Benefits and Challenges of Cloud Deployments

Part 3 of 8 – Application Hosting and Cloud Strategy

SIEMENS

This is part 3 of an introductory series of documents intended to assist your organisation in defining your Application Hosting and Cloud Strategy. Your organisation may already have such a strategy, in which case these documents will hopefully confirm you are on the right track or they may identify challenges your organisation faces.

Benefits and Challenges of Cloud Deployments

There are many reasons that organisations choose to deploy in the cloud and equally there are perceived challenges that prevent organisations from cloud deployment. This section details these perceived benefits and challenges, these have been generalised from analysts & vendors views.

Benefits of Cloud Deployments

Nine key benefits have been identified:

1. **Increased Agility**: Cloud services can be deployed and consumed very quickly allowing organisations to derive value from new initiatives much more rapidly;
2. **Consumption based delivery**: Cloud-based services have no up-front costs and are charged as they are used. When combined with resource elasticity organisations get an effective utilisation rate closer to 100%. Industry benchmarking has shown that a typical on-premise deployment runs at about 15% utilisation where-as a typical Cloud deployment runs at 65% utilisation. Usage metering by service allows full visibility of service costs to business units and therefore makes it simpler to make application cost vs benefit judgements;
3. **Elasticity**: Cloud environments can scale up and down rapidly (in many cases automatically) to handle workload volatility and imbalances in demand. This rapid scaling, when coupled with almost infinite supply, means that organisations no longer need to worry about building infrastructure to meet peak demand only to see it standing idle for much of its lifetime. Applications can be designed to 'burst' into the cloud to consume additional resources when needed and scale back down during 'normal' activity;
4. **TCO and Opex smoothing**: Utilising cloud environments allows organisations to avoid the peaks and troughs of CapEx based purchasing cycles in preference for a more constant and predictable OPEX cost model. This avoids the inefficiencies of having to size environments for future expansion, instead Cloud services can be "right-sized" with the flexibility to dial resource up and down as required. Other on-premise costs associated with managing infrastructure and expanding data centres can also be off-set. Up-front costs for new services can be dramatically reduced, helping to de-risk Proof of Concepts;
5. **Accessibility**: Cloud environments are built for browser and mobile delivery, providing flexibility of access and ease of use on almost any device from anywhere that has internet connectivity. Cloud also offers easier federation with partners enabling organisations to deliver services to anyone that needs to consume them;
6. **Innovation and Currency**: Cloud services are automatically kept current by the provider and typically have much more frequent new feature and function updates. This allows organisations to benefit from the latest innovations without the additional burden of traditional release, change and test management;
7. **Sustainability**: Hyper-Scale Cloud service providers (AWS, Azure, Google) need to drive out energy efficiency in their data centres to be profitable so they are typically much more efficient than on-premise equivalents. The sheer scale of CSP infrastructure means that any slight efficiencies deployed

are much more likely to be cost-effective and therefore implemented. Cloud deployment could provide an off-set for many organisations looking to reduce or minimise their environmental impact;

8. **Visible and contractual SLAs**: Cloud providers have well defined service levels associated with their services, these often include penalty payments when these are missed by the provider;
9. **Portability**: Applications are increasingly decoupled from the infrastructure and environments they run in. Containerisation and development of microservices also enable workloads to be portable between on-premise and Cloud as well as between cloud providers. This allows consumers to look for the best 'deal' and helps protect against vendor lock-in.

1. Increased Agility

One of the major features of cloud environments is the speed and ease of access that is provided to the consumer. This provisioning feature delivers infrastructure and services whenever they are required without being governed by traditional purchasing cycles. This means that new concepts can be quickly tested and either rejected, or taken forward to full development much faster. These features rely on deploying infrastructure/services as code which means that a cloud service catalogue needs to be created and maintained.

