The Business Case for Cloud IaaS

A leadership perspectives white paper
Recommended next steps for business and industry executives

Issue 10 in a series

وتحيا بها الحياة
add life to life

Executive Summary

Business needs, economic concerns and technology trends are driving many organisations to pursue IT strategies that include near and long-term plans for cloud services. The implementation of infrastructure-as-a-service (IaaS) is viewed by many adopters as a critical first step toward broader cloud adoption. In many cases, customers' experiences of working with providers to deploy IaaS services will determine their trust and comfort levels for expanding into other types of cloud services. So selecting a capable and trustworthy partner is critical.

UAE cloud market supply and demand

Major telcos like du have led the competitive march from telecoms into cloud computing, and are now widely seen as credible players in the provision of cloud computing services, including IaaS.

Telcos have a long heritage in providing data centre managed services, and have combined this with their networking and security expertise to address customer requirements for enterprise-grade cloud computing services. Their cloud computing focus has initially been primarily on IaaS, and this will likely continue to be the case as they build out services to appeal to a global customer base.

Cloud computing technologies that enterprises currently have or use and are prioritising for investment in the future (Source: Ovum – 54 UAE respondents)
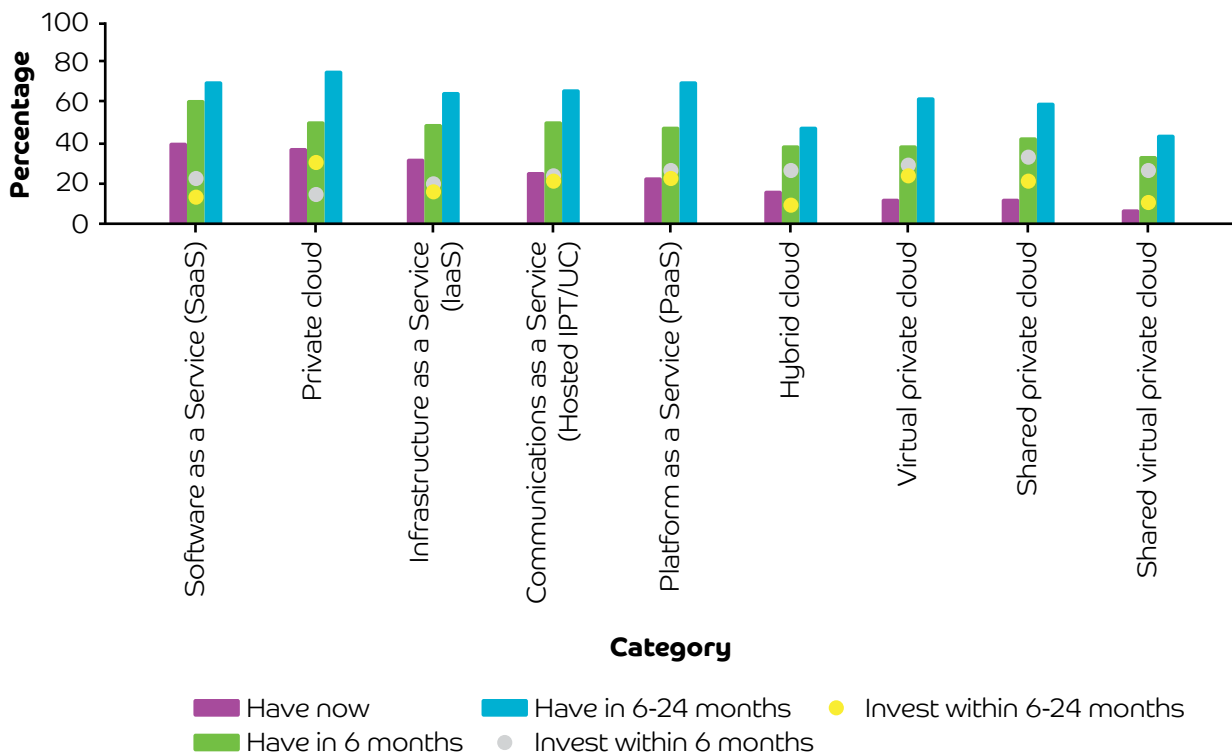
To date, the cloud focus of telcos has built around the importance of communications-as-a-service, variously referred to as collaboration-as-a-service or unified-communications-as-a-service.

