Cloud Strategy highlights the need to prevent the accumulation of future legacy technology through continuous improvement planning; that is, iterative or phased migration. Phased migration, however, is often interpreted as selecting a single service or dataset and moving it wholesale to a new permanent location — a cloud.

The danger of this line of thinking – beyond the sizable challenge of an all-in-one switch – is that it encourages looking at migration tasks in siloes. Moving chunks from A to Z generally means adopting the nearest digital platform in order to meet an immediate demand. This approach ultimately prevents data from being managed and used collectively, while opening up new potential risks and stifling innovation and digital progress in the long term. The new platform becomes tomorrow’s legacy system.

“Switch language” also implies -whether intentionally or not- that these legacy and cloud systems run in parallel rather than together, as is often the case with cloud migration language. It is positioned as a binary change from old to new, obsolete to useful.

However, iterative and phased migration should be thought of as a need for interim steps, making cloud migration a journey rather than a moment. This makes the whole process far more digestible through a “stepping stone” approach, while also allowing easier control of technical debt, up-front investment, and data management. It also allows consideration of the question “what next?” since there will inevitably be opportunities for future innovation as technology and services evolve.

Bridging old and new, one step at a time, with the cloud bridge approach

In a sense, what is needed is a bridge to the cloud — a structured approach that is not a binary A to Z, rather A, B, C. Beyond the benefits already mentioned, this approach keeps options open while meeting current needs and usage demands. Not everything needs to be located in the same cloud — or even any cloud. The goal isn’t 100% cloud; it’s an interoperable, accessible system that meets today’s needs and is future-proofed for tomorrow.

There is a diverse range of legacy infrastructure and services at different stages of their lifecycle. There are stable long-term production systems that are secure, and these naturally lend themselves to virtualised platforms like hyperconverged-infrastructure. Other services will continue to be delivered on premises due to security constraints or for reasons of data sovereignty. However, there will also be interim, run-out systems or transitory data that are better suited to an interim storage subscription model (not to be confused with a consumption model; that is, as a service in the cloud). The interim storage subscription model occupies fixed levels of rented resources until the cloud migration strategy, decisions, and personnel are available for the next step.

Managing change – infrastructure modernisation

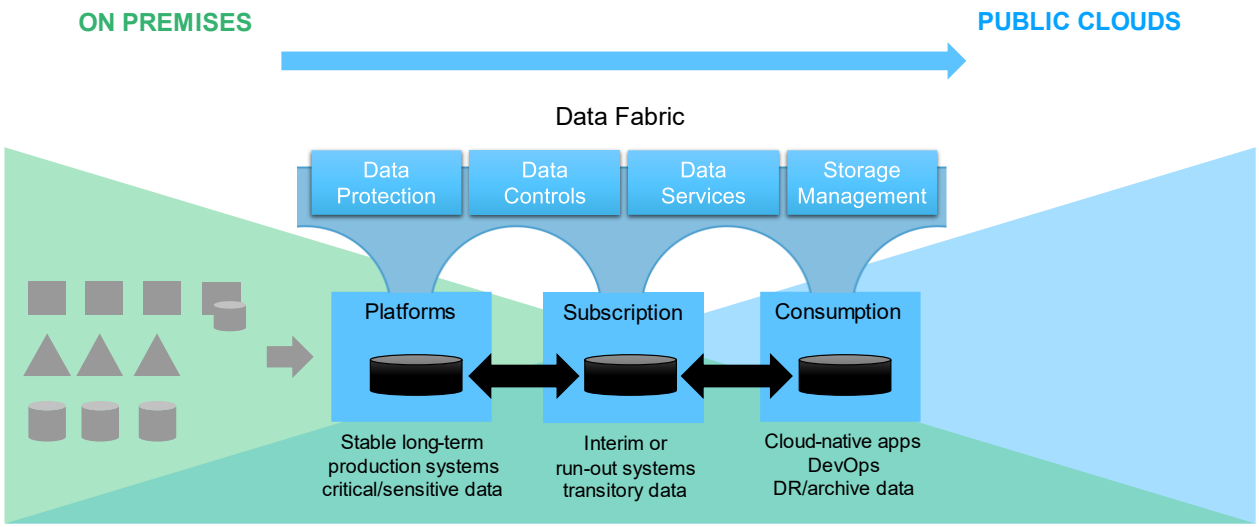

Figure 1) Bridging old and new.

