Disruptive Innovators Network

Adopting cloud computing:

a guide for executive teams

By Ross Fraser with Phil Brunkard

Disruptive Innovators Network

October 2020

In association with

aws rackspace

AWS UK & Ireland Migration Partner of the Year 2020

Introduction

Digital transformation is part of the business strategy of all social housing providers and has been for several years.

Increasingly, cloud computing is becoming a key driver and enabler of digital transformation.

For example, housing providers which have previously migrated operations to the cloud have, in responding to the COVID-19 epidemic, made the transition to remote working and digital service delivery more easily and effectively.

This guide is the first in-depth analysis of cloud computing produced for social housing and it written for housing association and ALMO Chief Executives, local authority housing directors and their respective management teams and Boards.

The objective of the guide is to enable housing sector business leaders to have a clear understanding of cloud computing – enough to enable them to engage with IT colleagues and cloud suppliers and ensure that their business develops a strategy in respect of the cloud.

This guide sets out to meet that challenge by explaining the potential of the cloud, the business case for adopting it, its costs and risks, cultural barriers to overcome and where to get support for the change process.

To make the report easy to read, technical jargon has been kept to a minimum.

The key theme of this guide is that the growth of cloud computing is inevitable and housing providers need to consider the opportunities that it brings to the sector.

Ross Fraser,
Director of Research, Disruptive Innovators Network

Contents

1. The future: digital transformation and cloud computing

1.1 Digital transformation and disruption

Few, if any, business sectors are untouched by digital transformation. Most organisations are increasing spend on IT with a view to growing the proportion of business operations delivered digitally.

The global research and advisory firm Forrester describe digital transformation as follows:

Digital transformation is not just about technology. It's the necessary but challenging journey of operating digital-first with the speed and nimbleness to change rapidly, exploit technology to create lean operations, and free people to do more complex tasks.

In other words, the true challenge of digital transformation is not the digitisation of existing processes but a whole transformation of business operations.

The Organisation for Economic Co-operation and Development (OECD) states[1] that:

Digital transformation is having a wide-ranging impact on the business environment, creating both opportunities and challenges. Inter-related trends such as e-commerce, big data, machine learning and artificial intelligence (AI), and the Internet of Things (IoT) could lead to large productivity gains for the economy.

Disruption to existing business and social models, as well as established markets, will disrupt the lives of millions of citizens.

Digital transformation and the future of work

The World Economic Forum states[2] that:

- digitalisation could create up to 6 million jobs worldwide by 2025
- elsewhere, automation will displace many human beings. Current estimates of global job losses due to digitalisation range from 2 million to 2 billion by 2030

The Forum goes on to say that, following digital transformation, there will be three types of jobs, categorised by the percentage of codifiable tasks within the role:

- those that will disappear (lost the race against the machine). For example, clerks and administrative staff
- those that are in collaboration with machines/algorithms (run with the machine). For example, those professions that rely on cognitive and social capabilities, such as doctors/surgeons
- those jobs that are completely new or remain largely untouched (running faster than the machine or running a different race). For example, roles in the creative arts are unlikely to be automated soon

The World Economic Forum concludes that:

- *"there will be both winners and losers –while the net impact on jobs in some industries could be positive, many sectors will experience job losses*
- *a big premium rests on the near-term ability of businesses to upskill employees and shape the next generation of talent for the machine age"*

Environmental sustainability

The World Economic Forum predicts that, in theory, *"digital initiatives in the industries we have examined could deliver an estimated 26 billion tonnes of net avoided CO2 emissions from 2016 to 2025. This is almost equivalent to the CO2 emitted by all of Europe across that time period."*

However, the Forum is cautious about whether this impact will be achieved, stating that "*ensuring this potential value can be realised and scaled means overcoming hurdles relating to the acceptance of new, circular business models, customer adoption and the environmental impact of digital technology itself."*

[1] http://www.oecd.org/sti/ind/digital-transformation-business-sector-summary.pdf

[2] https://reports.weforum.org/digital-transformation/understanding-the-impact-of-digitalization-on-society

1.2 The advance of cloud

A key enabler of the digital revolution is cloud computing.

The term 'cloud computing' refers to the on-demand delivery of IT resources via the Internet with pay-as-you-go pricing. Instead of buying, owning, and maintaining data centres and servers, organisations can acquire technology such as compute power, storage, databases, and other services on an 'as-needed' basis. It is similar to how consumers flip a switch to turn on the lights in their home, and the power company sends electricity. With cloud computing, providers such as AWS manage and maintain the infrastructure in a secure environment and businesses access these resources via the Internet to develop and run their applications. Capacity can grow or shrink instantly and businesses only pay for what they use.

Cloud computing usage is set to increase exponentially in business over the coming years. Research and advisory firm Gartner predict that worldwide, the public cloud service market will grow from £140 billion in 2018 to £254 billion in 2022, attaining a compound annual growth rate (CAGR) of 12.6%.

The cloud service provider market leaders are Amazon Web Services (AWS), Google and Microsoft Azure - often referred to as the *hyperscalers.*

Fig 1 - Magic Quadrant for Cloud Infrastructure as a Service, Worldwide[3]

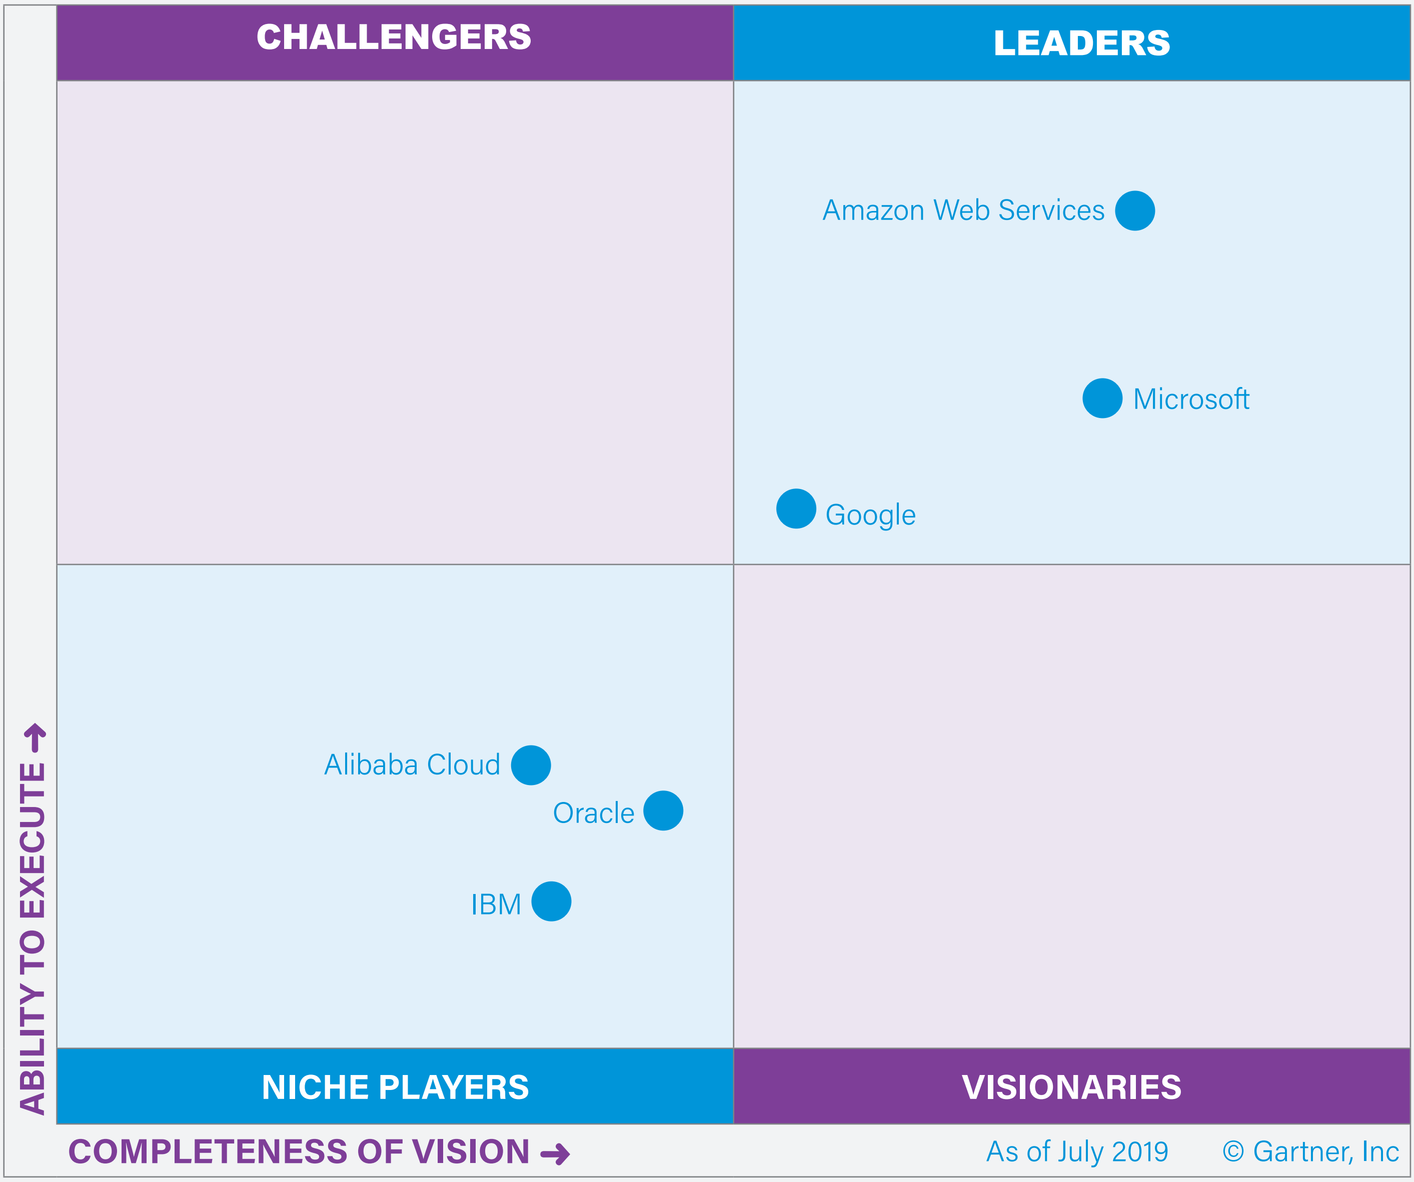

[3] https://www.gartner.com/doc

There is a secondary market in support services:

- cloud providers and partner ***managed service providers***, like Rackspace, offer support services which guide organisations through the design, decision-making, implementation and future use of cloud services
- ***technology development partners*** build solutions in a cloud provider's environment. These solutions can cover hardware, connectivity and applications

Figure 2 – Magic Quadrant for Public Cloud Infrastructure Professional and Managed Services, Worldwide[4]

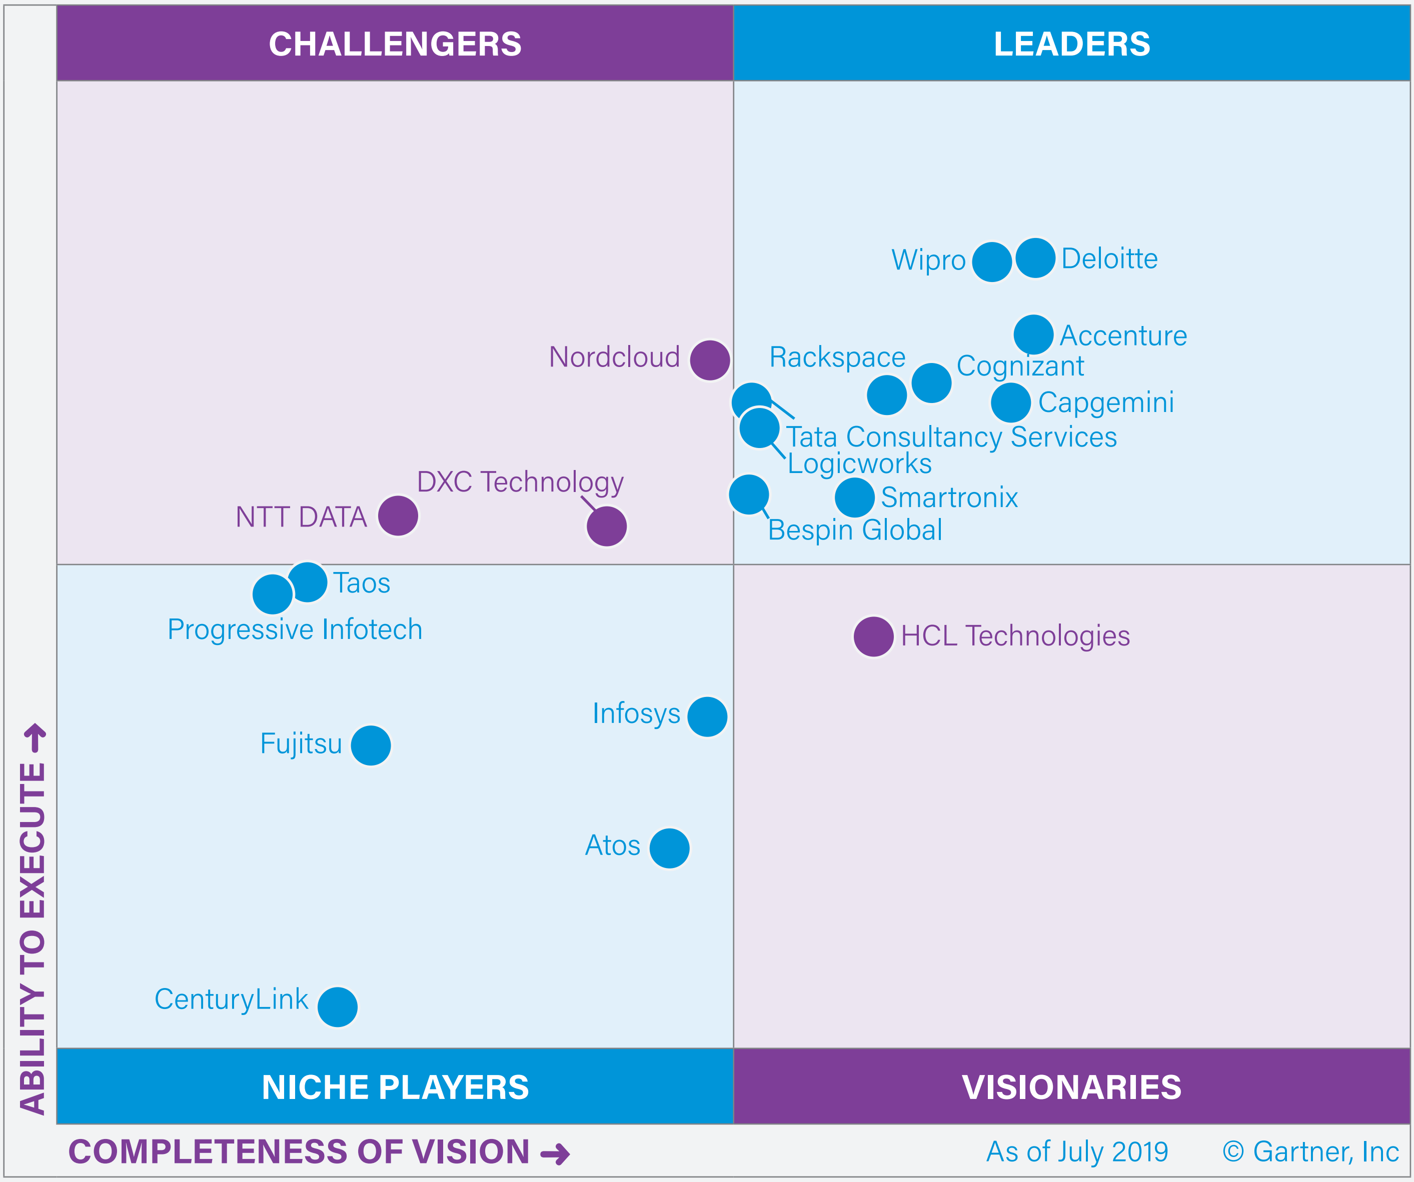

[4] https://www.gartner.com/doc

By working with accredited managed service and technical partners, cloud providers expand sales channels and stimulate new uses for their products without getting entangled in client deployment and delivery.

Most businesses will already be consuming cloud-related IT services. For example, Microsoft Office 365 and most of the recently developed customer relationship management systems (CRMs) and Customer Access Platforms (CAP's) operate within the cloud. Most applications accessed on mobile devices (such as Google or Facebook) are delivered from the cloud.

Businesses are likely to be drawn further into the cloud in future, through the increased development of 'cloud-native' applications - new business applications designed within a specific cloud environment. For example, Office 365 is accessed via the Microsoft Cloud. Sagemaker and its AI capability is accessed via the AWS cloud.

The growth in cloud computing is based on a series of end-user business benefits:

- *pay as you go* - sold on demand, typically via a specific time period, usually by the hour, the minute or (in the case of AWS) by the second. Charges cover central processing unit (CPU) usage, data transferred and so on
- *elasticity* - organisations can buy as much or as little of a service as they want at any given time and for any duration
- *enhanced data security* - improved ability to meet core security and compliance requirements, such as data locality, protection, and confidentiality with comprehensive services and features
- *improved business continuity* – serverless computing in the public cloud - see Appendix A.1 - enables the development of 'self-healing' applications
- *flexibility* - regarding the hosting or outsourcing of IT assets and their deployment or management
- *agility* – quick, easy and cheap to experiment with and deploy business applications in a cloud environment
- *automatic platform upgrades* – cloud users gain the benefits of automatic upgrades and improvements to the platform without having to invest time in doing this work themselves

Cloud providers contend that, as a result, businesses can achieve:

- increased agility to respond to changes in the operating environment
- improved service delivery experience for staff and customers
- reduced IT and operational costs

Businesses may wish to follow the lead of UK government and its Government Digital Service (GDS) and adopt a 'Cloud First' policy. 'Cloud First' means that organisations should evaluate cloud solutions first before considering any other option. Organisations are free to use other options but need to demonstrate that alternatives offer the right levels of security, flexibility and value for money. Since adoption in 2013, many government technology strategies have been directly built on a Cloud First approach, including the Ministry of Justice, the Department of International Development and the Department of Health and Social Care. The latter is particularly significant, as it mandates the NHS to follow the Government's lead whilst allowing NHS organisations to adopt compliant technology that suits local needs and reflects the nuances of local healthcare delivery.

In 2012, the UK Government developed **G-Cloud**, an initiative targeted at easing procurement by public-sector bodies of commodity information technology services that use cloud computing. Housing associations, although independent and not public sector, can procure through G-Cloud. For more detail on G-Cloud, see Chapter 6.

1.3 What is cloud computing?

Cloud service models

Cloud providers – directly or via technology partners - offer three cloud service models which can be deployed individually or together.

IaaS (infrastructure as a service)

The cloud provider enables businesses to provision processing, storage, networks and other fundamental computing tasks - within a cloud environment owned, managed and operated by the cloud provider. The business retains control over its operating systems and deployed applications.