Telcos also recognise the position of business ready service-provider-hosted voice services (IP PBX and UC capabilities) in terms of having some of the attributes of cloud services. This is becoming another important focus for them, led by a new generation of on-demand voice and unified communications services. The telco positioning for IaaS is firmly in the area of 'enterprise-grade' services, to distinguish the proposition from that of the likes of Amazon, which is seen as the player that most IT managers associate with the early adoption of cloud computing (and specifically IaaS). The term 'virtual private cloud' best describes the telco's IaaS offering, and is well understood by IT managers because of its similarity with the concept of virtual private networks. This means that cloud services are delivered over secure public IP (usually IPsec or SSL encryption) or over MPLS VPNs.

The concept of enterprise-grade can also be reinforced by capabilities that go beyond those usually seen in an IaaS offering. For example, telcos can offer physical servers on demand in addition to virtual servers. Through a service provisioning engine, it is possible to arrange services so that a customer does not share a server blade with other customers, as an additional reassurance regarding security, for instance.

One other topic that dominates many cloud discussions is the service providers' ability to maintain so-called hybrid environments, where there is a mix of traditional and cloud delivery models. Customers can keep their infrastructure and applications in-house, access traditional managed services, and/or move workloads to the cloud.

According to Cisco, IaaS is an area of cloud computing that currently receives the most market attention, with many enterprises planning to adopt IaaS via an external provider. In the UAE, a survey by Ovum has revealed that the use of IaaS is already on the radar for 30% of organisations, and that this level looks set to double over the next two years.

Infrastructure-as-a-service (IaaS) defined

IaaS delivers the three basic resources of a computer: processing, storage, and network I/O resources. Amazon Web Services is the best-known IaaS public cloud that combines all three and adds others. Many IaaS offerings are not as broad. Some, such as Nirvanix for example only provide storage-as-a-service.

IaaS mixes standardised hardware (servers and routers) and software components (hypervisor, operating system, management software) with associated services such as Internet addressing, directory, and security into a wide variety of offerings. These vary in terms of:

- scope – some are broad, such as Amazon Web Services; others are narrowly focused, such as on storage.

- technology – most are virtualisation-centric, but an increasing minority are not.

- target – some target IT operations; some other developers.

- configuration options.

In addition to generic cloud computing and public cloud benefits, IaaS enables enterprises to:

i. Reuse existing technology, code, and skills, and minimise lock-in. It does this by making it relatively easy to move IT resources (usually packaged as virtual software appliances) on to and away from an IaaS infrastructure.

ii. Have a good level of configuration control over the infrastructure resources they use such as CPU type, memory amount, and disk configuration, and the software (virtual appliance) they run on top of these resources within the constraints of the vendor's service catalogue.

IaaS consumers are responsible for the creation, deployment, and management of the virtual machines (VMs) and/or applications that use IaaS resources, and for defining and configuring the resources required by their VMs/applications, as well as for adding more resources if required. There is a range of management services available to cater for a range of enterprise requirements for 'hands-on' versus 'hands-off' operations. However, consumers do not manage the infrastructure that underpins IaaS resources. This is the responsibility of an IaaS provider like du and its partners.

IaaS is mostly for data storage, backup, and disaster-recovery purposes for data that requires limited levels of security, and to meet ad hoc compute-intensive batch-processing and analytics needs that internal data centres would not be able to cater for from a resource or cost-effectiveness point of view.

The market started with a focus on storage services then moved its attention to compute services. Now network issues in terms of bandwidth, latency, response time, and throughput are becoming more prominent because they prevent a lot of applications from being moved to the cloud. Vendors address this via a variety of technologies such as content and application-delivery networks, WAN optimisation, or a move from Internet to MPLS network technology, to deliver the end-to-end SLAs that enterprises are asking for.