The subscribed storage locations allow the legacy infrastructural challenges to be overcome in the migration process by abstracting the services above the storage layer. At the same time, the model prepares systems for the cloud migration and negates rushed cloud decisions, technical system dependencies build-up, and hyperscaler lock-in. Subscription models also help manage and effectively reduce technical debt build-up, so they make both technical and commercial sense compared with investing in capital or forcing a wholesale cloud migration.

The case for subscription over moving immediately to a consumption model (that is, the cloud) is that, until full systems are established with advanced optimisation capabilities, it's difficult to monitor and manage consumption costs. Consumption-based compute and storage are leading sources of overprovisioning, accounting for considerable overspending, with an entire service industry emerging to help customers tackle this challenge. Subscription costs, on the other hand, are controlled within certain bounds, and although they are not as optimised as a fully controlled consumption model, they are far more easily managed. In the interim, they enable cloud-native and cloud-ready apps and services to be developed.

Rather than imagining the journey as an evolution of legacy system to cloud, it is more helpful to think of this transition from a commercial point of view. That means an evolution from capital expenditure, through to subscription (fixed-term resource access), maturing into consumption (as a service) and ultimately automatic optimisation of resources across each of these pillars (see Figure 1).

With a data fabric architecture, the idea is that you can move data between each of these points, maintaining access controls and security, but abstracted from the underlying storage. Key to this process is delivering all the data protection, controls, services, and management in a consistent fashion across a hybrid cloud environment — that is, the bridge level, not the storage level.

Applications and services are then deployed at the bridge level, using industry-standard protocols and APIs, and can be moved across location as necessary.

Orchestrating the data fabric / bridge level

The bridge level, or rather the data fabric, is an architecture and set of data services that provide consistent data services that are independent of the underlying infrastructure, whether that is on premises, subscription, and/or in multiple cloud environments. In other words, it is a virtualised data abstraction layer with wide protocol and API support that provides a platform on which to deploy applications and services.

To deliver the greatest possible value to the organisation, the platform should be infrastructure agnostic where possible and fully application agnostic. This is the layer where data access controls, data services, and storage management are disaggregated from the underlying storage layer.

Deploying a data fabric as part of a standardised architecture also allows data to flow seamlessly from edge to core to cloud. This means that data from IoT devices, streaming services, and other connected entities can be securely ingested, distributed, analysed, and stored, regardless of the point of creation and across a geographically and operationally distributed workspace.

Ensuring long term optimisation

Cloud compute and storage expenses make up approximately 70% of any customer's spend. Once the move from subscription to consumption-based models is complete, managing expenses becomes an even higher priority. The FOI survey found that 43% of respondents were prioritising operating-cost control and flexibility of service. Unsurprisingly, a large proportion of compute and storage is wasted due to poor optimisation capabilities, accounting for as much as $17 billion in wasted global spend in 2020.

"With this form of continuous optimisation, the government could make incremental savings of more than £30 million per annum."

Among the tools and services available to help manage this waste, a leading solution is the use of application-driven infrastructure. ADI uses analytics and machine learning to translate and understand applications' workload patterns. This makes it possible to automatically provision and scale back resources, delivering optimal performance and availability while minimising the cost of storage and compute. ADI does this within the pre-established bounds of the contracted SLAs and SLOs.

This approach in turn helps application teams shorten development lifecycles and run more applications in their choice of cloud. Customers can save up to 90% of their compute and storage infrastructure expenses through this model. NetApp's own internal modelling suggests that with this form of continuous optimisation the government could make incremental savings of more than £30 million annually.

Conclusion

Cloud migration may seem like a destination, but it is actually a journey. Yes, there are easy wins and potential quick switches. However, in the context of the government's broader data strategy and vision for our public service's infrastructure, rushing or forcing cloud migrations could be detrimental in the long term.

The cloud bridge approach allows a more manageable and cost-effective pathway to the cloud, with clear steps to ensure maximum roll-out flexibility and minimal impact on service delivery. By adopting this model, our public sector can be confident that its infrastructure is meeting the expectations of today and is future-proofed for tomorrow.