The cloud environment is located in cloud provider data centres using virtual machines – software that runs an operating system or application environment, just as physical hardware would.

IT staff no longer need to manage and maintain the underlying physical infrastructure (usually hardware as operating systems are still managed by the client under IaaS) which allows them to focus on the management and deployment of their applications.

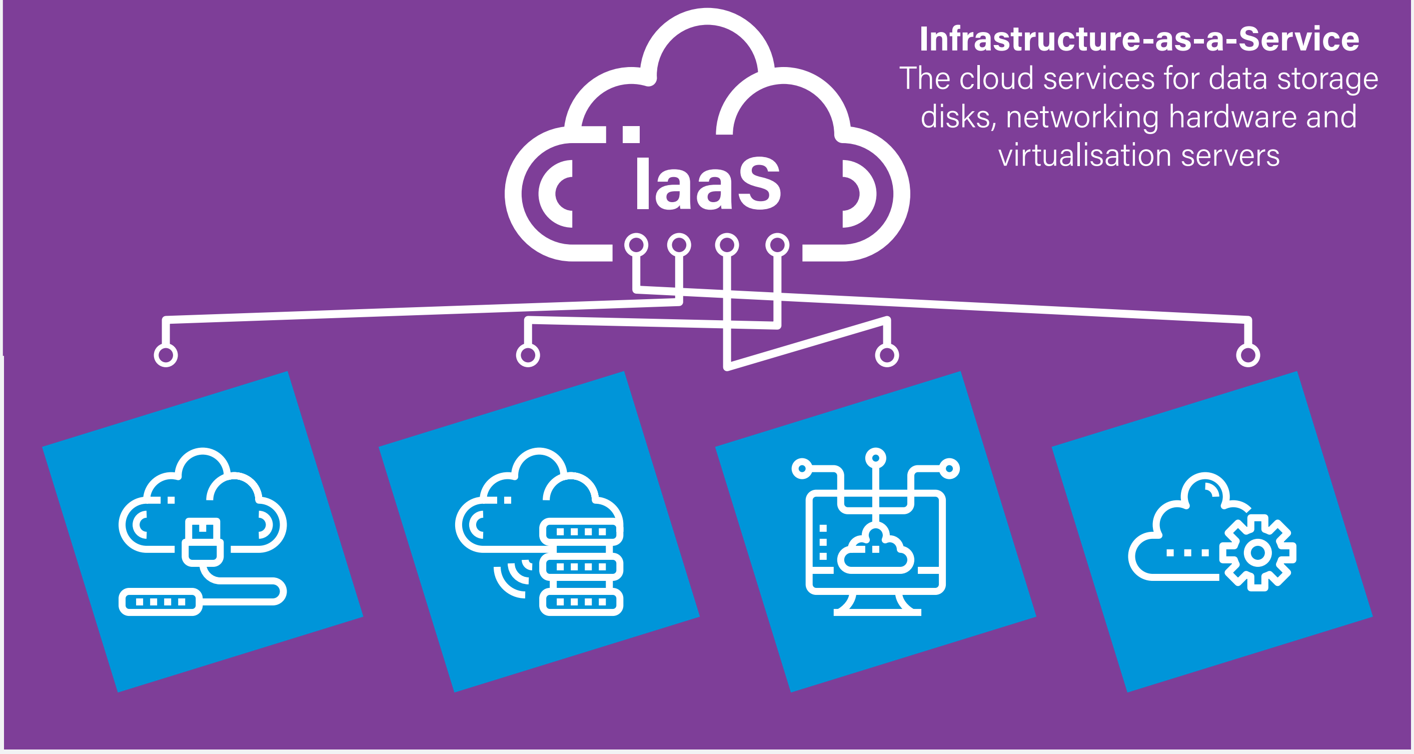

Infrastructure-as-a-Service
The cloud services for data storage disks, networking hardware and virtualisation servers

PaaS (platform as a service)

The cloud provider's offer goes further than IaaS by:

- including middleware (software that acts as a bridge between an operating system or database and applications, especially on a network), development tools, business intelligence services, database management systems etc
- enabling businesses to deploy and operate applications that they create themselves - provided that the business uses programming languages and tools supported by the cloud provider

As with IaaS, the cloud provider owns, manages and controls the underlying cloud infrastructure including network, servers and storage. Unlike IaaS, the cloud provider also manages and controls the operating systems.

PaaS is likely to be of interest to businesses that want to develop their own cloud-based applications in future.

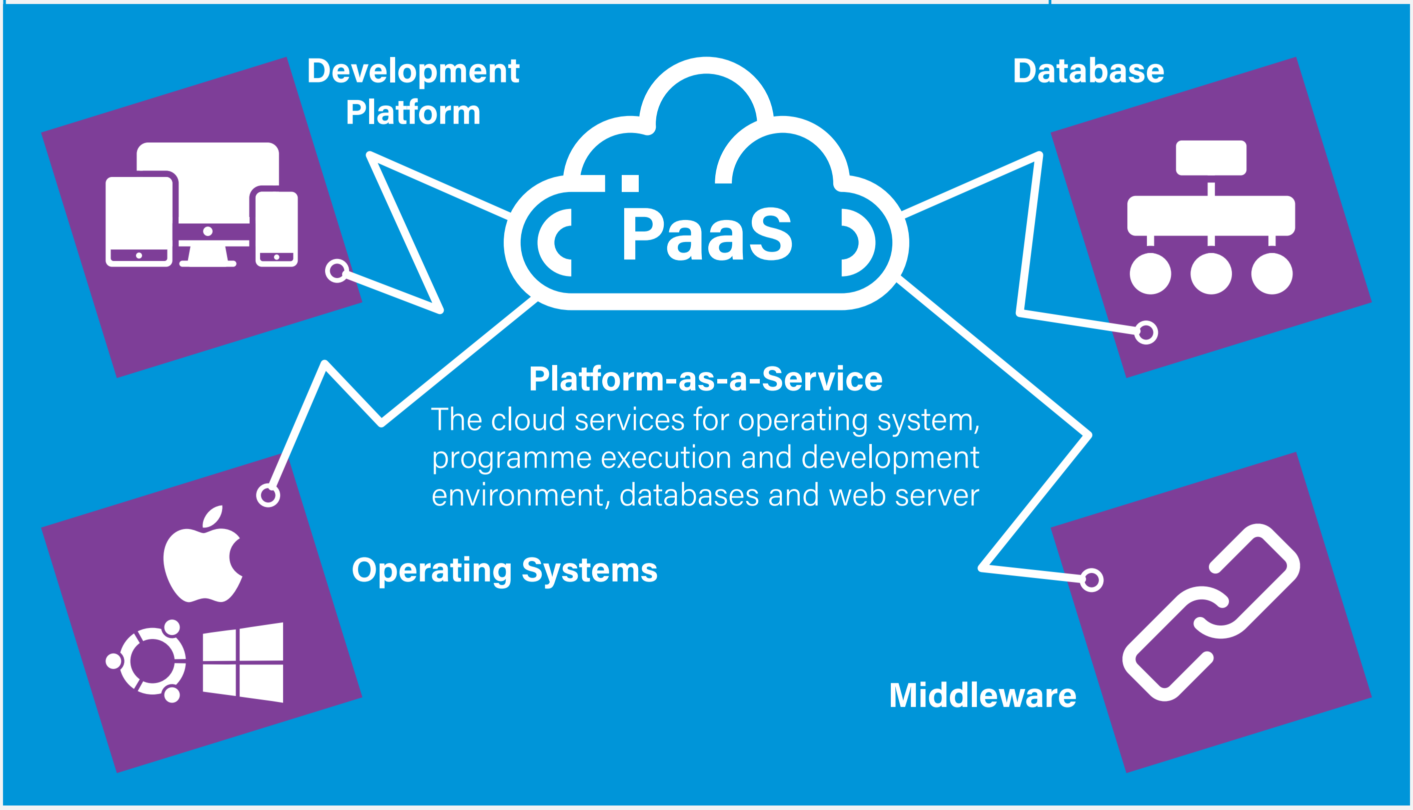

Platform-as-a-Service
The cloud services for operating system, programme execution and development environment, databases and web server

SaaS (software as a service)

Cloud providers, directly or via their technology partners, offer a complete software product (usually an end-user application) which they run and manage. As noted, applications offered in a specific cloud environment are known as 'cloud native.'

Where the business uses applications provided by the cloud provider or its technical partners, it no longer needs to manage software maintenance, patching or any other 'heavy lifting' involved in running software. The business only needs to think about how it will use those applications.

SaaS is already having an impact on the social housing sector and this is likely to increase where legacy housing and asset management system providers repackage their offering as a SaaS product.

Software-as-a-Service
Includes packaged software applications accessed by client business

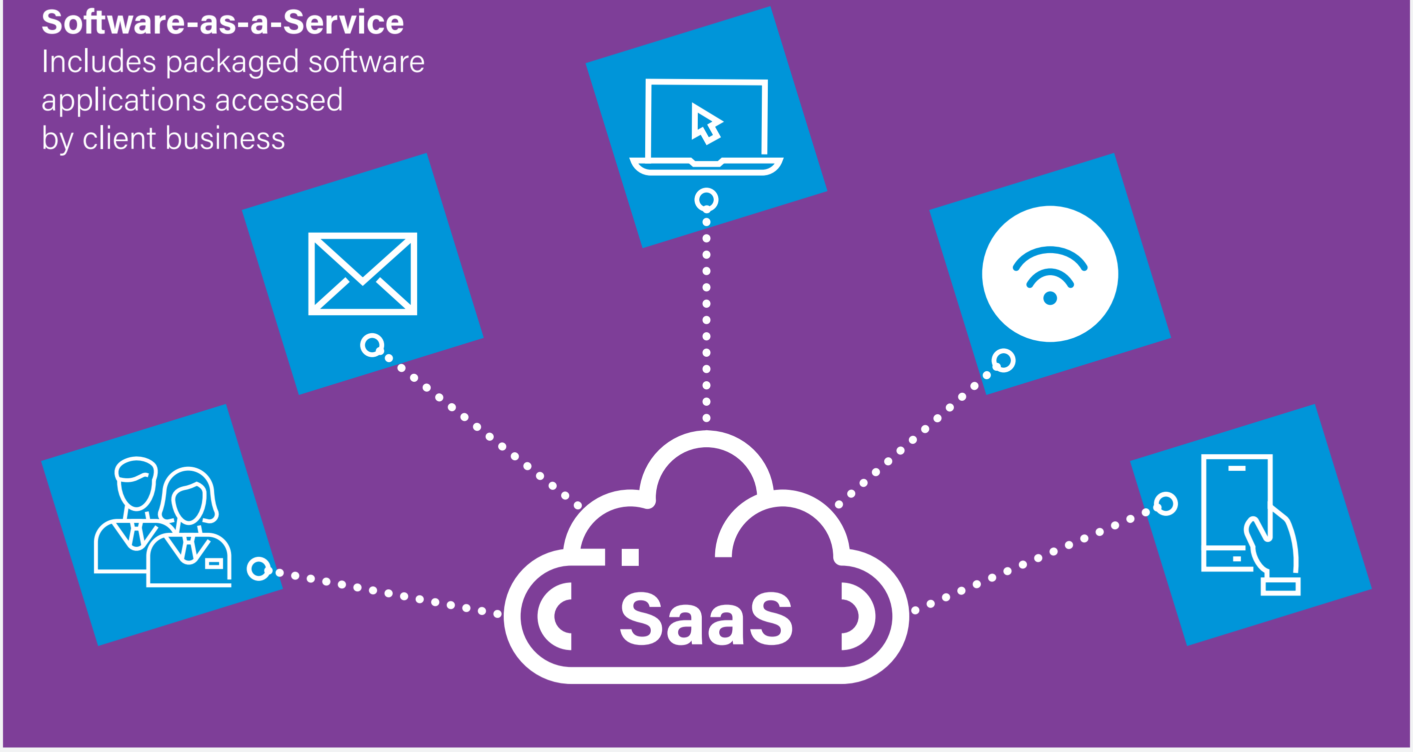

Distribution of management responsibilities

The distribution of management responsibility between the client business and the cloud and/or managed services provider varies by cloud service model and can be summarised as follows.

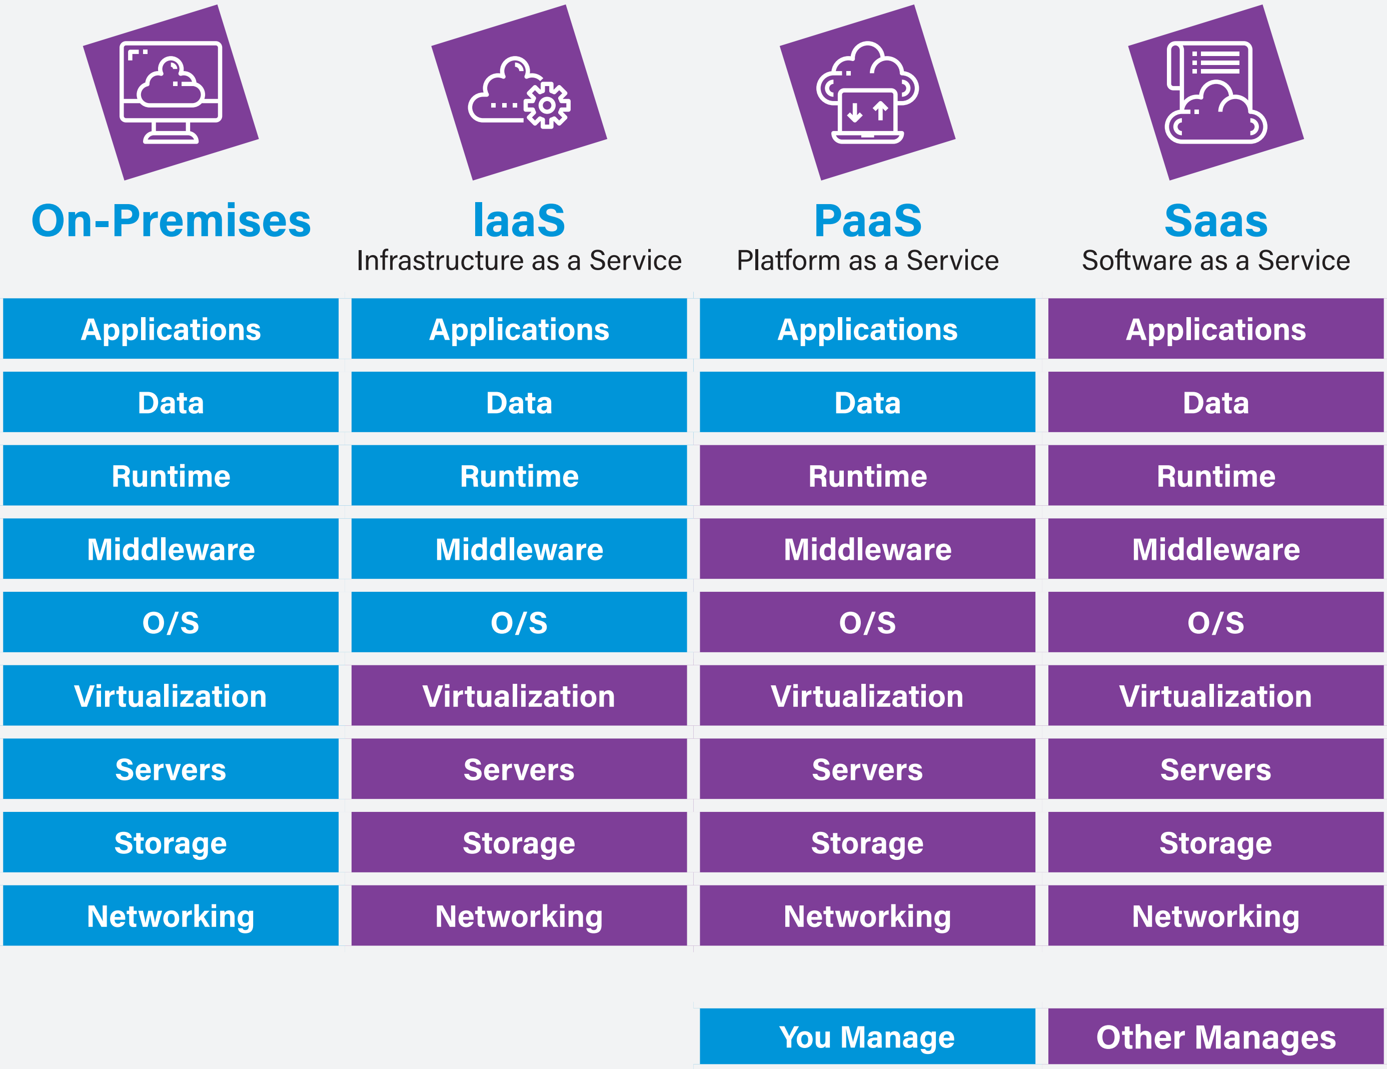

On-Premises	IaaS Infrastructure as a Service	PaaS Platform as a Service	Saas Software as a Service
🟦 Applications	🟦 Applications	🟦 Applications	🟪 Applications
🟦 Data	🟦 Data	🟦 Data	🟪 Data
🟦 Runtime	🟦 Runtime	🟪 Runtime	🟪 Runtime
🟦 Middleware	🟦 Middleware	🟪 Middleware	🟪 Middleware
🟦 O/S	🟦 O/S	🟪 O/S	🟪 O/S
🟦 Virtualization	🟪 Virtualization	🟪 Virtualization	🟪 Virtualization
🟦 Servers	🟪 Servers	🟪 Servers	🟪 Servers
🟦 Storage	🟪 Storage	🟪 Storage	🟪 Storage
🟦 Networking	🟪 Networking	🟪 Networking	🟪 Networking

🟦 You Manage

🟪 Other Manages

Cloud deployment models

Cloud providers offer different cloud deployment models.

Private cloud

Cloud infrastructure is provided for exclusive use by a single organisation. The private cloud may be owned and operated on premise by the client business or hosted by a managed service provider, such as Rackspace, within its own data centre.

Public cloud

The cloud infrastructure is provided for open use by the general public, including businesses, and is owned by the cloud provider and operates from its data centre. The cloud provider, for example AWS, distributes data storage over a number of data centres and client businesses share servers and storage capacity with security being underpinned by firewalls. Businesses access cloud services using a web browser and/or application programming interfaces (API's)

Hybrid cloud

The cloud infrastructure is a composed of two or more distinct cloud infrastructures (public or private) that remain unique entities but are bound together by technology that enables data and application portability between them.

Hybrid cloud can also mean the combination of public and private clouds with an existing on-premises non-cloud environment.