The rise of cheap IaaS-based analytics-centric services is currently boosting the business-intelligence sector.

The economic arguments, for and against IaaS

Amazon has become the yardstick by which every other IaaS offering is measured. IaaS providers define their pricing strategy against Amazon's. A wide ecosystem is expanding around it, with independent software vendors (ISVs), for example, keen to:

- add value on top of it.

- use it as a new software delivery channel.

- build a business on top of it.

Most cloud strategies have an initial focus on cost reduction or improved cost efficiencies. In fact, cost savings are now often an explicit and definite objective. Many organisations are strengthening benefit realisation processes that increase executive accountability for cost-saving targets and this is flowing through to an increasing interest in cloud computing.

With about 50% of IT budgets dedicated to infrastructure acquisition and management, IaaS could represent a significant source of savings.

Clouds reduce not only the cost and pain, but also the risk, of using IT assets. This is certainly the case for IaaS, with its pay-as-you-go approach to pricing and licensing, than for SaaS (software as a service like Salesforce.com or enterprise versions of Google Apps), which favours the less flexible subscription approach.

IaaS has a usage-based/metered approach with no long-term commitment or contract and the ability to scale usage up or down on demand. IaaS service providers do not penalise users if their level of usage changes, which allows users to do away with capacity planning. (However, this will require internal controls to ensure that users do not go overboard.)

It makes it easier not just to use IaaS services, but to:

- **Meet unpredictable demand:** the ability to do this cost-effectively and with an acceptable QoS is, according to some, the key value that public clouds bring to the IT table.

- **Innovate:** supporting innovation is not only one of the two main challenges facing IT today (besides cutting costs), but also the most important. Therefore, many enterprises consider this flexibility to be more important than the cost savings enabled by public clouds.

Top IaaS offerings – relative business advantages (Source: Cisco)
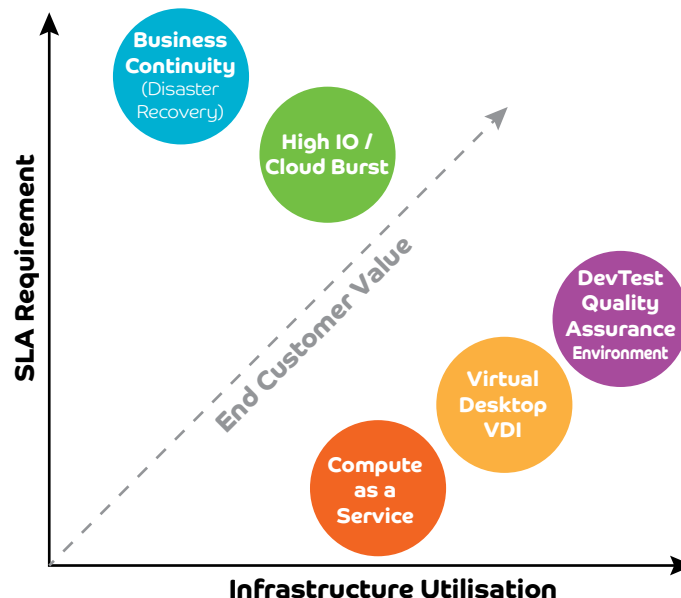
The balance of risk versus cost

Research produced for Cisco suggests that IaaS services offers are already proving attractive to both service providers and enterprise customers. Notably, business continuity and disaster recovery represents a low-risk service for the first adopter (since they still own primary infrastructure) and for the service provider, it offers a high-value service.

Persuading senior executives to loosen purse strings and spend scarce resource to improve data backup and disaster recovery, for instance, can be difficult, because that spending will subtract directly from company profits. This is frustrating for CIOs and IT staff, especially as they are the people who will be responsible for fire-fighting to recover data and applications after a major data loss or disaster. After such an event, the ability to say that the company should have heeded their calls for more spending on backup and disaster recovery will be of little value, and will make nobody happy.

There is a strong case to be made here for IaaS.

