

# Applied Cloud Computing

**A practical guide to identifying the potential  
in your environment**

**Dale Vile and Tony Lock**

Freeform Dynamics Ltd

October 2010

*Sponsored by:*

**Microsoft®**

## Introduction

The term 'cloud computing' is one of the most widely used in the IT industry today. To some, it means revolution and a fundamental shift in the way IT services are delivered to both businesses and consumers. To others, it's just about marketeers dressing up old ideas in new clothes in an attempt to encourage another round of spending on technology and services.

The truth is probably somewhere in between. While there is no denying that a bandwagon is rolling, and many suppliers and pundits are jumping on it whether they are justified in doing so or not, some very interesting developments are coming together that are extremely significant. Whether these ideas are totally transforming the world of IT is debatable; perhaps a more accurate description would be to say that a range of significant developments are taking place that together will enable a step change in the ongoing evolution of IT service delivery.

Within this paper, we will be reviewing some of these developments, and with a view to the more practical concept of evolution rather than revolution, will be looking at how things can be worked into mainstream IT activities in an incremental manner. Our aim is not to 'sell cloud', but to simply provide guidance on what it translates to in terms of tangible and actionable specifics.

As we get into the discussion, you may pick up ideas on where some of the capability often referred to under the cloud computing umbrella might be used to deal with immediate needs. Given that cloud computing, as we shall see, is as much an approach as it is a range of technologies and services, you may even discover that you are already 'doing it' at one level or another.

Beyond immediate needs and tactical starting points, however, we will be looking at how cloud computing concepts, services and technologies are likely to impact the way IT departments are organised and operate over the medium to long term. Whether we are talking one year, three years or five years until significant impact occurs, depends partly on your circumstances, and partly on how quickly the market and IT professional community moves forward in terms of solutions, skills and experience. Whatever view is taken on this, indications are that things are already moving forward in an irreversible manner in the cloud computing domain, so it's probably best for IT professionals to be prepared wherever they are starting from.

## Inputs into this paper

This paper draws upon the experience of the Freeform Dynamics analyst team advising mainstream IT departments on the practical application of emerging ideas and technologies in a business context. We also refer to a number of research studies (see Further Reading section) in which information was gathered from IT professionals, including early adopters of cloud solutions, on the context and implementation of the various ideas and approaches we shall be discussing.

One such study looked at what exactly people understood by the term 'cloud computing', and to make sure we are properly aligned on terminology before going any further, this is a good place to start our discussion.

## What do people mean by cloud computing?

Many people refer to cloud computing in a very natural and easy manner as if everyone automatically knows what they are talking about. The trouble is that 'cloud' means different things to different people, and all too frequently conversations occur in which the parties involved are actually referring to quite different interpretations.

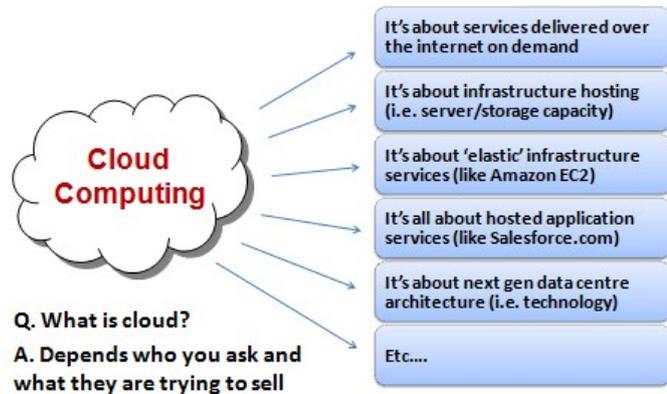
When one person says, for example, that the biggest problem with cloud is knowing where your data resides, and the other one disagrees saying that hardware budgeting and procurement is the most difficult challenge, it can get a little confusing. Until, that is, you figure out that the first person is talking about hosted application services in the so called 'public cloud', whilst the second is referring to the creation of a more dynamic 'private cloud' infrastructure within their own data centre.

If you are not clear on those terms we have just used, i.e. 'public cloud' and 'private cloud', don't worry for now as we'll be walking through them shortly. We'll also pin down some of the other

terminology that is in common use in this area such as Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS), and so on, as well as touching on 'elastic clouds', 'hybrid clouds' and some of the other language that suppliers, consultants and analysts seem to use quite freely and assumptively.