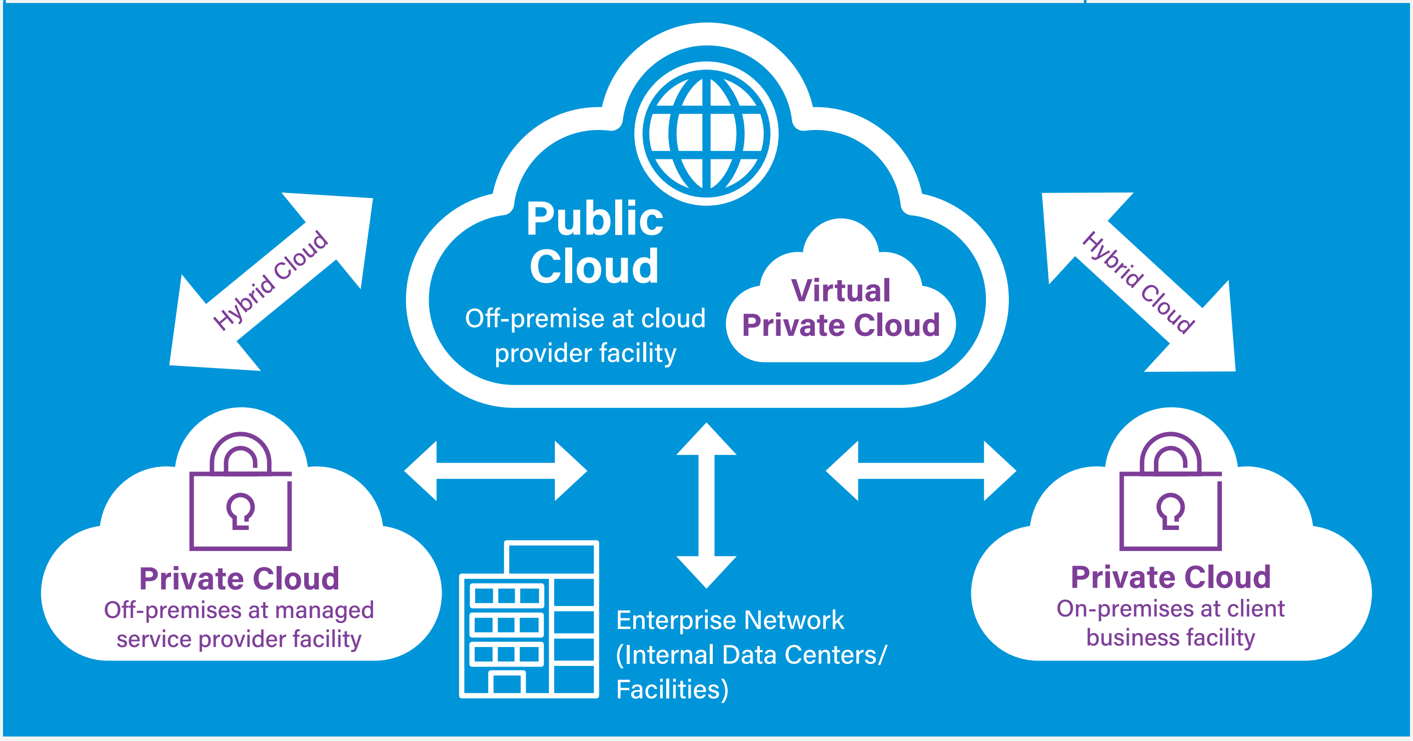
Public Cloud
Off-premise at cloud provider facility
Virtual Private Cloud
Hybrid Cloud
Hybrid Cloud
Private Cloud
Off-premises at managed service provider facility
Private Cloud
On-premises at client business facility
Enterprise Network (Internal Data Centers/ Facilities)

2. Cloud business case: meeting objectives

Any business will require a solid business case before adopting the cloud. The first step is to assess how cloud adoption can help it meet its key objectives – in a format akin to the following *generic* sector cloud benefit assessment *developed specifically for this guide*.

2.1 Social housing cloud benefit assessment

Cloud benefits	How delivered	Business impact	Business application
Reduced cost	» Rent not buy – no up-front capital costs » Pay as you go » Eliminate hardware and data centre management	» Better IT value for money » Resources available to spend on other activity	» Channel savings into surplus or additional activity » Invest in exploring and testing innovative cloud capabilities
Increased efficiency	» Minimise IT servicing roles » Automation of routine admin tasks » Increased capacity for value added work	» Leaner, more specialised and up-skilled IT function	» Redeploy IT staff into cloud management and developing new business solutions » Redeploy admin staff into new customer-facing roles » Reduce overall staffing levels if required
Improved security, resilience and continuity	» Shared responsibility with cloud provider » Cloud data centres staffed by experienced specialist security analysts » Cloud providers have capacity and agility to keep pace with newer security technology and with the level of patch updates	» Greater business assurance » Reduced downtime and faster recovery at lower cost » Few businesses will be able to afford to maintain an equivalent team of cybersecurity professionals	» Outsourcing of security, resilience and continuity releases staff and budgets for other activity » Business continuity testing is quicker, easier and more reliable where Disaster Recovery is migrated to the cloud
Enhanced business flexibility	» Elasticity - matching supply to demand	» Flexibility can help meet peaks in demand » Businesses can easily test out new approaches and scale them up	» Better manage peak demand resulting from seasonal business triggers e.g: ▪ Annual rent statements ▪ Responsive maintenance in adverse weather
Greater agility and innovation	» Easy experimentation » Access to eco system of cloud application developers » Fast release of new applications and services » Future technical developments may only be available in the cloud	» Wider access to new technological developments » Greater agility in responding to changes in operating environment » Greater meaningful engagement with customers » Speed of change enhanced	» Experiment with AI, robotics, Big Data etc with minimum cost, procurement difficulty or business disruption » Fast release of new applications - e.g. for managing and reporting on health and safety » Easier to engage customers in service co-creation » Easier to develop 'in house' applications

2.2 Potential for reduced cost

Cloud adoption can reduce IT infrastructure and applications costs as follows

Infrastructure savings:

- businesses only pay for the IT resources that they use
- cost of managing on-premises servers, data centres etc is removed
- cost of hardware renewal is eliminated - along with the associated management costs and risks
- power costs fall (often significantly) via elimination of on-premises infrastructure
- infrastructure support costs fall as legacy systems are retired

Application savings:

- consolidation of existing applications – where duplicate applications exist as a result of housing association mergers etc
- retirement of barely-used applications identified in preparation for 'cloud readiness' – saving on the cost of licenses and dedicated servers
- elimination of cost of manual upgrades to applications

The extent of these savings will depend, in part, on how many workloads are migrated to the cloud. Savings will be reduced by the cost of aligning cloud-based applications and non-cloud legacy systems. Savings will also reduce if legacy core-business applications like housing and asset management systems cannot be moved to the cloud, as that means that only part of the infrastructure can be reduced, or the data centre can't be shut down. However, most legacy system providers now support virtualisation of their applications, so this problem is becoming less significant over time.

Once a business is in the cloud, financial risk mainly relates to the variable nature of operating costs.

There are several ways to keep cloud costs in check, including conducting better financial analytics and reporting and automating policies for governance. Businesses should use the tools provided by the cloud service provider or managed services provider, or both. Additionally, as we shall see in Chapters 3 and 4, good planning and governance is essential to control costs and optimise savings.

2.3 Potential for increased productivity

By eliminating the tasks of maintaining and managing infrastructure and updating applications, businesses can create a leaner IT function. The cloud also supports the enhanced automation of routine admin tasks across the whole of the business.

Staff previously involved in these areas can be redeployed to focus on delivering value-added work that safeguards the business or enhances its commercial opportunities.

The **extent of increased productivity** will depend on:

- how many workloads a business migrates to the cloud
- the willingness of business to adopt cloud practices across all areas, not just in IT
- whether the business wants to 'cash in' on role or task elimination by making actual redundancies or whether it wants to minimise real redundancies via re-skilling and redeployment
- the effectiveness of the reskilling programme

2.4 Potential for enhanced data security and business continuity

Resilience is defined as maintaining operations in the light of unexpected events. Few if any businesses are willing to pay for 100% resilience (even if this were technically available). The key focus is on how quickly systems can be restored to operational capability.

Risk is managed through the cloud provider hosting and managing data over multiple availability spaces in a region – often, but not always, a country. Resilience is impressively high in the cloud.

The cloud can be used as a disaster recovery capability. Backup and service continuity is maintained whilst removing the need for redundant infrastructure on-site.

2.5 Potential for improved business flexibility

Cloud provision is based on actual demand, allowing organisations to scale up/down the services they need (automatically if required) for business peaks/troughs and only pay for what they consume.

This is known as 'cloud elasticity.'

For housing providers, this flexibility can help meet peaks in demand (e.g. annual rent statements and responsive maintenance response to adverse weather conditions) or new regulatory requirements (e.g. health and safety). Businesses can test out new approaches and scale them up with no restrictions other than their ability to pay for the work.

2.6 Potential for greater agility and innovation

Cloud computing can help businesses accelerate the release cycle for new services. This is important because CEOs and their executive teams are often frustrated by the pace of digital transformation.

Businesses can discover and prototype new capabilities (including IoT, AI/machine learning, local app development) without a large up-front investment or lengthy procurement processes. They can do so via access to tools and algorithms developed externally in the 'ecosystem' of cloud provider technology partners. For example, AWS provide pre-trained (out-of-the box) AI services with on-demand pricing that can be integrated into applications and workflows. One such example would be Amazon Lex, which delivers the capability to build conversational chatbots with Amazon's Alexa technology.

Businesses can, subject to inter-dependencies with non-cloud legacy systems, easily 'spin up' and 'close down' development testing environments and test and launch cloud-native applications within hours under this model.

If businesses want to develop new applications, in house and/or with technical developer support, they can test and launch these within weeks.

Resources, which otherwise would have been spent on buying, installing and configuring hardware and software can be deployed to build a proof-of-concept.

In future, as the cloud-native technology market increases, an increasing number of new technology solutions will be developed first or only in the cloud.

In developing the business case, organisations should seek opportunities for where cloud:

- lowers IT costs and increases IT agility
- can improve business agility and speed of release of business-critical service innovations
- can provide better resilience
- (as a process transformation tool) improves integration and collaboration across the organisation

For housing providers, one of the most exciting benefits of cloud migration is how it can aid 'service co-creation' or 'user-centric service design.' The ability to test new approaches to service delivery, modify them quickly following cloud-enabled customer consultation and deploy them within weeks means that housing providers can respond more quickly to emerging problems before customer patience runs out.

The extent of these business efficiencies will depend on the:

- readiness of the business to adopt the cloud and the skills to do so
- active scanning of emerging cloud-based technologies and new SaaS applications
- effective collaboration between the business and its IT function to procure or develop new cloud-based business solutions

3. Cloud business case: testing assumptions

If a business is excited or simply intrigued by the benefits that cloud computing can bring it will want to test key assumptions in the business case.

Some of these tests will be about the 'current state' on premises IT environment, how well it works and how much it really costs to operate. Other tests will be about how likely the business is to achieve the business flexibility and efficiencies identified in Chapter 2.

This process is often referred to as the 'discovery phase' and involves four introductory steps:

- accommodating the changes in IT accounting treatment and business processes required to invest in cloud computing
- reviewing current IT asset usage ahead of cloud adoption
- comparing the cost and performance of current on premises IT with that of cloud-based computing
- deciding what IT infrastructure and applications to prioritise for cloud migration. This is sometimes referred to as a migration readiness assessment

These actions will also help the business procure cloud services and undertake the migration of hardware, systems and data to the cloud.

3.1 Accommodating a different accounting treatment

Investment in cloud computing will involve a significant change in IT accounting treatment – from CAPEX to OPEX.

Cloud procurement largely operates via an OPEX model. When procuring IaaS, PaaS or SaaS there are no up-front capital costs. Instead, costs are billed as revenue on a monthly or annual pay-as-you-go basis.

Cloud services are also often priced by cloud and managed services providers as a percentage of IT infrastructure spend and flexed to reflect the level of support (low touch or high touch) procured by the client.

Key accounting differences

Options	**Capex** Capital Expenditure	**Opex** Operating Expenditure
Definition	Funds spent on acquiring, upgrading or maintaining fixed assets like land, buildings and equipment	Expenses incurred through ordinary business activities such as sales, rent, inventory costs, marketing, etc.
When paid	Lump sum payment up front (or financed with extra costs)	Recurring monthly or annual payment
Accounting	Accounted for over a 3 to 5 year lifespan as the asset depreciates	In the current month or year
Tax treatment	Deducted over time as the asset depreciates	Deducted in the current tax year
Listed as	Depreciation, equipment or property	Operating cost/expense
Example	Investing in a new data centre	Software as a Service Infrastructure as a Service

Key differences for IT procurement and its operation

CAPEX is not a cost-efficient accounting approach for investing in IT facilities. It leads to the purchase of server, data storage and disaster recovery capacity that isn't needed today in order to safeguard against tomorrow's uncertainties. It encourages housing providers to enter long-term IT approaches and vendor contracts that create business dependencies which are hard to break, limit business agility and maintain reliance on IT that becomes outdated or outgrown before it has paid for itself. CAPEX decisions also involve high-risk, often internally politicised, discussions and can provide organisations with paralysis when deciding to commit significant sums needed. Worst of all, CAPEX lets equipment dictate business approach, rather than business needs driving IT infrastructure investment.

The OPEX model ***can*** be more cost efficient than CAPEX, as it allows businesses to:

- scale computing and storage resources and application development as business needs fluctuate (no underutilised hardware, overprovisioning or emergency hardware purchases)
- more quickly and easily reconfigure existing and future IT architecture and applications – or develop new systems - to reflect changes in the operating environment

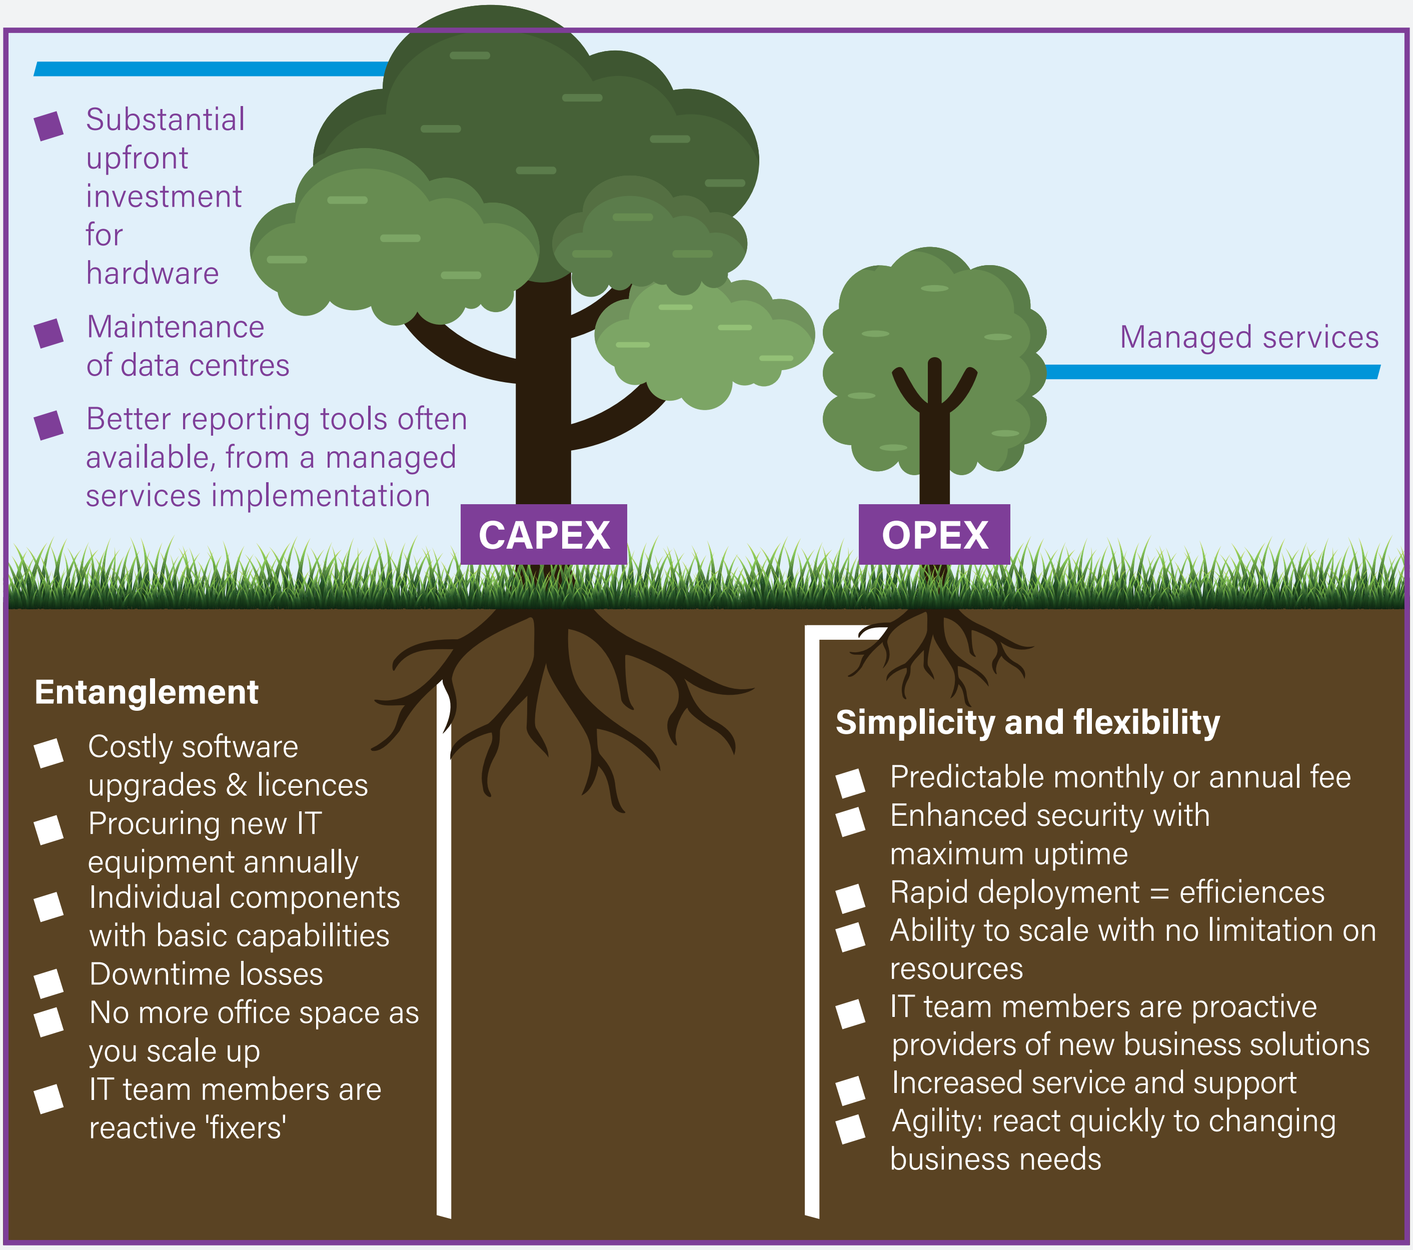

- Substantial upfront investment for hardware
- Maintenance of data centres
- Better reporting tools often available, from a managed services implementation

Managed services

Entanglement

- Costly software upgrades & licences
- Procuring new IT equipment annually
- Individual components with basic capabilities
- Downtime losses
- No more office space as you scale up
- IT team members are reactive 'fixers'

Simplicity and flexibility

- Predictable monthly or annual fee
- Enhanced security with maximum uptime
- Rapid deployment = efficiences
- Ability to scale with no limitation on resources
- IT team members are proactive providers of new business solutions
- Increased service and support
- Agility: react quickly to changing business needs

3.2 Reviewing current IT usage ahead of cloud adoption

Successful transformation to cloud computing must be undertaken in a structured manner.