Principle	Public or Private Cloud deployment will be the default where speed to value is important.
IT Organisation Implications	IT needs to establish itself as a 'cloud broker' that is trusted by business units to help solution selection and manage risks and maximise the perceived benefits of cloud.
	Appropriate governance will need be in place to ensure that cloud services are deployed and managed in a cost-effective way to avoid "bill shock".
	A cloud service catalogue will need to be created with appropriate governance in place to ensure that deployments are using the catalogue appropriately.

2. Consumption Based Delivery

The two key features of consumption based delivery are:

- The potential to re-deploy services in a much more cost-effective way in the cloud by "right-sizing" infrastructure components and consuming platform services;
- Providing much more granular and accurate costs for each service (show-back) so that business units can see the true running costs.

Show-back is a key tool in driving more value from IT investment and the flexibility that cloud provides around scalability could help further streamline costs by allowing business units to scale down to balance the value vs cost equation. These figures could ultimately be used to move IT costs back into business unit budgets via 'charge-back' but at this time there isn't the intention to pursue this.

Principles	Public or Private Cloud deployment will be the default option for new services to gain visibility of cost.
	Migrating existing on-premise solutions into Public or Private Cloud should be considered where there are potential cost savings.
IT Organisation Implication	New cost "show-back" processes will need to be developed to ensure business units have visibility of the cost of Private and Public Cloud based solutions.

3. Elasticity

Elasticity is most likely to provide cost and efficiency savings for workloads that vary considerably over a period (e.g. weekend peeks, end of month processes, seasonal variants).

Testing environments are also only typically required for short periods and could be shut down or scaled right back when they are not in use.

Public Cloud is preferred over Private Cloud as this is where elasticity is most practical, Private Clouds are run on dedicated infrastructure that has an upper capacity limit so savings are more difficult to derive from elasticity.

Principles	Public Cloud will be the default option for non-pre-production test environments.
	Scaling back services over-night will be considered where cost savings can be made.
	Auto-elasticity will be considered for cloud native solutions to drive out cost savings.

4. TCO and OPEX Smoothing (Benefit) & Financial (Challenge)

Reduced TCO is not an automatic product of moving to the cloud, in fact evidence suggests that most 'migrations' to the cloud are broadly cost neutral. There are however many opportunities to increase the likelihood of reduced TCO, these can be summarised as:

- Looking for opportunities across the current estate to migrate applications to the Cloud;
- "Right-sizing" or thinly provisioning applications that are migrated or deployed in the Cloud;
- Reducing the scale of the on-premise deployment and saving on associated operating costs and avoiding the costs associated with scaling these facilities out;
- Leveraging other benefits to drive out savings (e.g. elasticity);
- Close financial governance and effective configuration management of Cloud deployments.

Comparing TCO between on-premise and Cloud environments is not straight-forward as the comparison is not like for like. In order to maximise the benefits and address the challenges in this area the following principles will be established:

Principle	Cloud deployments will be right-sized in the Cloud based on current requirements, monitoring and scaling out when necessary will replace the need to over-provision.
IT Organisation Implications	Mechanisms for realising savings associated with shrinking the on-premise environment will need to be established with new IT outsource providers.
	IT costs will be increasingly OPEX focused and this will need to be incorporated into financial plans.
	Appropriate governance, including effective configuration management, will need be in place to ensure that cloud services are deployed and managed in a cost-effective way to avoid "bill shock".

5. Accessibility

Principles	Public Cloud will be the default option for services that require delivery to multiple partners.
	Identities related to 3rd parties will be managed in Azure Active Directory (or similar Identity Management software) through B2B and B2C services.

6. Innovation and Currency

From a commercial perspective the latest innovations could lead directly to revenue streams and for back-office functions new features often assist automation and efficiencies.