Spending on backup or disaster recovery only delivers a return on investment when there is a data loss or disaster. Even then, the return is only to keep an application or organisation afloat. Those two facts invalidate standard return on investment (ROI) calculations. However, value in return for cost must still underpin all decisions about spending on backup and disaster recovery. The fundamental decision is about the balance between the level of protection, and the cost of achieving it. While there are technical considerations, ultimately that is a business decision.

Choosing the balance between cost and level of protection will be based on an estimate of the financial impact of data losses or application outages. Although this estimate should be framed in the IT terms of recovery point objective (RPO) and recovery time objective (RTO), it is a business calculation, which can only be completed by application owners or lines of business chiefs. This will not only perform the essential step of quantifying the penalty suffered after data loss or application outage; it will also encourage non-IT staff to fully appreciate the need to spend money on data protection.

Business chiefs should also estimate the financial impact of data loss or application outages for at least one set of RPOs and RTOs, and ideally for a handful of different RPOs and RTOs. As well as giving the IT department a better understanding of the impact of outages, performing this calculation for multiple pairs of RPO and RTO values will help non-IT business managers to understand the differences between levels of protection. Similarly, businesses should estimate the impact of total site outage or disasters. An organisation that has no idea of how it would operate after a disaster is flying blind. Exactly as for data backup, asking business chiefs to quantify the impact of a disaster will sharpen their focus on the issue.

While non-IT management should make the ultimate decisions about the balance of cost and risk or levels of data and disaster protection, decisions about how best to achieve that balance are in the domain of the IT department. This requires an analysis of the costs and benefits of continuing with existing in-house systems, or switching to an online or cloud backup service. The line items should include:

1. IT staff hours managing backups, within data centres and remote offices

2. IT staff hours handling requests for data recovery

3. software maintenance and renewal costs

4. hardware maintenance and depreciation costs

5. hardware space, power, and cooling requirements

6. tape media costs

7. offsite tape storage costs, as applicable

8. offsite data replication, as applicable

9. disaster recovery colocation costs, as applicable.

IaaS stacks up well against this. Any changes in data protection levels achieved by switching to cloud-based IaaS backup may affect operational costs and productivity for non-IT staff, it has to be said, because it may alter the time taken to wait for data recoveries. And those changes should be estimated. But there is a strong financial case for IaaS-based backup or disaster recovery.

Charges for online backup services vary widely, beginning from around $0.25GB per month for basic services suitable for small organisations or remote offices. Comparisons of costs between services need to account also for the following factors:

•   Charging is commonly based on the volume of the backup copies of data after compression and data-deduplication – but sometimes can be based on the volume of the original data.

- Charges may include fees for recovery of data beyond specific monthly limits on data volume, and these are more likely for enterprise services.

- Seeding or reverse-seeding – the shipment of data on physical media at the beginning of a service contract, or for recovery of large data volumes – carries varying charges.

Business of all sizes are recommended to carry out an assessment of Iaas-based backup or disaster recovery.

The Pros and Cons of IaaS-based Back up

Model	Example suppliers	Pros	Cons
Existing backup tools linked to service providers	Symantec , EMC, CommVault, Acronis	Ease of implementation Strong backup features, e.g. granular recovery, application-level backups Reduced lock-in	Ineffective de-duplication may cause heavy usage of bandwidth and cloud storage
Existing backup tools linked to public clouds via gateways	StorSimple, Panzura, Riverbed, Ctera	Ease of implementation Ability to also use gateway for primary data storage	Ineffective de-duplication may cause heavy usage of bandwidth and cloud storage
SaaS backup	Autonomy, EVault, IBM, HP and various service providers	Low management overhead Ability to switch between providers that use the same software platform Option to license software for in-house use	Need to replace existing backup systems Bandwidth reduction may also be an issue

Summary View

Moving ahead with cloud as part of an enterprise IT strategy, including something as relatively core as IaaS offerings, is neither as difficult as it may sound nor as easy as some vendors would lead customers to believe. Customers that have existing contracts or prior experience with specific service providers will naturally gravitate toward investigating their cloud services, and their preferred local telco supplier is a very good place to start.