For example, it is important to fully understand current IT usage by the business to ensure cloud solutions are not adopted in an ad hoc, tactical manner, resulting in a more complex and fragmented IT landscape, and increased costs.

As a critical first step, the business needs to gain key insights into its 'current state' IT environment:

- CAPEX depreciation, amortisation and renewal cycles for hardware and software
- current system usage, data storage requirements and central processing unit (CPU) usage compared to planned provision levels
- application interfaces and inter-dependencies
- IOPs (input/output operations per second)
- data storage
- current running costs – including the 'hidden costs' of current 'on-premises' provision such as the requirement to continually upgrade and patch systems, or tracing unused licenses or applications
- applications that are rarely used and can be retired

3.3 Comparing the cost of current on-premises IT with that of cloud-based computing

The business case should calculate the total cost of migrating to the cloud - the cost of moving systems over and the cost of running them in the cloud after migration - and then compare the total figure to the cost of keeping systems in-house.

This will require a keen understanding of both the existing infrastructure and associated costs (the 'current state' analysis) and the likely future costs following cloud adoption (the 'cloud future' analysis).

The 'current state' analysis involves understanding total usage of resources to operate current infrastructure and applications.

Current infrastructure costs should include:

- the run cost of in-house server rooms or data centres - including management costs, cooling, backup, power utilisation, floorspace rental, disaster recovery facilities and so on
- compliance and certification, plus end-of-life and decommissioning costs
- actual server utilisation and under-utilisation
- cost of connectivity, such as leased lines
- the number of physical servers, virtual servers, and details of specifications like memory and processor configurations
- cost of storage devices used for database and file stores
- licenses for operating systems (e.g. Windows), network and IT security products and IT support management tools
- potential upcoming hardware refresh plans that could be avoided
- depreciation of physical and amortisation of intangible IT assets
- in-house and third-party support costs (and contract termination costs where appropriate)

Current application costs should include:

- licensing
- in-house support and administration
- third party support
- updates, patching and bug fixes
- application support tools
- integration dependency

The 'cloud future' analysis involves estimating total usage of resources to operate cloud-enabled infrastructure and applications, with focus on:

- forecasting total use over some time period, typically a month. In terms of an application, how many hours will it be running per month? How much storage will it use? How much network traffic will be associated with it?
- changing rental rates based on usage tiers. A gigabyte of storage that costs a certain rate if total storage used is 10 gigabytes may be less expensive if total storage used is 10 terabytes
- variable usage patterns — low use for some periods and very high use for others

Cloud infrastructure costs should include:

- networking - bandwidth will vary dependent on how the organisation accesses the cloud
- data clean up (though the business should be doing this anyway)
- cost of migrating data to the cloud
- data egress and ingress transit costs
- administration tools for migrating current apps into a cloud environment
- data security – testing during and after migration
- staff - retraining of personnel
- relationship management – cloud providers and managed services or technical development partners

Cloud application costs should include:

- networking - bandwidth to enable access to hosted applications
- administration apps - for managing application performance and for managing integration with other apps
- testing data migration and security
- cost of migrating applications to the cloud – including licence transferability and upcoming licence renewals
- application design and development (where required)
- licenses for new cloud-based applications

Other additional cloud-related costs should include:

- management of cloud providers and their partners
- admin costs of managing monthly invoices from suppliers (if a public or hybrid cloud option is selected)
- exit charges for changing cloud service provider (should it prove necessary)
- further data egress and ingress transit costs
- write-off of existing ICT that will no longer be required or that will need to be run in parallel while the migration takes place

The cost of migration needs to be realistically costed and factored into the 'cloud future' calculations.

Migration is likely to be a phased transformation programme requiring both internal and external resources (delivery partner and possibly additional cloud specialists). Costs for delivery, testing and on-boarding services into production and training will need to be accounted for.

After the cloud goes live, running costs will depend on several business decisions:

- the choice of cloud provider or managed services/technical partner and their respective charging models
- the configuration of the cloud environment - private cloud, public cloud, hybrid cloud and multi-cloud
- the number of legacy systems that migrate to the cloud and those that can be retired
- the extent of procurement of new cloud-based applications
- optimising the use of IT resources in a cloud environment
- eliminating out-of-budget spending on cloud applications by removing 'shadow IT' from the business

'Cloud future' costs can then be compared to the 'current state' environment.

When determining the value for money of cloud migration, value-added benefits should be considered, including:

- improved disaster recovery and business continuity in the cloud
- increased financial efficiency of IT provision in the cloud
- increased financial efficiency of business operations in the cloud
- increased staff and customer satisfaction arising from cloud enabled improvements in back-office operations and front-line services

The traditional forecasting period of five years was tied to the lifecycle or depreciation schedule for hardware. In today's IT world, cloud providers suggest that three years is about as far as the crystal ball can see, as there are too many variables and emerging technologies. The forecasting term is critical to any cost-benefit analysis of cloud adoption. A business case which maps the cost of cloud against a three-year depreciation cycle looks more financially compelling than where the client is depreciating over five, seven or eight years.

Nonetheless, estimation time frames should be:

- long enough to recover the initial cloud investment, to fully exploit its benefits and to exit from existing contractual relationships
- short enough to provide reliable and realistic estimates of future IT capacity requirements

Tools to help define the business case

It can sometimes be viewed as an expensive and lengthy exercise to conduct a full total cost of ownership analysis on your current infrastructure to assess the business case for Cloud migration. However, there are now many sophisticated tools available, some totally free of charge, which can help gather some of the key data required for the start of a business case. One such tool is TSO Logic which provides data–driven analysis of total cost of ownership and cost modelling for future state post migration. The tool ingests millions of data points from your current infrastructure and then creates a statistical model of compute patterns for all operating systems, identifying which are over-provisioned and highlighting the potential savings when compared to a cloud provider's platform.

4. Cloud business case: risks and mitigations

No business decision is risk-free, especially where it can involve a radical transformation of the whole business.

Summary of the risks involved and how to mitigate them

Issue	Issue Detail	Risk	Mitigation
Overcoming cultural resistance to change	Impact on staff	► Role redundancy ► Flight of staff	► Cloud education ► Workforce management plan ► Upskilling IT staff ► Redeployment
	Loss of control of IT	► Perception of risk deters cloud adoption	► Effective cloud governance ► Cloud management audit and control tools
	Business is not culturally 'cloud ready'	► Implementation delayed or aborted ► Cloud benefits not fully realised ► Cloud IT costs not under control	► Create staff capacity for change ► Procure a managed services provider to provide guidance and strategic support ► Staff engagement plan
Developing the skills to manage in a cloud environment	Business is not technically ready	► Implementation delayed or aborted ► Cloud benefits not fully realised	► Training and redeployment programme ► Upskilling IT staff ► Employ a managed services provider to provide technical support ► Outsource to managed services provider
Legacy system dependancy	Business retains legacy systems on premises Business accesses legacy system in provider's private cloud	► Cloud benefits not fully realised ► IT usage not optimised ► Cloud cost savings reduced ► System provider private cloud is a strategic and technical 'cul de sac'	► Procure technical support from a cloud or managed services provider ► Adopt 'lift and shift' approach ► Utilise 'Desktops as a Service' in the Cloud to accelerate deployment ► Seek a SaaS alternative where appropriate
Overcoming concerns regarding cloud 'lock in'	Cloud provider performance consistently falls short of expectations	► Difficulty in switching cloud provider ► Difficult to return to 'on premises' provision	► Design contracts to require cloud provider support with migration ► Use of 'containerisation' to facilitate switch ► Procure technical support from managed services provider ► Why would business want to return to 'on premises' provision?
Data security	Data security diminishes when transferred to cloud	► Apocryphal risk — data is far more secure in the cloud ► Client business fails to deliver on its security obligations	► Shared responsibility agreement ► Focus on understanding client responsibilities ► Adopt Cyber Essentials or NIST protocols ► Adopt 'real time' data security monitoring tools ► Managed services provider support

4.1 Overcoming cultural resistance to change

Successful cloud adoption — like any other digital transformation project — extends beyond technical decisions. All business functions need to change the way they govern, operate and consume cloud services.

It is thus important for a business to understand the cultural concerns and barriers which can cause digital transformation to fail and sometimes prevent it from taking place at all.

4.1.1 Overcoming concerns about cloud impact on roles and employment

Social housing providers typically spend around a third of their IT costs on staff. Typically, between 50% to 75% of IT staff are employed to maintain and service IT platforms, servers, data centres, software updates. The upskilling of these individuals is a key component in driving cloud adoption.

A study by the IDC[1] found that comprehensively trained organisations are:

- 80% faster to adopt cloud
- 2.7x more likely to realise that cloud can help jump-start innovation
- 4.7x more likely to agree cloud can improve IT staff productivity
- 3.8x more likely to meet cloud return on investment requirements
- 4.4x more likely to overcome operational/ performance concerns

Most important, the report suggests that concerns related to cloud adoption are far more likely to be overcome with the support of comprehensive training.

The rapid growth of 'cloud careers' is also visible in the job market where demand for staff with experience of the major cloud platforms continues to be high, even in COVID-19 times. A current search for 'Cloud Engineer' on the job-board Indeed.co.uk returns close to 5,000 listings across a wide range of business domains. This begs the question; how do we retain our staff once they are fully trained? It also presents a risk of staff leaving the organisation if there is no investment in training.

When establishing a cloud strategy, people remain the key ingredient in driving transformation and their skills will determine the pace at which that adoption happens.

At this point, organisations will be faced with key decisions around whether to invest in their own staff or outsource to one of the many managed services providers to accelerate cloud transformation.

If an outsourcing strategy is selected, there is the potential for role redundancies, especially in the areas of core infrastructure management, where duties such as server patching, updates and monitoring could be perceived as manual tasks and not adding value to the overall mission of the organisation. Moreover, these tasks can be automated and performed at lower cost with cloud technologies and the support of a managed services provider.

Businesses will need to plan and deliver a reskilling programme to maximise the numbers of staff that can be redeployed into roles managing the cloud environment or in new 'value added' roles elsewhere in the business.

[1] To request this report go to https://pages.awscloud.com/tc-enterprise-idc-whitepaper.html

The cost of hiring ready-made cloud management specialists is likely to be prohibitive — so businesses may wish to retrain existing IT staff. Moreover, retaining their knowledge of the existing platform and systems will reduce risk in cloud migration. Skills transfer to redeployed staff is a key aspect of the offer of managed service providers.

Businesses will need to decide whether to 'cash in' these 'role redundancies' by making actual redundancies where staff do not have the aptitude for redeployment.

Mitigation of people concerns

All of this means a major role for HR in developing a cloud adoption ***workforce management plan****:*

- identifying required changes to processes and roles
- changing job descriptions as required
- identifying roles at risk of redundancy (if applicable)
- identifying skills gaps and new competencies required
- identifying staff with the aptitude for retraining and redeployment
- organising training programmes

It is both logical and necessary to undertake formal structured engagement with staff *and their trade unions* early on, and throughout cloud decision-making and implementation, to manage the workforce implications of cloud-driven digital transformation.[1]

Keeping employees informed throughout the project will make them feel more positive about cloud transformation and create opportunities for business units to contribute to cloud readiness, strategy and governance. A skills-retraining and redeployment programme is vital.

4.1.2 Overcoming concerns about loss of control over organisational IT

Loss of control is a primary concern of many organisations considering the cloud.

Reliance upon a third-party organisation — the cloud platform provider — for data security, business continuity and disaster recovery can cause frisson amongst executives, Boards and IT managers.

With PaaS and SaaS, the IT team does not always have full control over the provisioning, de-provisioning, and operations of *infrastructure*.

The IT team will also cede varying levels of control to the cloud or managed services provider over the management of *applications* — depending upon the PaaS or SaaS strategy the business adopts. **See page 9 for a diagrammatic representation of this topic.**

Consequently the business may be anxious about a 'loss of control' over governance, compliance, risks and data quality management.

Mitigating concerns about loss of control over organisational IT

The business is not 'losing control' when it shares responsibility for operational IT with a could or managed services provider. It is sharing responsibility with companies that can undertake functions more cheaply, more quickly and with greater expertise.

For most organisations this is a benefit, freeing up IT staff to focus more on the needs of the business rather than maintaining the underlying technology.

Moreover, cloud and managed services providers will provide the client IT team with a suite of management auditing and control tools which enable the business to monitor the delivery of outsourced (and on-premises) functions.

Cloud auditing and control tools include:

- monitoring of system health and performance with alerts
- visibility of configuration and resource utilisation
- user access control, compliance and policy management
- inventory management security rules, configuration, controls and alarms
- service catalogue creation
- license control and management

[1] A two year commission ***organised by the trade union community and the Fabian Society*** will identify the immediate actions that government, employers and trade unions need to take to support workers as technology impacts on jobs over the next 10 years. The commission's report is due for publication in Autumn 2020.

4.1.3 Ensuring 'cloud readiness'

Cloud migration reflects both the technical transfer of IT assets and data to the cloud and the reality that all business units need to change the way they govern, operate and consume cloud services if cloud adoption is to be successful.

'Cloud readiness' is the common term where a business successfully develops the necessary technical, security, business, commercial and operational capabilities.

Where a business fails to achieve cloud readiness, *negative outcomes* are likely:

- systems and data migration activities are under-costed
- implementation is delayed, repeated or aborted causing cost overruns
- use of cloud is not optimised causing more operational cost
- variable cloud IT costs are not controlled
- business improvements available via the cloud are not realised

For housing providers, cloud readiness requires a plan for aligning the, often inter-dependent, housing management, asset and finance systems that deliver core business operations.

To avoid unnecessary business disruption, it is advisable to adopt a phased approach:

- start teams off with low-risk projects
- then move to medium-risk projects, migrating non-customer facing apps to the cloud; and finally
- undertake high-risk projects, such as moving over mission-critical systems

AWS supports customer transformations with what is termed a 'Cloud Adoption Framework (CAF)'.[2] It helps guide a cloud journey by focusing on six perspectives: Business, Governance and People for business capabilities, and Platform, Security and Operations for technical capabilities.

Mitigation of concerns regarding cloud readiness

Ensuring that staff have the capacity to execute change is basic good business practice in any transformation programme and is not unique to the cloud.

By procuring a managed services partner, the business will be far better placed to manage cloud migration and minimise business disruption. These organisations have vast experience of this work and can help ensure the success of cloud adoption. External support for the change process is generally flexed to reflect the level of 'in house' capacity and confidence.

A staff engagement plan is the best way of overcoming concerns about achieving cloud readiness. The plan might involve:

- demonstrating the benefits and risks (and mitigations) of cloud migration for each tier or business unit in the organisation
- 'visioning' exercises such as 'how will the business operate after cloud migration?'
- learning from others via visits to cloud and managed service providers and 'early cloud adopters'

[2] https://aws.amazon.com/professional-services/CAF/

Figure 1 — The 4 stages of Cloud Adoption

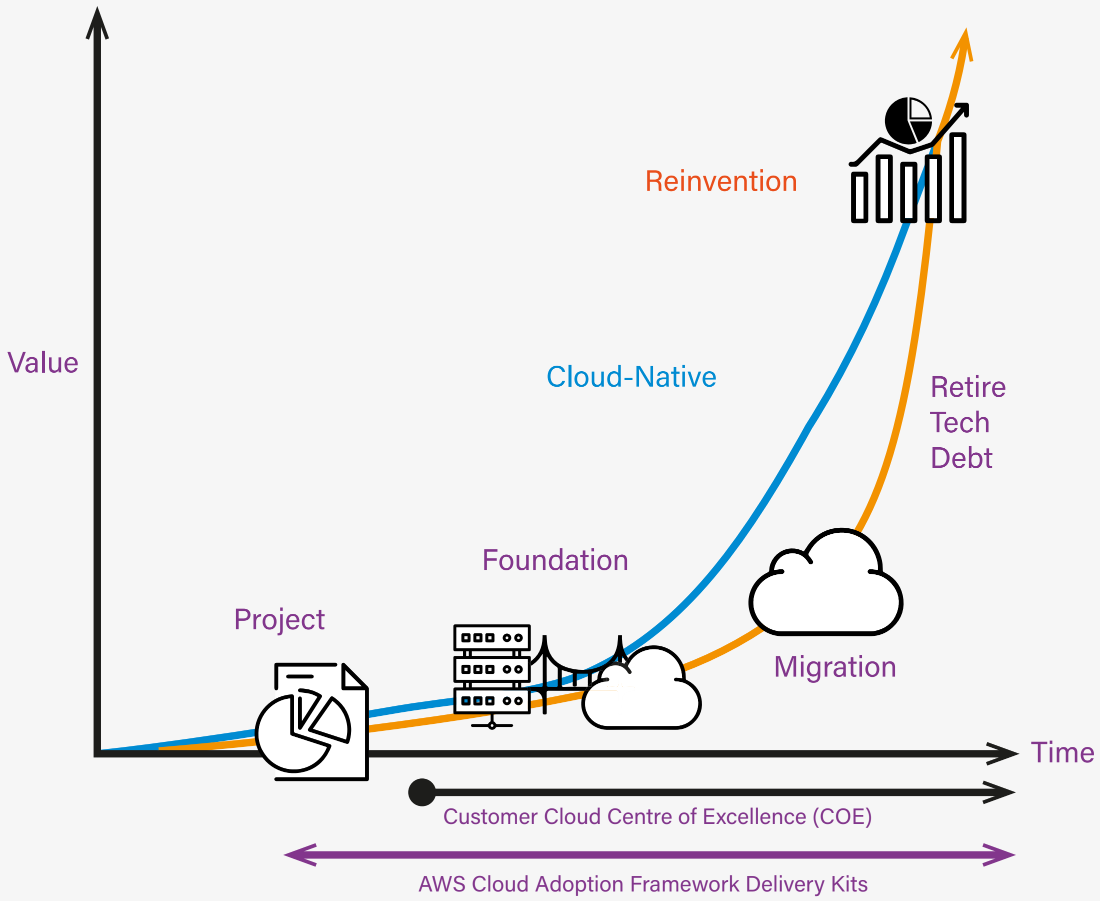
Reinvention
Value
Cloud-Native
Retire
Tech
Debt
Foundation
Project
Migration
Time
Customer Cloud Centre of Excellence (COE)
AWS Cloud Adoption Framework Delivery Kits

4.2 Developing the skills to manage in a cloud environment

The onus is on the whole business (and especially IT) to refocus its priorities and resources and adopt the skills needed to operate effectively in the cloud:

- business relationship management
- agile application delivery
- customer co-design of services
- cloud environment automation, monitoring and governance
- shift to a DevOps environment (*see Appendix A.1*) to deliver new business solutions

Unless the business acquires cloud-based technical skills at an early stage of the migration programme, and supplements these incrementally as more systems and applications migrate, cloud migration will fail to deliver the expected benefits.