Many organisations have key areas of the operational business where change is controlled very tightly to protect operational stability. For these areas, the "evergreen" Cloud model presents some challenges as changes to services are not optional, they are simply part of a rolling programme where consumers are informed of upcoming change and given little flexibility on when these are put live. In some cases, consumers can choose to be in later phases of a release and stability can be easier to guarantee as the service is common amongst all consumers so issues can be flushed out by early adopters. Sometimes changes may require business change (e.g. testing, re-training, comms) and it is the ability for business units to respond quickly to these and prepare the way for new updates that could be particularly challenging.

One of the downsides of constrained change is that systems gradually get more and more out of date. This leads to supportability and security issues over time which is arguably just as destabilising. It is therefore a balancing act between staying up to date and minimising the impact of change.

Principles	The ability for business units to absorb change associated with "evergreen" cloud services will form part of the options assessment process.
	By default, cloud services will be preferred in order to benefit from staying current and getting access to the latest features and innovation.

IT Organisation Implication	Alternative change, release, configuration management and testing processes will need to be developed for cloud based services and a similar shift in attitude towards change within business units will need to be established to account for “evergreen” services.

7. Sustainability

In order to deliver IT expansion sustainably your organisation must become a much more efficient user of IT, significantly increasing resource utilisation and reducing aggregate energy consumption.

The main Hyperscale Cloud providers have been expanding and developing data centres at an incredible rate over recent years. From a financial perspective, it is critical for these providers to be as energy efficient as possible and to that end they have driven down the power usage efficiency (PUE) of data centres well beyond what is achievable for most on-premise data centres. However, these providers also replace hardware at a rapid rate and therefore need to be responsible in the sourcing and disposal of this equipment so the overall sustainability picture needs to be considered. These providers have recognised the importance of sustainability to consumers and have strong sustainability targets themselves so compared to on-premise equivalents, Hyperscale Public Cloud services represent a net gain for sustainability.

Principles	Public Cloud services from Hyperscale providers will be preferred to leverage sustainability benefits.
	The sustainability credentials of cloud service providers will be considered as part of the selection criteria for new services.

8. Visible and Contractual SLAs

While cloud providers often offer service credits or financial penalties for not attaining published SLAs, this is more of a token gesture as these credits rarely cover financial or reputational damage that could be incurred by organisations. Actual availability of services offered by the Hyperscale providers is far better than published SLAs which makes published SLAs less relevant (although this may not always be the case). They may be more relevant for smaller SaaS providers and your organisation will need to continue to assess SLAs and associated terms and conditions around penalties and credits.

Principles	Public Cloud services from Hyperscale providers such as Microsoft’s Azure, AWS, Google, IBM will be preferred as these services consistently exceed SLAs.
	Service Levels and associated penalty and credits will be considered as part of the selection criteria for new services.
IT Organisation Implication	Processes for service management will need to be developed to account for the different ways cloud services are supported and underpinned with fixed SLAs.

9. Portability (Benefit) & Vendor Lock-in (Challenge)

The portability of services underpins the risk of vendor lock-in, so the easier it is to move services the less the risk of getting locked in to one vendor. It should be noted that different types of cloud services have different portability traits:

- IaaS services tend to be highly portable with many tools available to move workloads between providers;
- PaaS services are typically much less portable, they tend to be vendor specific. This is slightly mitigated by the fact that these services have quite granular functions which helps break down the complexity of migration;
- PaaS services that provide containerisation for code or applications (which typically represent microservices in Cloud Native solutions, see section **Error! Reference source not found.** on Cloud Native) are highly portable, there are several market leading technologies (e.g. Docker), many of which are supported across all of the hyperscale cloud providers;
- SaaS services are no worse in terms of portability than traditional on-premise applications as data migration and functional/logical mapping still represent the biggest challenge.

With a majority of IaaS and PaaS there are options to deploy with code, commonly referred to as 'Instructure as code', as well as being useful for deployment this also assists with portability as you have an abstracted description of what is deployed that could be ported to other platforms.