Mitigation of risk of lack of cloud management skills

The **retraining and redeployment programme** should focus on:

- business and financial skills — to keep the change programme aligned to business objectives and on budget
- project management skills — to ensure that a multi-year change programme remains on track
- security and compliance skills — to ensure that cloud migration enhances rather than undermines existing protections and is sufficiently agile to respond to future regulatory requirements
- technical skills — to develop the capability to optimise cloud benefits for the organisation, such as a shift to a DevOps environment (see technical appendix) and ensure agile application delivery
- negotiation and contractual skills – the business relationship management skills to be an effective client when dealing with cloud providers and their partners, able to secure redress if contractual obligations are not being delivered

Businesses may find adding cloud specialists to their IT teams to be prohibitively costly, so it is critical to ensure that IT teams work closely with the cloud delivery partner to upskill. These IT teams should see this as a great opportunity to develop up-to-date skills that are in demand.

Collaboration with a managed services provider should be based on the 'knowledge transfer' of learning from pilot and subsequent migration projects to the 'in house' team. In this way, the organisation can develop the skills needed to optimise its use of the cloud.

One model is to procure onsite specialists from cloud providers and managed services and technical partners to work on-site in the business and to phase out these roles as in-house technical expertise is developed.

An alternative model is for businesses to share the cost of employing independent cloud specialists.

4.3 Legacy system dependency

The business of social landlords is underpinned by core business systems — housing management, asset management, rent accounting, finance and HR. User access to these legacy systems normally requires on premises deployment of software onto a desktop PC or device.

Retaining legacy applications, hosted on-premises, will mean that a business will:

- continue to require specialist skills for managing these applications
- fail to realise the benefits of the cloud in terms of economy, efficiency and flexibility
- lack the service delivery agility that staff and customers are demanding
- wait longer for application updates and improvements

Housing providers often utilise desktop virtualisation technology, such as Citrix, to make it easier to access legacy applications from any location.

Nonetheless, there is often a retained dependency on the housing provider to manage the underlying infrastructure on premises.

Some legacy systems have already migrated to the cloud. Suppliers such as Civica, Northgate and Oracle are now offering a hosted version of their application — as a SaaS product.

However, the legacy provider's private cloud may not be aligned to, or even compatible with, primary cloud environments (such as AWS, Microsoft or Google) — thus limiting future flexibility and innovation in these systems and reducing the economies of scale of cloud migration.

Moreover, the application version in the supplier's private cloud may be its latest — any business which has not adopted and system-tested every previous upgrade will need to do so before it is ready to engage with the cloud version.

Current housing sector dependency on legacy systems suggests that when a legacy supplier does migrate its application to their private cloud, clients may follow.

Businesses are therefore faced with either a hybrid of applications (cloud-based and not) or a multi-cloud environment (primary and legacy clouds).

Mitigation of risks of legacy system dependency

Two approaches can be adopted:

- migrating legacy systems to the cloud
- building and rebuilding applications in a public cloud environment such as Dynamics or Salesforce

Migrating legacy to the cloud

Where a legacy application is not cloud-enabled, one option is to work with a cloud or managed services provider to 'lift and shift' that application to the cloud. This *IaaS option* removes the problem of maintaining on premises infrastructure, eliminating the need for power and cooling and other physical maintenance dependencies, but retaining user responsibility for managing and provisioning virtual servers in the cloud.

As an alternative to Citrix and to help accelerate cloud deployment, businesses can take advantage of the Desktops as a Service (DaaS) facility that cloud providers offer on a pay-as-you-go basis. This virtual desktop service removes the need for on-premises servers whilst offering greater agility, speed of deployment and seamless integration with other cloud services. Amazon Workspaces is one such offering. A Desktop as a Service (DaaS) solution makes most sense if the provider is already managing a large on-premises virtual desktop environment.

Another option is to buy a packaged SaaS replacement system in the primary cloud environment that the business is aligned to. At the time of writing, few if any such applications are available.

However, organisations can move to SaaS based applications over time and as they emerge in the marketplace.

Building or rebuilding applications in a cloud environment

A business can build or rebuild an application from scratch in a cloud environment. However, this will require the recruitment of 'in house' developers and, as such, is unlikely to be affordable to all but the largest organisations.

A more cost-effective approach is build, or rebuild, a legacy application in a SaaS environment. A packaged SaaS solution should accommodate most of the client business's processes but will still require an element of bespoke configuration.

In social housing this commonly involves utilising Microsoft Dynamics or Salesforce as a development platform. For example, newer suppliers into the housing sector (such as Hitachi, TSG, RedKite Systems, Esusasive and HCL Power Objects) are helping clients to build new housing management systems on a Microsoft Dynamics platform.

However, with either approach, the business may not be able to retire its legacy housing management system completely because the new application will still be dependent on the old system for data inputs from highly complex subsidiary systems such as rents.

Choosing to refactor or rebuild either a legacy application, or a suite of inter-dependent services, is a major business decision. Cost is likely to be significant, as help from an experienced cloud technical services provider will be essential.

However, if the business can afford to do so (and address all data dependency issues) it will become far more agile and better equipped to deal with new or persistent service challenges.

Where there is a suite of independent services, this is a technology construct usually referred to as microservices

Microservices enable businesses to rebuild functionality in lightweight loosely coupled applications on that can be scaled horizontally on the cloud. Their use can be an effective strategy when legacy application architecture wont scale further, e.g. if the database grows too large or there are too many millions of lines of code or the business simply can't add features quickly enough and cost-effectively to meet its objectives. However, this approach involves major investment, especially if the process needs to be replicated across multiple core business applications. See Appendix A1 for more detail.

4.4 Cloud ‘lock-in’

A common fear around cloud migration is that the move will lock the business into a specific provider.

This risk is apocryphal. If cloud is built to enable migration, there are no long-term commitments and as there is an active supplier market there is little risk of ‘lock in.’

There are, however, active steps that businesses can take to reduce that risk still further:

> The most effective, and perhaps radical, way to avoid supplier lock-in is for a business to build its own housing and asset management systems supported by a technical development partner operating within a hyperscale cloud environment.

Active steps to mitigate the risk of ‘lock in’

- ensure that design decisions made at the outset of the cloud migration are supplier-agnostic
- assess possible contractual constraints, identify the exit arrangements and cost of exit (including data egress and ingress transit costs)
- ensure that all applications and associated data can be easily transferred, and data restored from backups
- ensure that the business — or its cloud or managed services provider — can unbundle cloud operational processes and procedures and map them to a different cloud environment if required
- take advantage of ‘containerisation’ which enables applications to be developed in a cloud-neutral manner but deployed in any cloud environment via APIs and configuration changes (see Appendix)
- be clear how the provider facilitates the movement of heavy amounts of data either through services or data migration tools

4.5 Data security

Assessing the risk of storing confidential or sensitive data in the cloud is essential.

The EU General Data Protection Regulation (GDPR) imposes extensive requirements for any organisation. GDPR compliance requires a strong level of tracking data and related consent.

Businesses are required to disclose any actual or potential breach to the Information Commissioners Office and to their customers and suppliers. They must report breaches in the security of their adopted cloud environment, *even if their organisation was not a target*. Housing providers must also report actual or potential breaches to the Regulator for Social Housing.

For any type of cloud environment, businesses will need to 'map over' existing GDPR controls.

The cloud model of distributing data storage over numerous data centres presents, in theory, a challenge to one of the key facets of GDPR compliance: knowing exactly where data is. In reality, for UK-based businesses, all hyperscale cloud providers provide 'sovereign' cloud solutions which means that UK data is held in a UK data centre.

Shared responsibility for data security

When a business runs and manages its own IT infrastructure on premises, within its own data centre, it is responsible for the security of that infrastructure, as well as the applications and data that run on it.

When the business moves to a **public cloud model,** it hands off some, but not all, of these IT security responsibilities to its cloud provider. Each party — the cloud provider and cloud user — is accountable for different aspects of security and must work together to ensure full coverage.

Cloud service providers and their managed service and technical development partners view data security as a shared responsibility — with both the client and the service provider making sure that their respective activities align with rather than undermine each other's endeavours.

The type of cloud service model — IaaS, PaaS, and SaaS — dictates who is responsible for which security tasks.

For example, in IaaS, the cloud provider is responsible for securing basic cloud infrastructure components, such as virtual machines, disks and networks. The provider is also responsible for the physical security of the data centres that house its infrastructure.

IaaS users, on the other hand, are generally responsible for the security of the operating system and software stack required to run their applications, as well as their data.

In a PaaS model, the cloud provider will look after the security of the operating system, software development platform and business intelligence services. The PaaS user will look after the security of the bespoke systems built on top of PaaS infrastructure. This is particularly important because many cyber attacks target poorly configured or unpatched servers

In a SaaS model, the provider is primarily responsible for the infrastructure and software stack, as the user has less control over these components.

Cloud solutions are much more secure than traditional data centres

Although 100% data security is impossible, cloud solutions are much more secure than traditional data centres.

First, they need to be. The profitability and reputation of cloud providers depends on ensuring that customers' data remains secure within the cloud. They refer to data security as 'job zero.'

Second, the hyperscale cloud providers have the resources to invest in the latest security technology and management techniques.

Deep and extensive technical security expertise is core to their business. They have mature access controls, vulnerability assessment practices, data protection and patch and configuration management controls. A typical cloud or managed service provider data centre is staffed 24x7x365 by experienced GCIA, GCIH and CISSP - certified specialist security analysts.

Few businesses will be able to afford to maintain an equivalent, robust, experienced team of cybersecurity professionals at their disposal to properly protect their on-premises data.

Concern may arise because different businesses share servers and storage capacity in cloud data centres.

However, cloud provider data centres have rigorous and stringent levels of physical security controls, perimeter protection access control and firewall separation. There is no evidence that shared business use of cloud architecture makes their data less secure. In fact, the controls within cloud data centres are significantly higher and better implemented than other data centres, because they are shared multi-tenant services.

It is essential that the client understands its responsibilities and how to execute them.

Data security should be the bedrock of all digital activity, including the cloud. Before committing to any path of action, businesses should have regard to guidance published by:

Cyber Essentials — the UK government information assurance scheme operated by the National Cyber Security Centre (NCSC) that encourages organisations to adopt good practice in information security. It includes an assurance framework and a simple set of security controls to protect information from threats coming from the internet.

NIST — a non-regulatory agency of the private sector organisations in the United States can assess and improve their ability to prevent, detect, and respond to cyber-attacks. The framework is often cited as global 'best practice'.

Patching remains the single most important thing a business can do to secure its technology. By using automated serverless components, cloud and managed services providers can take over the management of patching leaving the business with only its own code to manage. Serverless security patching, linked to cloud providers up-to-date knowledge of cyber threats, can produce greater security than even the most diligent local administrators.

Mitigating cloud data security risks

Ensure that the *cloud data governance plan* focuses on the following issues:

- identifying what service levels, security policies, governance and auditing are currently in place to protect the business
- recognising that the business will remain the data controller under GDPR and that the cloud or managed services provider will only be a data processor
- ensuring that the cloud provider provides a 'sovereign' cloud solution — thus ensuring that all data is stored in the UK (ensuring regional alignment to GDPR)
- limiting data access to only users and systems that need it
- ensuring that current client Data Protection Officer responsibilities extend to the cloud
- understanding of 'shared security' responsibilities in each cloud model and ensuring that the business meets its obligations
- utilising cloud provider alert systems to monitor for any exposure of cloud data to the public
- managing the 'attack surface' via file integrity monitoring and minimise risk with secure configuration and vulnerability management
- determining what security features an application has, such as data encryption, and ensuring that data are encrypted and protected when in transit across the network
- periodically testing the cloud governance plan — verifying the integrity of backups and ensuring that archived data is stored securely and that any unnecessary data is destroyed

[1] https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf

5. Developing a cloud business strategy

Where an organisation decides to adopt the cloud it then needs to develop a cloud business strategy as a component of its overall IT or digital transformation strategy

Specifically, a cloud business strategy will involve making decisions about:

- the preferred model for cloud governance
- which IT assets and business applications to migrate
- which cloud services to procure
- which cloud deployment model best suits the business

5.1 Developing cloud governance controls

Lack of basic cloud governance often leads to mismanagement and friction between the business and IT. For example, organisations often lack the necessary understanding of spending on cloud-based applications by business units.

Moreover, without business-wide controls over the procurement of cloud-based applications, businesses cannot contain server sprawl and properly control procurement and running costs.

Effective cloud governance requires an executive-sponsored, cross-functional team, empowered to drive the adoption of the cloud across the business through a common governance model.

The governance method needs to involve all business units — perhaps through representation on a cloud governance board but certainly via structured engagement. In order to benefit from agile cloud technology in delivering services to customers, intelligent businesses will seek structured engagement with customers at this strategic level. For example, decisions on which service applications to migrate to the cloud can only benefit from customer views on prioritisation.

The importance of the cloud governance function is highlighted by the extent of the activities that it must monitor:

- cost visibility and controls
- decisions on how the organisation will manage and operate cloud services
- business-wide policies that control the lifecycle of cloud investments and ensure a proportionate service application catalogue
- cross-functional organisational alignment — ensuring that cloud migration meets the needs of each business unit
- reporting to the executive team on implementation progress and alignment to budget
- controlling the influence of cloud service providers — as it is their technology, they can significantly influence the client organisation's change roadmap

Successful cloud governance is evidenced where:

- new and rigorous oversight is implemented — building on and enhancing existing IT governance
- the strategy for controlling procurement of cloud-based applications works in practice
- board, staff and involved customers are consulted and brought 'on board' with cloud adoption
- cloud migration delivers the expected benefits

An important aspect of cloud governance is the availability of independent advice

Cloud providers will want to sell their own product. Managed services providers are usually 'vendor-neutral'. However, where referrals are made to them by cloud providers there can be a natural expectation that the referring supplier will be recommended.

There is, therefore, value in undertaking 'pre-transformation' work to give a company the time to build a common understanding of what technology can achieve and its impact on the business.

The company will understand technology in a 'business-first' rather than a 'technology-first' context and the new business models that might be open to it by embracing this technology.

The adviser can assist the business in gaining awareness of tech alternatives, likely levels of business disruption that will follow their introduction and the changes to culture and process necessary to make the transformation successful.

This 'pre-transformation' work should be a short 'kick start' process and should finish before — and neither conflict nor overlap — with the role of the managed services provider.

The independent adviser helps the business to understand what it wants from the cloud, having explored the options.

With this clarity, the business can then engage with a managed services provider to put that vision into practice — procuring cloud infrastructure that best fits its needs and migrating current infrastructure and applications to it.

5.2 Deciding what to migrate to the cloud

Most infrastructure — servers, data centres, disaster recovery centres — will be replaced when a business migrates to the cloud under the IaaS model.

Full infrastructure migration will depend on whether the business adopts a PaaS or a SaaS model for applications and the extent of legacy system migration. Decisions about the pace of infrastructure or application migration are likely to be determined by where a business is at in terms of its depreciation, amortisation and renewal cycles.

Key issues to consider when deciding upon priorities for the migration of infrastructure or applications:

- changes in business objectives or operating environment
- impact on business users and their expectations
- impact on customer service delivery
- number/ type of workloads and ease of mapping them to an IaaS, PaaS or SaaS model
- integration dependency with other applications

- Retire — permanently
- Replace — buying an entirely new software-as-a-service (SaaS) package
- Re-host — 'lift and shift' existing applications in the cloud unchanged (IaaS)
- Re-platform — 'lift and shift' existing applications in the cloud with some changes (IaaS or PaaS)
- Rebuild — completely rebuilding them for the cloud (PaaS)

Key considerations for each approach are summarised as follows. For a more detailed explanation see Appendix A5.

Migration Strategy	Timescales	Complexity (cost, effort, skills)	Possible Pros	Possible Cons
Retain	Intermediate	Low	▶ Minimal business disruption	▶ Lost savings opportunity ▶ Application may no longer meet business needs
Retire	Short	Low	▶ Cost savings from systems shutdown ▶ Reduce IT support sprawl	▶ Data retention requirements may constrain retirement of some systems
Replace	Short	Low	▶ Accelerate cloud adoption buy-in	▶ It may be difficult to map data from the old app to the new one. Not necessarily bad dependent on the original data quality ▶ SaaS cloud-based alternative may not exist ▶ Access/security issues ▶ Lock-in to new supplier
Rehost	Short	Low	▶ Speed of response to business changes ▶ Quick return on investment ▶ Legacy supplier lock-in continues	▶ Unsuitable for complex workloads ▶ Cloud provider will need to build new hosting environment for legacy application ▶ May not be optimal solution long-term if application becomes out-dated and obsolete
Replatform	Medium	Medium	▶ Savings on infrastructure and licenses ▶ Reduced operational overhead ▶ Reduced lock-in to legacy supplier	▶ May cost more than re-host ▶ Requires complex and thorough integration testing
Rebuild	Long	High	▶ Innovate and re-invent business processes ▶ Greater longer-term efficiency ▶ Ability to build in cloud portability ▶ No vendor lock-in	▶ More expensive and complex short to medium term ▶ Complex migration planning and implementation

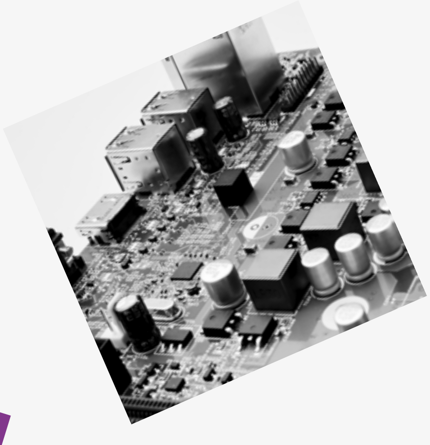

5.3 Deciding what type(s) of cloud service to buy

As noted in Chapter 1, cloud providers primarily offer three different service models:

- Infrastructure as a Service (IaaS)
- Platform as a Service (PaaS)
- Software as a Service (SaaS)

Businesses will need to assess the suitability of each model for the applications and processes they need to migrate or the new ones that they wish to deliver. A generic assessment follows.

Model	Description	Pros	Cons
IaaS	▸ Businesses get the basic infrastructure building blocks covering computing (physical/ virtual), storage, processing and network ▸ The business manages the operating system layer and above	▸ More control and flexibility ▸ Beneficial for quick wins and speed of migration to cloud	▸ More resource intensive to manage ▸ Fewer cost savings than other models
PaaS	▸ Businesses don't have to manage the underlying infrastructure so IT can focus on building business functionality	▸ Flexible and speedier service development ▸ Enables bespoke 'in house' application build ▸ Less IT resources required	▸ Restricted ability of business to modify infrastructure and its functionality ▸ Provider lock-in (unless built by client business)
SaaS	▸ Businesses get a complete product, run and managed by the service provider and they only need to think about how to use the software	▸ Quick deployment ▸ Fewer IT resources required	▸ Provider lock-in ▸ Limited flexibility for 'in house' bespoke application modification ▸ Forced upgrades at an inconvenient time to the business

When evaluating the options, one approach is to map out whether an application is a core business process (e.g. housing management) or a supporting business process (e.g. HR, payroll, finance) and whether that process can be considered as a:

- non-value adding process — does not add or enhance any benefit for the customer
- value-enabling process — does not directly add value to customer but must be performed to enable improved customer service
- value-creating process — core business process which directly impacts on customer service

The choice of cloud service model should reflect how critical an application is to the business. It will also reflect the appetite of the business for building its own applications or simply buying 'off the shelf.'

	Core process	Example	Support process	Example
Non-value adding process	► SaaS	► Audit and compliance	► SaaS	► HR (time booking)
Value-enabling process	► SaaS, PaaS	► CRM	► SaaS	► Repairs scheduling
Value creating process	► Paas, IaaS	► Housing management — rents ► Property management	► SaaS, PaaS	► Property sales

All these services can be procured via cloud providers or their technical development and managed services partners. SaaS solutions for a private cloud tend to be procured direct from the software vendor.

5.4 Deciding upon a cloud deployment model

Businesses will need to decide which cloud deployment model to adopt. The table below summarises their advantages and constraint.

Deployment Model	Advantages	Constraints
Public cloud	► No infrastructure investments ► Scale and flex IT resources to meet demand ► Reduce IT infrastructure management — IT staff can be more business focused ► Pricing based on usage consumption so no over-provisioning of IT ► Access to cloud-native developers and their service catalogue ► Easy to 'spin up' and 'close down' development testing environments	► Obligations under the 'shared responsibility' data security model need to be carefully managed ► Strong finance governance and management required to control costs
Private Cloud	► Dedicated and secure environment where the business controls: ▪ service levels and performance ▪ data protection and privacy ▪ compliance and regulatory requirements	► Usually (if not always) more expensive than public cloud due to retained maintenance of IT infrastructure, unless a full IaaS or PaaS model is adopted ► Limited business agility unless a full IaaS or PaaS model is adopted

Deployment Model	Advantages	Constraints
Hybrid cloud	► Flexible approach to cloud migration ► Balanced trade-off between using public cloud for capacity and cost drivers versus private cloud for security and/or regulatory/ compliance drivers	► Complex to manage two cloud environments and more expensive ► Running infrastructure and applications across public and private cloud may create hidden interdependencies between systems ► It may thus become more difficult to control and fix — exacerbating the disruption caused by a single point of failure ► Enhanced security considerations for traffic handling etc. ► These issues can be mitigated by using containers (see Appendix A1)
Multi cloud	► Select the best cloud services and best price on a vendor-neutral basis ► Spread service continuity and resilience across multiple providers	► Running a variety of cloud native applications will be more expensive and complex to manage ► As with hybrid, containers can make multi-cloud easier to manage ► More challenging to control costs — each provider may have different pricing models ► Can lead to organisations needing to use lower level services if all cloud providers don't have the same breadth and depth of capability

When deciding which applications to move to the public cloud, an organisation should start with ones that are less mission-critical and can be cost-effectively maintained at the cloud provider's data centre.

Some businesses take the view that private cloud is more appropriate for applications and data that — if rendered inoperable or stolen — could affect the organisation's ability to function.

A key question is under which cloud deployment model will business-critical applications be most resilient. The locality of on-premises solutions ostensibly provides greater peace of mind, but the cost of making private clouds resilient may push organisations to go to a public cloud provider with resilience embedded within its offering.

5.5 Developing a cloud migration plan

Housing providers will need to devise a structured cloud migration plan and then deliver on it. Typically, a business will use the services of a managed services provider to develop the detailed plan.

The migration plan should focus on:

- applications, security, data
- finance and its governance
- gap analysis, transformation and design — mapping apps and data codes to establish and tackle service dependencies
- future requirements regarding capacity and utilisation
- future requirements regarding staff skills, service management and support

An efficient migration pathway reduces the risk of data being lost during the process or accessed inappropriately.

Businesses will need a consistent toolset, processes and security posture operating across hybrid IT deployments to provision, view, access and manage public and private cloud resources with a single set of credentials.

Managed services partners use specific tools to guide a business in making transformation decisions for cloud migration. These discovery tools assess existing IT infrastructure to identify utilisation, server and application workloads, consumption, peaks and capacity.

From this insight, the tools provide guidance on the best solutions to adopt - based on factors such as performance, application dependencies, suitability and readiness as well as predicted costs compared to current costs. The tools also provide recommendations for optimal cloud configuration, including server compute instances, storage options, network settings and pricing plan.

In taking its decisions regarding cloud services and deployment models, the business will have determined its target IT/business operating model and its appetite for outsourcing via the cloud.

The business will be in a good position to determine how much of its IT infrastructure and operations that it wants to retain 'in house' and how much responsibility it wants to transfer to the cloud provider and managed services partner.

After cloud migration goes 'live', the business will need to deliver on its structured cloud implementation plan. This will involve:

- problem source identification and resolution
- managing risk in terms of business disruption, disaster recovery, data protection etc
- optimising cloud computing resource usage
- optimising the use of PaaS and SaaS in the cloud environment
- controlling costs
- further staff training and cultural realignment

Cloud migration checklist

PLAN AND ASSESS

Lay a solid foundation by gathering important information about your current environment and requirements

☐ Map your business services and applications, including dependencies

☐ Calculate your operational costs, including hardware, software and labour as well as hidden costs like facilities, recruitment and training

☐ Assemble security and compliance requirements

☐ Identify your critical-to-quality (CTQ) metrics, such as current downtime, your acceptable downtime, and your ability to observe and respond to downtime when it occurs

☐ Determine how much risk is acceptable

☐ Identify individuals from your team who will be involved in the migration (including a project manager familiar with technical project delivery)

Application Profiling

Use application profiling to gather and organise information about your applications, and identify a corresponding migration strategy

☐ Itemise and collect key information about your applications, such as application role, level of criticality, and level of infrastructure automation

☐ Where applicable, determine the workload categories you currently support (for example, your business may primarily run database-intensive, compute-intensive and user-interface-intensive workloads)

☐ Determine your baseline current-state environment

☐ Explore how end users run your applications

☐ Make sure you understand the critical-to-quality metrics

☐ Match the requirements for each workload category, with the right cloud platform and configuration

☐ Identify which workloads can be most-easily migrated, without the need to re-platform or refactor

DESIGN

Now that you have identified the workloads you wish to migrate, it's time to design the migration

☐ Align your platform-build and migration activities, and be sure you have access to both teams for creating operational schedules

☐ Build out your organisational strategy into a runbook, detailing the procedures and operations you will follow throughout the migration

☐ Be sure all of your stakeholders are aware of known risks and issues associated with the migration

☐ Assemble a migration team to execute the migration

☐ Develop a comprehensive project plan to complement the runbook

Migrate

At this point, you understand how your application is constructed and what its dependencies are, you've identified a migration strategy and you've exploded that out into a detailed migration runbook. Now, with runbook in hand — and a team of experts behind the wheel — you should be all set for a successful migration

☐ Perform a dress rehearsal of the migration runbook and remediate as needed

☐ Identify ideal migration window

☐ Communicate migration plan to the business

☐ Establish a contingency plan for all of the 'what-ifs'

☐ Ensure the migration testing and rollback criteria are agreed upon and captures

☐ Execute the migration based on the plan

☐ Conduct post-migration validation of applications, data and network accessibility

☐ Cut-over and go-live

6: Procuring cloud product and support

6.1 Procuring cloud computing services

An organisation can buy cloud computing product and services direct from a cloud provider, procure external support for systems integration and migration from a managed services provider — but then manage ongoing cloud delivery in house.

This is often referred to as the ‘DIY option.’ It is likely to suit those housing providers that have a high-level IT and project management skill base — capable of managing digital services, optimising their usage and alert to the business opportunities than can be unlocked by the cloud.

Alternatively, a housing provider can buy cloud computing product and services from a managed service provider. These providers typically sell on cloud storage, applications, software licenses, server stacks and so on. Their primary role is to manage migration delivery and systems integration.

Procurement of applications can benefit from the vast array of technical development partner offerings within each hyperscale cloud environment.

Generally, housing providers will go through a full OJEU procurement process (directly or via an established OJEU-compliant framework) when selecting a cloud or managed services provider. However, where a business is seeking merely to rent space in a development testing environment no OJEU process is required unless the value is over the OJEO procurement threshold.

In theory, when a business has a framework agreement with a cloud provider, it should be able to procure (for example) packaged cloud-native SaaS applications – via the UK government Digital Marketplace (see box) — without undertaking an OJEU process. For example, Microsoft365 or anti-virus tools are often purchased this way.

However, for more significant application procurement, businesses often voluntarily undertake OJEU procurement in order to maximise competition and choice and minimise the risk of challenge from alternative suppliers.

As noted in Chapter 1, the UK Government has developed G-Cloud, an initiative targeted at easing procurement by public-sector bodies of commodity information technology services that use cloud computing. Housing associations, although independent and not public sector, can procure through G-Cloud.

The G-Cloud consists of:

- a series of framework agreements with suppliers, from which public sector organisations can buy services without needing to run a full tender or competition procurement process
- an online store — the ‘Digital Marketplace’ — that allows public sector bodies to search for services that are covered by the G-Cloud frameworks

As noted, managed services providers typically assist the business in deciding which cloud deployment model and services to procure.

When managing cloud operations after ‘go live’: Businesses often commission support from cloud providers or (more commonly) from managed services providers. A contingency or specific budget should be reserved for this type of support. If internal ICT skills are not available, serious over-spend can result.

This support can involve assistance to:

- manage the public cloud platform on behalf of the business
- host a private cloud platform – in either the data centre of the cloud provider or managed services partner or on the client’s premises
- manage the alignment of a hybrid of private and public cloud operations and non-cloud systems

6.2 Selecting a cloud provider

When selecting a provider, organisations will need to determine 'best fit' with their future business objectives. Procurement award criteria may include:

- functional fit — e.g. how well the supplier's proposal meets the needs of the business in terms of coverage, network capacity, performance and innovation
- technical competence — e.g. the supplier's skills and experience, how it will identify risks and dependencies and manage them
- cultural fit — e.g. how the supplier will work in the business, how it will solve problems and share knowledge and experience with 'in house' staff
- non-functional characteristics — e.g. supplier terms, help with onboarding and offboarding, scalability, reliability and automatic disaster recovery, catalogue of partner technical services
- after-sales service management — helpdesk, account management, business continuity and data protection assurance
- whole life cost — the cost effectiveness, price and running costs of the service — and overall value for money
- on-going support and maintenance

Traditional data centre procurement requirements are no longer relevant. Recycling them might result in cloud providers being unwilling to bid and/or poorly designed contracts which fail to optimise benefits.

For example, it is no longer necessary to dictate customised specifications for equipment, operations, and procedures (e.g. racks, server types, and distances between data centres).

Instead, successful cloud procurement focuses on application-level, performance-based requirements that prioritise workloads and outcomes.

Commercial cloud industry standards and industry-recognised accreditations and certifications can help provide assurance for businesses procuring cloud services.

The cost of licenses can be the biggest single cost in an IT budget.

Software licensing can be a complex and often frustrating challenge for many IT departments, given the complexity and rules around licensing models, usage and compliance — particularly if the IT environment is split between outsourced services and internally delivered systems. This issue tends to be more complex for IaaS and PaaS models compared to SaaS, where the software license is part of the provision.

When it has migrated key processes to the cloud, the business may still need to continue with some existing licenses where it:

- 'lifts and shifts' a legacy application to the cloud
- it has procured an alternative application but needs to retain the legacy application as a system of record, especially where it is the repository of data drawn from its other non-cloud legacy systems

Cloud providers generally enable client businesses to port these licenses to their operating environment under 'bring your own license' schemes.

6.3 Procuring support for cloud migration and beyond

Businesses will need to decide whether their current IT skills base is sufficient to undertake the cloud transformation project or whether and where external support is required.

Organisations adopting the 'DIY option' for ongoing cloud management are, nonetheless, still likely to procure support for systems integration and migration.

Most organisations will be unable to pursue the 'DIY option', so the support role is likely to cover integration and migration but also extend to the management of the cloud environment(s) after the initial 'go live' date.

A 'managed service' typically involves augmenting the client's existing operational capacity to deliver:

- strategic support — in decision making, sign off, or alignment to business goals
- design and delivery of a structured cloud migration strategy
- going 'live' and beyond

Typically, this support will be provided by a cloud provider or one of its managed service partners. The choice of which will depend on whether:

- the cloud provider views the organisation as of sufficient scale and commercial opportunity to deliver support services direct or whether it refers the organisation to one of its managed service 'channel partners'
- the extent of the managed service required — for example cloud providers do not offer a full managed private cloud service based on dedicated hardware. However, that provision is a core element of every managed services provider offer

Effectively, only the larger housing providers will have the 'enterprise scale' to justify cloud providers delivering support services directly. Most will be directed by cloud providers to their managed services partners.

Managed services can be for dedicated or shared hardware. Providers offer two base procurement options:

- manage the client's hardware on the client's site or in the managed service provider's own data centres, **or**
- rent hardware to the client organisation, with that hardware being hosted in the provider's data centre

The managed service can be scaled at two core levels:

- the managed services provider manages the infrastructure and the data virtualisation platform, **or**
- the provider manages the entire IT stack and related dependencies up to operating system level — with the client organisation only delivering its responsibilities regarding data security and compliance and operating cloud-based applications

A typical business client/managed services provider operational interface might look like this.

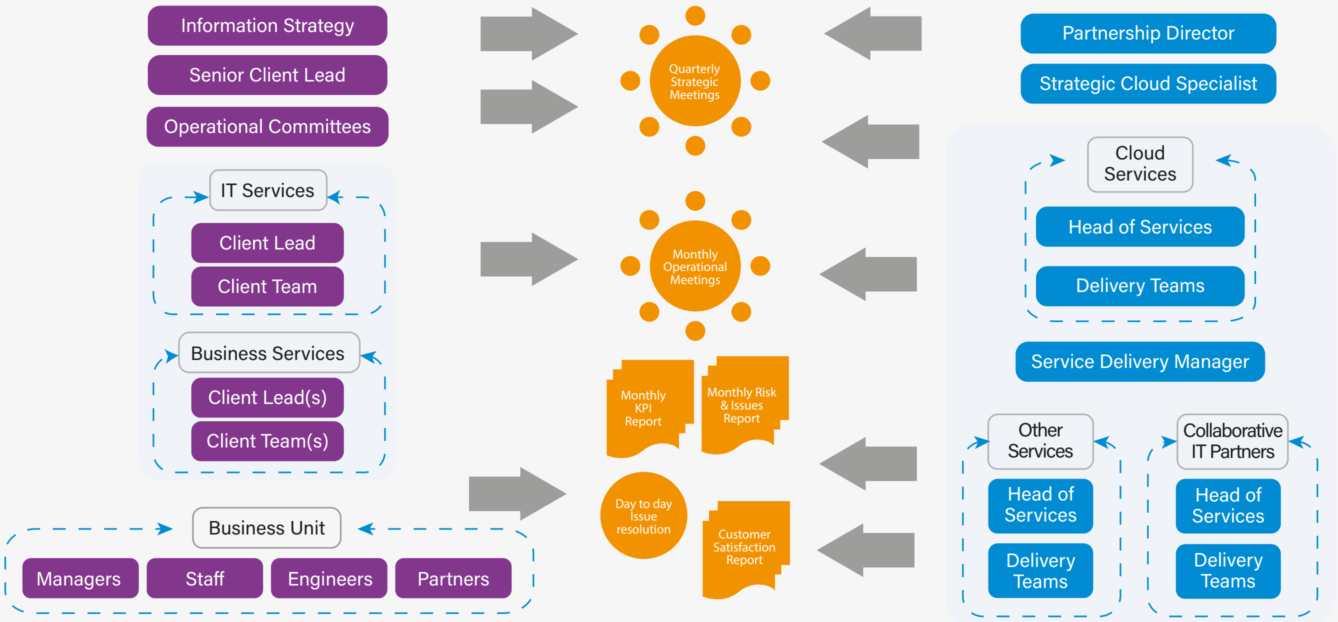

Selecting a managed service provider

Having determined the need to procure the services of a managed service provider, organisations then need to select one.

Key procurement considerations include:

- track record in understanding and tackling the business challenges of cloud migration
- understanding of the client's needs and how cloud technology should be configured to meet them
- technical knowledge and vendor neutrality required to advise client on 'best in class' and 'best buy' in terms of IaaS, SaaS and PiaS — added value can be provided where the cloud or managed service provider has a wide range of SaaS partner relationships
- technical skills required to plan, scale and optimise cloud migration
- technical knowledge required to integrate cloud infrastructure with existing systems and applications
- technical capability to provide cloud services (including external applications incorporated into the provider's environment)
- ability to engage with a multi-cloud environment and work with 'open source software (see Appendix A.2) to facilitate 'portability of supply' and thus enable clients to avoid being 'locked in' to proprietary systems and software
- ability to train existing staff and undertake a comprehensive knowledge transfer

Selecting a technical development provider

A technical development provider offers the expertise to migrate applications to cloud or rebuild them in the cloud. Some managed service providers also offer technical development services.

Key procurement considerations include:

- similar checks as with a managed service provider
- technical expertise in cloud-native technologies such as microservices and containers
- technical expertise in core-business applications
- ability to adopt 'open source' software and development tools - see Technical Appendix A.2
- examples of innovation based on latest technological developments
- proven case studies and references from similar deployments (including learning from failed projects)
- (for Paas but not SaaS) commitment and ability to train on-premises staff and undertake a comprehensive knowledge transfer
- application of best practice software development methods

Cloud procurement and climate change

It takes a tremendous amount of energy to manufacture and power our devices, data centres, and related infrastructural needs. The energy footprint of the IT sector is already estimated to consume approximately 7% of global electricity.

The internet's energy footprint is expected to rise further, fuelled both by our individual or business consumption of data and by the spread of the digital age to more of the world's population, from 3 billion to over 4 billion globally. The European Commission-funded Eureca project found that data centres in EU countries consumed 25% more energy in 2017 compared with 2014.

The leading cloud providers all recognise the issue and have made public commitments to convert their data centres to run on renewable energy in order to reduce emissions. In September 2019, Amazon pledged to be 'net zero carbon' by 2040. At Davos in January 2020, Microsoft pledged to become net zero carbon by 2030 and (without saying how) to remove as much carbon as it has emitted in its 45-year history.

Cloud computing is more energy efficient than on-premises provision. Shared capacity in cloud provider data centres on a multi-tenant basis, coupled with the retirement of client organisation data centres, can reduce data centre emissions by 25%.