Principles	Select the primary platform for PaaS & IaaS Public Cloud deployments.
	Develop and maintain an "exit plan" that will detail how services could be migrated in the event that the CSP's pricing becomes uncompetitive or services are scaled back.
	Infrastructure as code principles will be adopted to aid deployment and portability.

Challenges of Deploying Cloud Solutions

Seven key challenges have been identified:

1. **IT Operating Model**: Most IT organisations that are built around managing on premise systems are ill prepared for managing cloud based services. A change to managing a hybrid environment is very significant and requires new skills, processes, policies and potentially new service partners. This represents a large scale IT transformation for most organisations;
2. **Security and Compliance**: The security of data is a major concern to most organisations and having complete control over where the data is and how it is protected is preferable for some over trusting 3rd party controls. In some cases regulatory conformance and compliance constrains cloud deployments;
3. **Service assurance**: Service Levels offered by cloud providers are in most cases fixed and may not be equivalent to internal service levels where there is more control over people, process and deployed technologies. Even if they are equivalent, some organisations are concerned about the lack of flexibility to step up service levels if future demand warrants it;
4. **Performance**: In some cases, it may be difficult to guarantee performance of cloud deployed services, particularly where workloads are latency sensitive or tight integration with on-premise deployed solutions are involved. This makes organisations wary of deploying mission critical applications in the cloud as the tuning and optimising of services and the resolution of performance issues may be beyond their control;
5. **Financial**: The CAPEX to OPEX shift necessary for migration of services to the cloud does not suit the financial models of some organisations, particularly where the choice to "sweat assets" based on current business performance is important;
6. **Integration**: For most organisations managing a traditional on-premise environment, the existing application estate has a complex set of integrations and inter-dependencies. Teasing these apart in order to enable deployment of appropriate services in the cloud creates some technical challenges around legacy integration methods or data volume/latency that might not be supportable in the Cloud;
7. **Vendor lock-in**: Although CSPs provide many portable services (as above) there are other services (particularly PaaS offerings) with proprietary features that are difficult to replicate elsewhere. Organisations may be concerned with providers increasing the consumption costs of these services over time with no simple option for moving these services to a cheaper provider.

1. IT Operating Model

Managing a Hybrid-Cloud environment is considerably different from managing a traditional on-premise environment, many established processes will not be fit for purpose and will require different skill sets and approaches. None of the benefits of cloud adoption can be realised without transforming the IT operating model.

2. Security and Compliance

Hyperscale Public Cloud providers have had to respond to consumers concerns about the security of their data in the Cloud as this has been one of the main reasons that organisations have held back. The general perception is that these providers now have the processes, accreditations, physical environments, services and technical mechanisms to meet a very high standard (e.g. Government approval to deploy 'confidential' level systems). They also have the scale and agility to respond to new & emerging threats and employ an evergreen principle to keep services up to date. Therefore, this is becoming less of a constraint for organisations to invest.

Governance of how solutions are deployed will be crucial to avoiding security risk, although Public Cloud providers have the backend security it is still incumbent on customers to consume the appropriate services and tools to configure & deploy solutions securely.

Principles	Reference architectures for cloud deployments will include how to correctly secure services.
	Your organisation will govern & scrutinise the security of cloud based solutions, it will not be assumed that by deploying in the Public Cloud a service is automatically secure.

3. Service Assurance & Performance

Ensuring good levels of performance, service and availability for IT services is critical for your organisation, the end-user experience and 'service excellence' is a real focus and runs through the vision, principles and purpose of IT and the broader organisation. IT are often judged on the performance and availability of services as this is the most visible to end-users.

Two main issues represent potential constraints for organisations. Firstly, the increased reliance on external network connectivity and secondly whether the reasonably conservative SLA's associated with Public Cloud services will meet your organisation's need.