Greenpeace's 2017 report *Clicking Clean: Who Is Winning the Race to Build A Green Internet* (see link below) provides an independent assessment of cloud provider performance. Greenpeace argues that companies entering the race to build a renewably powered internet are motivated by:

- ► customers who have carbon or renewable energy goals demanding that their digital infrastructure is powered by clean sources of electricity
- ► the rising cost competitiveness of renewable energy, with long-term contracts increasingly at cost parity or even beating fossil fuels in many markets, while also providing long-term price security
- ► competitiveness among IT companies and the linkage of brand identity with a renewable supply of energy, given the growing concern on climate change among employees and customers

Therefore, when procuring cloud product, housing providers may wish to stipulate that the data centres hosting their systems, applications and data should be powered by renewable energy.

https://www.greenpeace.de/sites/www.greenpeace.de/files/publications/20170110_greenpeace_clicking_clean.pdf

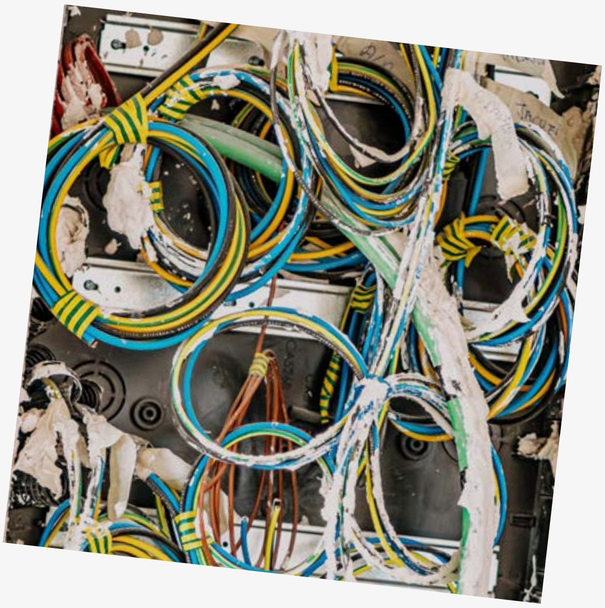

7. Conclusion: strategic options regarding cloud adoption

We have seen in Chapters 2 and 3 that whilst cloud-related cost savings can be considerable, they are not guaranteed.

Future costs will depend on procurement decisions, the repayment profile of legacy applications and infrastructure, the optimisation of cloud resource usage, the number of new SaaS applications purchased and the extent of application development work using cloud agile capabilities.

The business case for the cloud — its 'cost justifiability' — is that it's intrinsic flexibility and development speed enables any organisation to:

- flex IT resources and scale services to meet demand on a 'pay as you go' basis
- improve core IT resilience, disaster recovery and business continuity
- use cloud applications and development environments to, quickly and economically, deliver business solutions to changes in the housing sector operating environment
- transform into a more efficient and customer service-focused business

We have noted, in Chapters 3 and 4, that effective deployment of the cloud depends on managing and mitigating cultural barriers and risks — particularly on having the IT and business skills to manage the cloud environment and use it to deliver new business solutions.

In Chapters 5 and 6 we demonstrated the range of cloud deployment models and service options available to businesses and the available choice of suppliers.

We conclude by suggesting that, in determining its approach to the cloud, a business has three strategic options:

- Do nothing
- Lift and shift
- Cloud First/Hybrid

Do nothing

This is often expressed as a 'wait and see' strategy where a highly risk-averse business watches other peer organisations take the initiative on cloud adoption before deciding to engage with the cloud.

However, this approach comes with significant risks:

- the current IT suite becomes an increasing business risk — in terms of technical dependency, cost of maintenance, risk of failure and data loss
- the longer a business retains outmoded technology the more likely it will have to migrate to the cloud to resolve a 'burning platform' or constrained timeline issue thereby adding significant risk and cost to migration activities
- customer satisfaction declines — customers will increasingly be accessing services via the cloud in other aspects of their lives and will seek similar services from their housing provide
- attracting staff becomes more difficult — recruitment success, in IT and operations, will be increasingly determined by applicants being satisfied that they will be working in a modern cloud-enabled business environment
- the business goes stale — access to future business opportunities and the ability to compete with new disruptive entrants to the sector all potentially depend on its ability to act in an agile manner — with that agility being enhanced by the cloud

In short, from a financial, operations and customer perspective, Do Nothing is no longer a viable option.

Lift and shift

Housing providers need to actively consider their future strategy for core business legacy applications. At some point, these applications will need to migrate to the cloud or be replaced by new applications developed within a cloud environment.

The catalyst for taking these decisions can be:

- technical — hardware reaches the end of its useful life
- contractual — term contracts for existing legacy systems are set to expire
- performance-related — applications that are no longer fit-for-purpose are holding back the organisation from improving its services

Lift and shift, also known as 're-host', involves moving an application or workload 'as' is onto a new virtual server running in the cloud.

Organisations often adopt 'lift and shift' to address 'burning platform' issues — such as the failure of existing servers — within constrained timelines that preclude the more complex approach of building or rebuilding applications from first principles within the cloud.

Whilst this entry into a cloud environment may be enforced and unplanned, it is merely the first step in the cloud migration journey. Organisations can then plan for the functional optimisation of the migrated applications through 'right-sizing' their CPU usage and virtual server capacity, taking advantage of the intrinsic elasticity of provision and ultimately by moving to serverless computing over time.

Cloud First

Cloud First means that a business always considers operating a project, workflow, process, etc. in the cloud before considering any other option. Teams within the business are free to use other options but need to demonstrate that alternatives offer the right levels of security, flexibility and value for money compared to the cloud.

Cloud First does not mean 'cloud only'. While the cloud can handle a lot of tasks, many businesses still rely on on-premises hardware or software to perform various functions. Switching all those functions to the cloud may be more cost-efficient,but can sometimes be less productive.

Choosing to adopt a cloud solution, or not, should be undertaken on a case-by-case basis rather than being strictly forced. Instead, the expectation should be that when designing strategies and planning workflows and processes, the business will always contemplate a cloud-based solution in the first instance.

The challenge for social housing businesses is, as we have noted, the proliferation of non-cloud core business applications. Each of these will require a medium to long-term strategy for cloud migration. In practice, businesses may not wish to review current legacy provision until contracts come up for renewal.

In such cases, a housing provider may opt for a 'hybrid' approach, where cloud migration becomes a phased medium-term strategy linked to the investment cycle of the business. For example, 'quick wins' — such as disaster recovery — can be programmed for cloud migration whilst the legacy systems issues are resolved.

It is unlikely that housing providers will be able to migrate their entire environments over to a public cloud in the short-term. They are likely, therefore, to adopt a hybrid solution whilst they run down their pre-existing contracts or refactor or re-procure some of their applications.

A hybrid strategy also allows executive teams and Boards to retain sensitive processes and data on-premises until such time as they are comfortable with migrating these to the cloud.

The hybrid approach is active rather than passive. It requires awareness of the cloud and its benefits, a critical evaluation of risk and the energy and commitment to actively pursue cloud adoption.

Strategically, the hybrid option is the lowest risk for a business. Conversely, the 'do nothing' or 'wait and see' approach, constitutes the biggest risk. Given the ongoing issues with migrating legacy systems, discussed in Chapter 4, we expect the 'hybrid strategy' to be adopted by most businesses in the social housing sector.

Whatever option a housing provider chooses to take — and we do not recommend 'do nothing' — it cannot escape the need to evaluate the benefits of cloud computing.

There is no better time to begin than now.

Technical Appendix

by Phil Brunkard

In this section we consider some of the many emerging technical developments in the cloud and take a deeper look at cloud optimisation and migration options.

A.1 Modern public cloud services — cloud native computing, microservices and container frameworks

Cloud native computing

Over the last few years there has been a rapid evolution of cloud technologies — collectively known as 'cloud native' computing. This shift has fundamentally changed the way software development is done today, leading to the adoption of new practices such as DevOps that accelerate the release of new application services and capabilities.

Cloud-native is an approach to building and running applications that fully exploits the advantages of the cloud computing delivery model. Apps live in the public cloud, taking advantage of modern cloud techniques such as PaaS, microservices and containers.

Microservices

Microservices is an architectural approach that breaks large applications into lightweight apps that can be scaled horizontally. Businesses shift from having a complex internal architecture to smaller, independently scalable applications (the microservices). The idea is that each microservice is small and less complex to develop, update, and deploy. The apps are loosely coupled, meaning the code is not hard-wired to any of the infrastructure components, so that the app can scale up and down on demand.

Microservices can be an effective strategy when legacy application architecture wont scale further, e.g. if the database grows too large or there are too many millions of lines of code or the business simply can't add features quickly enough and cost-effectively to meet its objectives.

Microservices provide an opportunity to deal with some of the constraints of legacy applications and the complexity faced when new feature is needed.

Rather than face the challenges with upgrading functionality in these applications — and dealing with the associated costs, timescales, specialist supplier expertise and associated risks — businesses can create the functionality using microservices.

Choosing to refactor a legacy application in this way is a major business decision. The business will certainly need help from an experienced cloud managed services partner with the capability to deliver with minimal risk. The business will need to assess if will need to procure a managed service 'post-migration' and to what extent it can develop and maintain such cloud native skills in-house.

Despite their benefits, microservices bring a new set of problems that are hard to tackle — logging, monitoring, testing, and debugging any loosely coupled applications made up of many different microservices.

If a legacy application is still working well and doesn't need to be changed much, adopting microservices will be an unnecessary expense.

Container frameworks

Container frameworks such as Docker are gaining popularity for software deployment in the cloud. Docker acts as an orchestration layer across multiple containers. Like virtual machines, containers guarantee software will run regardless of the host environment but are far more lightweight and resource-efficient in deployment and operation. Containers hold everything an application needs to run. This could include files, libraries and environment variables.

As containers don't require their own operating systems to execute applications, they consume fewer resources. However, container frameworks are also complex to configure and implement and need specialist expertise and skills to be effective.

Serverless computing

Serverless computing enables developers to build and deploy applications faster without the need to manage the underlying infrastructure. With serverless applications, the cloud service provider automatically provisions, scales and manages the infrastructure required to run the code. It is like PaaS — with the exception that capacity and utilisation of servers is fully automated. Serverless architectures are scalable and event-driven meaning resources can be specifically assigned when a specific function or trigger occurs.

These technologies have been instrumental in the success of digital native companies disrupting existing business models like Uber and Netflix.

Devops

DevOps is a combination of the skills, processes and practices across software development, quality assurance and IT operations with the aim of releasing software updates to customers much quicker than traditional practices.

The growth of DevOps is a result of the advances in cloud technology enabling a more agile approach to IT delivery and operations. The principle is that by collaborating across both development and operational teams and automating the process of software delivery and infrastructure changes this will make an organisation more agile and efficient in responding to customer expectations.

Instead of one team building software and another separate team supporting it, everything is done continuously to shorten development, testing and release cycles.

Devops is a cultural change more than a technology or process change but it needs to be fully adopted if businesses are to realise its benefits. This can be challenging in an organisation with considerable legacy technology and clunky business processes that may be difficult to unbundle and modernise.

DevOps aims to deliver software quickly and effectively by fostering collaboration between different departments. This is supported by automation, and by analysing business-wide metrics, to validate what's working well and what's going wrong.

A key element is continuously obtaining customer feedback, real-time automated testing and immediate deployments and software updates and testing.

DevOps ensures development teams take responsibility for the quality of their code, listening to customer feedback for improvements and fixing bugs. With automation this can be more productive. Devops requires a range of tools to support its practice.

Devops processes can either be managed by the cloud provider or by in-house teams. As most organisations won't have these skills or be experienced in the cultural practices of DevOps, it is normal for the cloud provider to do this — with training and skills transfer over time to in-house IT teams.

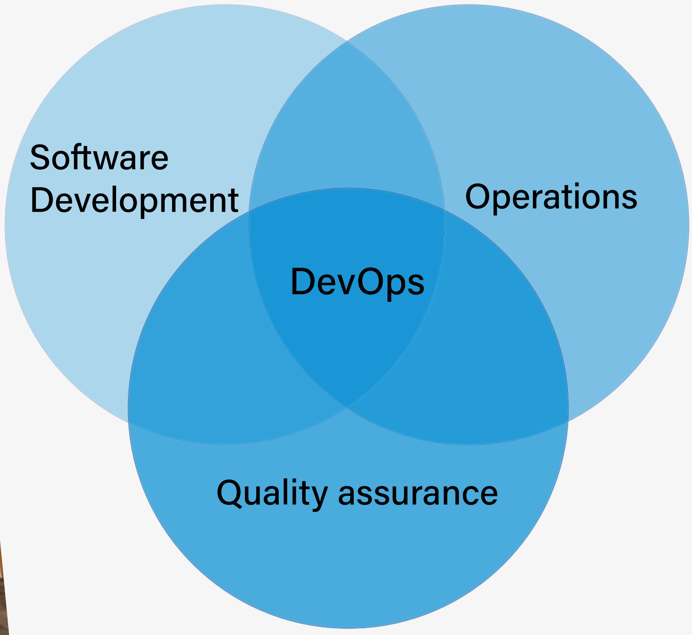

A.2 'Open source' and cloud

The use of open source tools for cloud is increasingly common. These tools enable cloud infrastructures with more flexibility and less cost in developing and delivering services.

An open source cloud service is any cloud solution that is developed using open source technologies and software. This covers any public, private or hybrid cloud models providing SaaS, IaaS and PaaS that have been built and operate entirely on open source technologies.

The principles of open source cloud are:

- not locked-in to any single vendor
- source code is made available to the community and is free from royalty. Users can modify and redistribute copies of either the original or a modified version
- enables open standards and open data formats allowing greater data sharing and integration for new open solutions
- enables greater interoperability with other solutions through reusing recommended software stacks, libraries and components

There are several aspects that businesses need to consider before they integrate any open source technologies into their cloud environment.

Consideration	Details
Communities and support	Most open source projects have dedicated user communities, consisting of experts who share resources on how to build and develop solutions using the open source technology. Community-led projects also create new features, push out updates and fix bugs. This is effectively an unofficial support model but without any service levels or obligations. When an open source software developer releases its code to the public, it generally leaves it to the user community to take care of further developments. Users rely on other users for support. For open source cloud computing, users might need to become active community members in order to diagnose and fix issues. If they can't solve the problem, they need to ensure the community can.
Costs	Open source cloud solutions are not entirely free. Traditionally, open source software is free to install and use, but depending on the software, there are likely costs to maintain and update it. Open source communities provide free resources for users, so building a cloud infrastructure on open source programs will still save money. Cloud providers have developed services using open source with a managed service wrap, commercialising their offerings without developing the original open source product. This has led to many open source solution providers having to license their products in different ways. On the other hand, there are examples of cloud providers developing open source projects and making these generally available, e.g. the Kubernetes containers framework from Google.
Knowledge/ expertise	Open source software allows developers to modify the tool's code to fit their needs. Therefore, to get the most benefit, businesses need technical know-how to build and alter code. While open source programs may have user-made guides and resources, these do not guarantee that a business can successfully achieve what it needs without significant effort.
Data ownership	Any user of an open source program maintains data ownership. Keeping data ownership in check is one of the principles of open source programs. However, businesses still need to bear in mind the risks of putting confidential data into the public code when using public cloud open source tools.
Maturity	Open source cloud tools only started appearing in the last few years and are still maturing. As a result, businesses are likely to experience bugs, non-optimised code or other issues.

There are many different open source products, tools and software across the different cloud computing layers. Examples include:

- open source databases: MySQL, MongoDB
- PaaS: OpenStack or Cloud Foundry — used for creating, configuring, building, testing and deploying applications on the cloud
- container orchestration tools: Docker or Kubernetes — Used to create, deploy, and run applications in containers. These allow a developer to wrap up an application, with everything it needs, into a single package. This ensures the application will run on any machine regardless of any customised settings which differ from the machine on which the application was created.

A.3 Cloud tooling

There are several useful tools on the marketplace that the cloud and managed services providers use to guide businesses in making the right cloud migration transformation decisions.

Examples include Cloudamize, Corent SurPaas, Movere (now part of Microsoft) and Device42. These are discovery tools that assess existing IT Infrastructure in order provide comprehensive data and insight to allow businesses to identify and plan what they can move to the cloud.

Cloud tools frequently scan the business's IT environment to identify utilisation, server and application workloads, consumption, peaks and capacity.

From this insight businesses can then explore different scenarios and options for cloud migration. The tools guide them on the best patterns and solutions to adopt based on factors such as performance, application dependencies, suitability and readiness as well as predicted costs compared against current costs. The tools provide recommendations for optimal cloud configuration, including server compute instances, storage options, network settings, and pricing plan. However, they can only provide recommendations based on infrastructure sizing and utilisation — they don't take into account application complexity, design etc.

Businesses can also compare options for moving to AWS, Microsoft Azure or the Google Cloud Platform.

A4. Cloud optimisation — elasticity and scalability

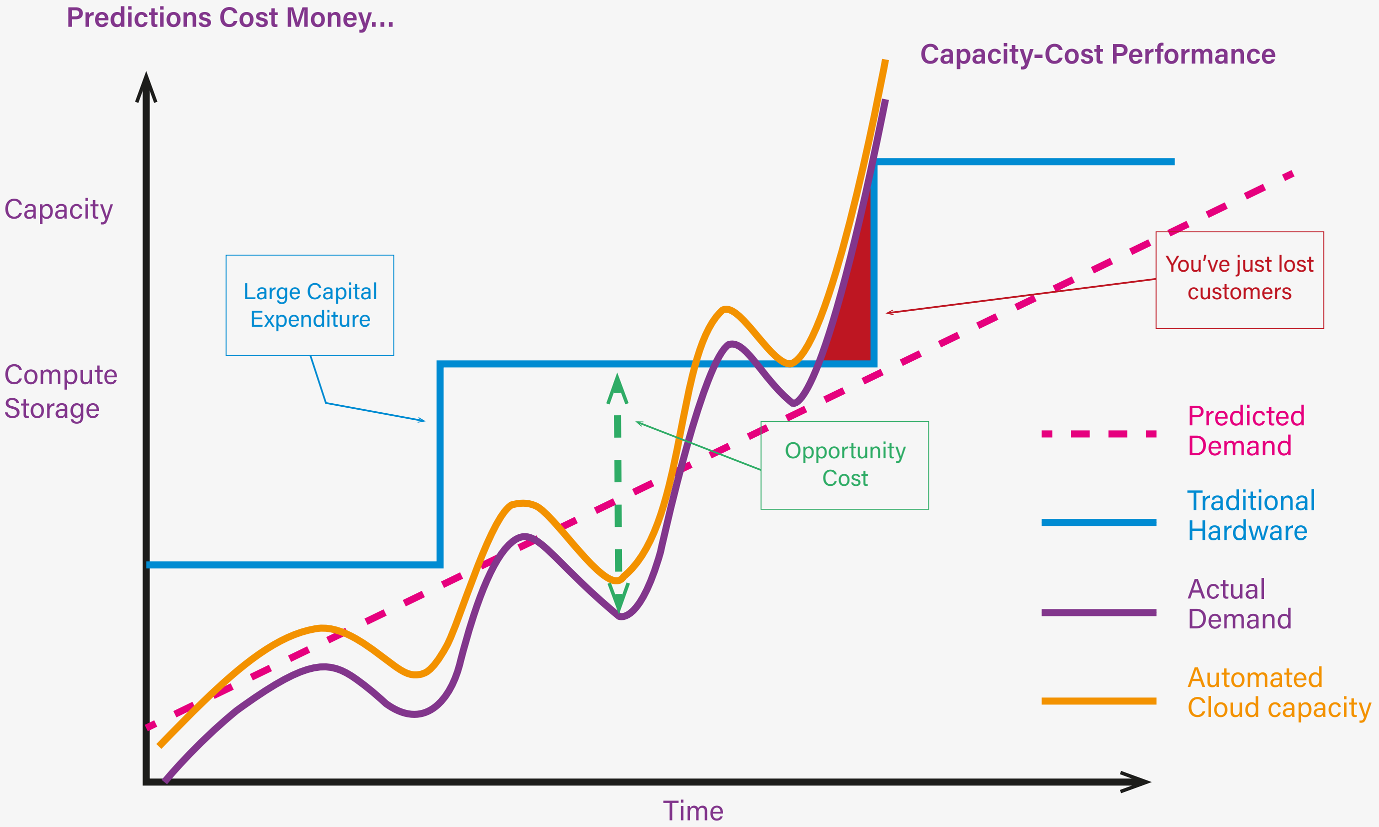

Source: Amazon Web Services

The capacity/utilisation curve used by Amazon Web Services to illustrate the advantages of cloud computing demonstrates how cloud optimises deployment of IT resources through on-demand provisioning of services to meet actual usage.

Without cloud, resources are typically sized to handle peak loads, which mean they are under-utilised at off-peak times. Often, servers and storage are dedicated to specific functions or departments and can be massively under-utilised. It may also be common for departments to request new servers for new projects, even though other teams have underused server capacity.

Organisations will invest (CAPEX) based on predicting demand for these peaks, leading to unused investment where actual demand differs. If the organisation fails to keep pace with actual demand due to CAPEX budget cycles, this can cause potential risk to overall business performance and customer experience.

Cloud computing enables resources to be shared by different loads and thus improves utilisation. The sharing can be between different organisations with public cloud, or within an organisation via a private cloud.

With IaaS, resource sharing means supporting many platforms and applications on the same physical resources, using virtualisation techniques. At the platform and application levels, providers can achieve resource sharing through multi-tenancy. For PaaS, this means applications from many clients running on the same operating system. For SaaS, it means clients sharing the same application instance.

This is how cloud service providers can achieve better utilisation of the underlying assets than if each client uses either a separate application instance on the same operating system or another virtualised instance sharing hardware.

One of the core concepts of cloud computing is to avoid the cost impact of over-provisioning and under-provisioning. The resulting savings are in addition to the financial efficiencies enabled by rapid deployment of cloud services with low entry cost. Cloud scaling and cloud elasticity principles enable this capability.

Cloud elasticity

The purpose of elasticity is to match the resources allocated with the actual amount of resources needed at any given point in time. The cloud provider will dynamically allocate resources to an organisation's application process based on the needs of that process, providing the exact amount of resources needed to meet KPIs.

Cloud providers have systems and tools in place to automatically deliver or remove resources in order to provide just the right amount of assets for each application. The cloud user is provided with enough power to run their workflows without wasting money on resources they do not need.

The main purpose of cloud elasticity is to avoid either over provisioning or under provisioning of resources. Giving a cloud user either too much or too little data and resources is a disadvantage. If the user has too many resources, they pay for assets they are not using. If they have too few resources, they cannot run their processes correctly — thus negatively impacting performance and customer experience. An example use for elasticity is where a housing provider has seasonal demands such annual rent statements.

However, businesses that do not experience sudden or cyclical changes in demand may not benefit from the cost savings elastic services offer — as demand is static, they don't need the additional functionality and automation that elastic services provide.

Cloud scalability

Scalability is the ability to increase workload size within existing infrastructure environment (hardware, software, etc.) without impacting performance. Usually resource capacity is pre-planned and will include a certain amount of headroom built in to handle peak demand.

This is where procuring reserved instances can be advantageous. When resources are no longer needed there is usually the option to manually scale down the infrastructure. Increasing or decreasing resources is a planned event and static for the 'worst case' workload scenario. The approach allows the organisation to scale up performance, without having to worry about not meeting SLAs in a steady pay-as-you-grow solution.

Both capabilities are important considerations where infrastructures are constantly changing. The decision to consider elasticity or scalability depends on the business need or use case to determine the best choice.

Cloud scalability is generally delivered more readily in private cloud environments while cloud elasticity is generally delivered more readily in public cloud environments.

A5. Cloud migration strategy options for applications

In 2011, Gartner published five cloud migration strategies for applications — referred to as the 5 Rs. These were initially based on options for how a business might consider IaaS, PaaS or SaaS options.

Since then, cloud and managed services providers have revised, developed or modified the principles of the 5 Rs to include modern cloud native capabilities such as micro services and containers. AWS, for example, refers to the 6 Rs. Each migration strategy has different pros and cons.

The method known as 'application sentencing' involves each application being assessed against each sentence type to guide the migration plan. This then informs the business case — but should be deployed within the context of determining the option that best aligns with the culture, business objectives and risk appetite of the business.

Retain

This is not usually offered as an option by cloud or managed service providers but is nevertheless a valid tactical short-term consideration.

The situation may be that the application is too difficult and/or expensive to migrate to a public cloud platform within the current roadmap. There could be technical, service, commercial or contractual dependencies that are challenging to resolve and need further analysis or negotiation. Factors may include:

- third party service gaps
- lack of in-depth understanding and documentation
- compliance or regulatory dependencies
- vague business requirements, or
- uncertainty on how to design and test in order to migrate the solution without considerable business disruption.

Retire

One of the benefits of undertaking a cloud readiness assessment is the opportunity to identify and switch off applications that are not being used but are consuming resources or costing more money to operate than the revenues they generate.

The application discovery (using tooling or other methodologies) will identify the candidate applications for retirement. Often the exercise will uncover unknown systems, systems that may be a legacy from a previous M&A activity (left in place just in case it is needed) or applications where the functionality could be better delivered elsewhere. Typically, this equates to about 10% of an IT portfolio. The assessment also pushes businesses to make decisions on current services due for retirement and will often provide an immediate cost saving as a result.

Repurchase/Replace

The principle is to discard an existing service and use commercial software, delivered as a service. This is normally based on a case when a SaaS company provides a better solution than something the business already has or can build itself.

Effectively, this is a licensing change. Instead of using a traditional on-premises license, a business can start using the same application as a cloud service. This is a smaller effort than 'lift and shift' because the business is not moving anything — just starting a new license agreement in the cloud — but still requires business process mapping, data migration and training. The obvious examples are migrating to Microsoft Office 365 or Salesforce.com.

Re-host

Also known as a 'lift and shift.' Re-host migration moves a current system to a cloud environment, with minimal change to its overall architecture.

If businesses are looking to speed up migration, re-hosting will potentially give them some immediate benefits — such as operational savings (once they have sized the target hosted cloud platform based on actual usage rather than provisioned patterns).

Revise | Re-platform

This is where a business might need to move some components of an application to the cloud. It could, for example, move its on-premises Oracle database to a public cloud database-as-a-service or move from a vendor specific web server product such as Oracle WebLogic to an open source solution such as Apache Tomcat.

As this may result in the application being deployed across a hybrid cloud configuration, it will require a good understanding of application workload, performance and traffic dependencies to validate the business impact and costs involved. Platform as a service (PaaS) options can reduce the operational costs that are associated with many applications.

Refactor | Re-architect

Often, aging applications aren't compatible with a cloud environment because of architectural decisions that were made when the application was originally built.

In these cases, the application might need to be re-architected before it can be migrated to the cloud.

In other cases, applications that are cloud-compatible, but not cloud-native, might engender cost and operational efficiencies if components are refactored into a cloud-native application.

This approach allows the business to evaluate weaknesses in existing applications and remedy these by applying cloud features to their architecture or code e.g. by deploying micro-services.

This involves considerable work, so businesses will need a strong business case to refactor existing applications to the cloud.

Often an application will need to be refactored to fit a PaaS-based model and might not be viable as a straightforward re-platform migration.

Re-build

In some scenarios, the work required to carry a legacy application forward can be too large to justify further investment.

This is especially true for applications that previously met the needs of a business but are now unsupported by software suppliers or where the applications no longer align to current business processes.

In such cases, businesses may not identify any alternative SaaS solution as being viable — so a new code base will need to be created to align with a cloud-native approach.

A business may therefore decide to completely discard the legacy application and develop a fresh capability using cloud services and features. This approach requires good familiarity with existing application and business processes as well as cloud services. However, rebuilding core business applications is highly complex and will require on significant business buy-in and potentially external technical support.

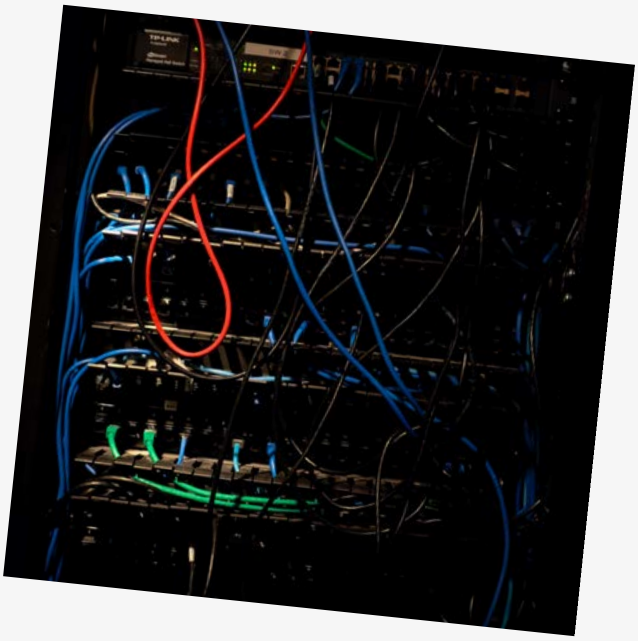

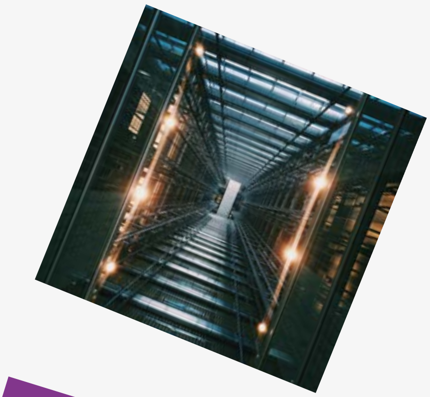

Glossary

- **Application** — An application is a group of computer programs designed to allow a user to perform a set of functions or tasks.
- **Application programming interface (API)** — an interface that allows the user to access information from another service and integrate this service into their own application.
- **Business continuity** — the daily activity performed by an organisation to ensure that critical business functions will be available to customers, suppliers, regulators, and other entities that must have access to those functions. These activities include many daily chores such as project management, system backups, change control, and help desk.
- **Containers** — Like virtual machines, containers guarantee software will run regardless of the host environment but are far more lightweight and resource-efficient in deployment and operation. Containers hold everything an application needs to run. This could include files, libraries and environment variables.
- **Cloud computing** — the delivery of information technology services over a network, usually the internet.
- **Cloud management platform** — a product that gives the user integrated management of public, private, and hybrid cloud environments.
- **Cloud marketplace** — an online marketplace, operated by a cloud service provider, where customers can browse and subscribe to software applications and developer services that are built on, integrate with, or supplement the cloud provider's main offering.
- **Cloud migration** — the process of transferring all or a piece of a company's data, applications, and services from on-premises to the cloud.
- **Cloud native** — Cloud-native is an approach to building and running applications that fully exploits the advantages of the cloud computing delivery model. Apps live in the public cloud, taking advantage of modern cloud techniques such as PaaS, microservices and containers.
- **Cloud portability** — the ability to move applications and data from one cloud provider to another.
- **Cloud Service Provider** — a company that offers a cloud computing service, such as PaaS, IaaS, or SaaS, to individuals or businesses.
- **Cloud storage** — a model of computer storage in which data is stored in facilities (often multiple facilities) managed by a hosting company (cloud service provider) and is accessed remotely by the user via a network.
- **Cyber Essentials** — a UK government information assurance scheme operated by the National Cyber Security Centre (NCSC) that encourages organisations to adopt good practice in information security. It includes an assurance framework and a simple set of security controls to protect information from threats coming from the internet.
- **Data migration** — the process of moving data between two or more storage systems, data formats, warehouses or servers.
- **Elasticity** — the ability of a system to adapt to changing workload demand by provisioning and deprovisioning pooled resources so that provisioned resources match current demand as well as possible.
- **Hybrid cloud** — a cloud computing environment that is comprised of a mix of private cloud, public cloud, and on-premises solutions. In a hybrid cloud, private and public cloud infrastructures remain distinct from one another but are bound together by technology that allows data and services portability between them.
- **Infrastructure** — Information technology (IT) infrastructure is a combined set of hardware and virtual resources that support an overall IT environment.
- **Infrastructure as a Service (IaaS)** — a model of cloud computing in which the vendor hosts virtualised computing resources, as well as network and storage resources, and provides them to the user as a service via the internet.

- **Managed Service Provider (MSP)** — an IT services provider that provides fully outsourced network, application, and system services across a network to clients.
- **Microservices** — a way of designing applications in which large complex applications are broken down into lightweight independently deployable apps that can be scaled horizontally.

  These 'microservices' run their own processes and communicate with one another using lightweight mechanisms such as language-agnostic APIs. Businesses shift from having a complex internal architecture to smaller, independently scalable applications (the microservices). The idea is that each microservice is small and less complex to develop, update, and deploy and scale up and down on demand.
- **Multi-cloud** — the concurrent use of separate cloud service providers for different infrastructure, platform, or software needs. A cloud management platform is recommended for this approach.
- **Multi-tenancy** — a mode of operation for software in which multiple instances of one or many applications run in a shared environment. In a cloud computing model, pooled physical and virtual resources are dynamically assigned and reassigned to tenants according to consumer demand.
- **On-premises** — software or infrastructure that is run on computers on the premises (in the building) of the person or organisation using the software or infrastructure. Microsoft Office Suite is an example of on-premises software because it needs to be installed on the computer that runs it, while Office365 is not, because it is accessed via the internet and run remotely.
- **Open Source** — a development model in which a product's source code is made openly available to the public. Open source products promote collaborative community development and rapid prototyping.
- **Platform** — a computer system that applications run on, or as a base of technologies on which other technologies (such as applications) are built.
- **Platform as a Service (PaaS)** — a model of cloud computing in which a vendor provides the hardware and software tools necessary to create, deploy and manage applications at scale to the user via the internet, as a service.
- **Private cloud** — a cloud infrastructure that is provisioned for use by a single organisation comprised of multiple users. A private cloud can be managed and operated by the organisation, a third party, or some combination of them, and it can exist on or off premises.
- **Public cloud** — a cloud infrastructure that is hosted by cloud services provider and is made available to the public via the internet.
- **Reserved instances** — Where a business commits to resources for an agreed capacity and term (e.g. 1 or 3 years) in exchange for a sizeable discount
- **Serverless computing** — Serverless computing is most easily understood from a developer's perspective. A serverless component performs a single task, such as storage, computation, or access control. Common examples include Azure Functions, Google Cloud Storage, and AWS IAM. Respectively, these are used to supply compute, storage and access control to applications built in the cloud. The developer interacts with a component via a published API and cannot modify the underlying platform.
- **Scalability** — the ability of a process, system, or framework to handle a growing workload. In other words, a scalable system is adaptable to increasing demands
- **Software as a Service (SaaS)** — a model of cloud computing in which applications (software) are hosted by a vendor and provided to the user as a service. SaaS applications are licensed on a subscription basis and are made available to users over a network, typically the internet. A well-known example of SaaS is Microsoft's Office 365
- **Vendor lock-in** — when a customer finds themselves 'locked-in' or stuck with a certain cloud service provider. Vendor lock-in is characterised by extreme difficulty in moving from one cloud vendor to another, usually due to lack of standardised protocols, APIs, data structures, and service models.
- **Virtual machine (VM)** — a software computer that runs an operating system or application environment, just as physical hardware would. By running VMs, a hardware computer can run multiple instances of the same operating system. Applications running on separate instances cannot interfere with each other, so if one app crashes, it will not affect apps on other VMs.

Author's acknowledgements

I owe the greatest debt of gratitude to Phil Brunkard, Senior CIO adviser at Forrester (and formerly Head of Technology Strategy & Innovation at BT).

Phil generously shared his extensive knowledge with me, and we co-authored the early drafts of this publication before Phil's full-time move to Forrester made it impractical to continue his involvement. Although the text has undergone many substantial revisions since then, the Technical Appendix survived that process intact and is credited solely to Phil.

I am also most grateful to **Paul Jackson**, Solutions Director at **Rackspace**, who was equally generous in sharing his extensive technical knowledge and made many critical contributions during the editing process.

Particular thanks are due to our sponsors **AWS** and **Rackspace** and specifically to **Simone Hume**, **Chris Masey** and **Tom Griffiths** at AWS and **Curtis Fielder** and **Shelley Wright** at Rackspace.

Warm thanks are also due to our Readers Group, which commented critically and helpfully on the first substantive draft, and whose members included:

Jonathan Creaser — Head of IT, Salvation Army Housing Association

Mark Lordon — CIO Origin Housing Group

Rob Peck — Director of Procurement Services, Inprova Group

Paul Jackson — Solution Director, Rackspace

Leon Blakely — Senior Manager, Commercial Sales, Rackspace

Colin Sales — Managing Director, 3C Consultants

Tony Smith — Managing Director, Acutance Consulting

Arturo Dell — Director of Technology and Innovation, HouseMark

Dominic Pride — Founder and Chief Instigator, Upstart Strategy

David Treanor — Author and founder of M3 Consulting

Finally, this report could not have been produced without the personal support and financial investment of **Ian Wright**, Chief Executive, **Disruptive Innovators Network**.

About the author

Ross Fraser is Research Director at the Disruptive Innovators Network (DIN), where he is responsible for all long and short-form research and for editing the DIN bulletin.

Ross is a highly experienced writer, researcher, and editor who specialises in the innovation of professional practice in the social housing sector. He is the author and editor of over a dozen major publications on social housing.

Ross is probably best known as the founding chief executive (2001 to 2016) of HouseMark, the cross-sector business intelligence provider owned by the Chartered Institute of Housing and the National Housing Federation.

About the Disruptive Innovators Network

The Disruptive Innovators Network is the membership network for social housing organisations investing in innovation.

Our mission is to ensure members to make sense of disruption, be more innovative and grab the opportunities to build back better and to create even more social impact.

Our objective is to help leaders to work collaboratively, learn from each other and out-of-sector disruptors and share the cost and risk of innovation.

How do we do this?

- By horizon scanning future developments, trends and technologies to assess their impact and value to the social housing sector
- Inspiring and engaging members to work 'on' the business to see the bigger picture by connecting with out of sector leaders.
- By expanding members' peer networks with leaders in other sectors and countries looking to tackle similar problems
- Through identifying the risks and opportunities that disruption will bring to our member organisations
- By working alongside our members to navigate the new business models that may disrupt the sector
- By creating new opportunities and a safe space for our members to work collaboratively to define problems better and identify innovative solutions

To find out more about the benefits of DIN membership, please contact Ian Wright, the Chief Executive of Disruptive Innovators Network, by calling 07946 509 322 or via email at ian.wright@disruptiveinnovatorsnetwork.co.uk

Design

Publication design by Grace Abell abelldesign.co.uk
grace@abelldesign.co.uk