Network connectivity is critical to ensure performance and service availability, currently cloud based services share the internet links and compete with browsing and other more general uses. Products such as Microsoft ExpressRoute (which allows more direct connectivity to Azure, typically provided by companies such as Equinix) can be sized to carry Office 365 traffic, and the links will be resilient and diverse. This will move traffic off the main internet links. Any solution built on Azure (e.g. O2P) could also be routed through ExpressRoute but the amount of bandwidth provisioned would need to be reviewed to safeguard performance. Latency is a potential issue for applications that are not based on web technologies (e.g. classic client/server architectures) or rely on proximity to edge devices.

Service levels provided for many of the PaaS services on Azure, AWS etc may not, on face value, meet business needs. It is possible to deploy services in a number of high availability configurations which mitigates losing any particular component and services can be deployed across multiple Azure, AWS, IBM

Cloud regions. The aggregate availability profile of a service can be considerably improved, although it is only the individual components that are service assured by the provider.

With on-premise deployments, your organisation has full control over service assurance and are able to build in any level of high availability and service which is deemed worth the investment.

Public Cloud could offer a good opportunity to provide additional failsafe facilities for on-premise solutions which may be particularly pertinent for critical services that require a disaster recovery option. The fact that the platform will not share any common infrastructure or be reliant on data centre facilities (with the exception of external network services) makes this a good option for certain failure scenarios. This is also relevant for individual components of a solution, for instance witness nodes for clustering technologies could sit in Public Cloud as a valid additional location. This hybrid approach may also be relevant when more critical systems are deployed in Public Cloud and fail-safe facilities could be provisioned on-premise.

Overall a majority of availability and performance needs can be met through aggregated Public Cloud services as described above and these are arguably better maintained and more stable than on-premise equivalents.

Principles	New Public Cloud (e.g. Azure, IBM Cloud) based solutions will consider bandwidth requirements and will drive scaling up the private connection between datacentres and infrastructure (e.g. ExpressRoute) to ensure existing solutions are not affected by increasing demand.
	Where service assurance and performance of a system is 'mission critical' to the operation of the organisation, on-premise deployment in a physically segregated environment will be preferred.
	Public Cloud deployment is suitable for services that have high availability requirements.
	Suitability for a service to be deployed in the Public Cloud will not be judged on individual component level SLA's but will be based on an aggregated view of components in high availability configurations.
	Deployment patterns that leverage Public Cloud based failover for on-premise solutions (or individual components of solutions) and vice versa will be explored.
Roadmap Item	Implement private connections between datacentres and infrastructure (e.g. ExpressRoute) with diverse and resilient connectivity to improve service assurance for Public Cloud (e.g. Microsoft Azure, AWS, IBM Cloud) based services.

4. Integration

An integration strategy that aims to simplify and consolidate on-premise and Cloud based solution integration has already been published and is based on Azure technologies. Some of the constraining factors include:

- Legacy patterns of integration (e.g. point to point);

- Legacy methods of integration (e.g. protocol/bespoke technology);
- Latency sensitive integration;
- High data volume integration.

These need to be addressed as a matter of course but may constrain what can be done proactively ahead of full deployment of new integration hubs. Generally, the view is that integration will be no more of an issue than it is currently.

Rheinberry specialises in assisting IT leaders in organisations with complex environments to successfully plan and deliver IT change across their organisation. We provide our clients what they really need, costs representing good value, flexible ways of working with proven IT professionals and most importantly successful Project and Programme outcomes.

To do this we have combined the proven methods, professionalism and structure of traditional IT consulting practice with the experience, pragmatism and independence of IT freelancers, interims and contractors, to create a new 'hybrid consulting' model. Our approach and working style is very much 'client-side', we collaborate and engage deeply with stakeholders and we strive for long term relationships. We believe in small teams of highly skilled experts with not only the right experience but also those that have worked together previously and know each-other's strengths and working styles that collectively deliver quicker than a team that has not worked together previously.

www.rheinberry.com

info@rheinberry.com

Rheinberry Ltd
International House
24 Holborn Viaduct
London
EC4A 4AN
United Kingdom

+44 (0)20 3000 2660