While going back to basics like this might seem a little tedious, it is important to be clear on what we are talking about when looking at the practical application of cloud computing. The truth is that the biggest challenge at the moment is one of understanding. As we have said, it's actually quite a broad domain, encompassing various ideas, approaches and solution types, which sometimes have nothing to do with each other, sometimes overlap, and sometimes even conflict. The latter is particularly true when suppliers and others with vested interests try to narrow and twist the definition of cloud to fit with what they have to sell, while simultaneously trying to exclude offerings from the competition. So depending on who you ask, and arguably what they are promoting, you are likely to get a range of definitions of cloud computing thrown at you as you explore the area.

If you are in any doubt about the ambiguity and inconsistency that exists, it is worth reviewing the findings of a recent study<sup>[1]</sup> we conducted at Freeform Dynamics to investigate the extent to which perceptions of cloud vary. In this, we described over 30 different offerings (stripped of their marketing speak and spin) and asked 400 IT professionals to tell us whether they regarded each as a legitimate example of cloud computing. The results are summarised in Appendix A and what's most interesting about them is that for each item, without exception, there were people who regarded the offering as a legitimate example of cloud, and others who disputed this.



This level of variation in views around cloud computing has created a number of challenges:

<b>Challenges around definitions and messaging</b>	
Confusion and uncertainly	Inconsistent and contradictory terminology can be confusing for those trying to understand the area and figure out what's relevant to them and how.
Oversight and missed opportunity	Artificially narrow definitions can be latched onto, creating the illusion of understanding and certainty when in fact some important ideas and offerings have not been considered.
Misunderstanding of practicalities	Generalisations proffered with regard to pros, cons and practicalities can be misleading as these often vary by the type of cloud computing being considered, and this is not made clear.

The other challenge with cloud computing is that supplier and pundit messaging tends to be long on vision and short on practical detail, with 'magic bullet' propositions often coming across as quite simplistic and detached from reality.

If you know or suspect that you have encountered any of the challenges we have highlighted here, then read on. We're going to start by breaking down the overall cloud computing proposition into its constituent parts as a precursor to exploring practical application.

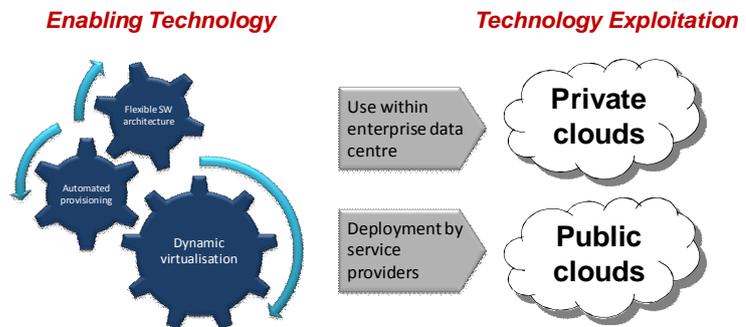
## Perspectives on cloud

A useful trick to help make sense of everything that is claimed and offered in this space is to realise that suppliers, consultants and commentators are often approaching the topic from different perspectives, and the key is to be clear on the perspective or dimension being discussed. In our experience, the three most useful perspectives to be aware of are as follows:

- Technology versus services
- Architectural stack perspective
- Functional service taxonomy

The first of these is about distinguishing between enabling technologies such as dynamic virtualisation, automated provisioning, flexible software architectures, etc, and the systems or services that result from their exploitation. In terms of exploitation itself, cloud technologies can be used by service providers to create hosted services delivered 'on demand' via 'public clouds', and by organisations directly to create the aforementioned 'private clouds', i.e. flexible pools of resource within their own data centres. Either way the end result is what we might describe as more of a 'utility' approach to providing the IT resources and/or the application functionality required by an organisation to support its business.

The second perspective on cloud has its roots in the traditional systems stack, and provides a useful way of understanding the various types of hosted services in particular from an architectural point of view.



As we look across the service provider community, we see suppliers operating at various levels. Some, for example, are simply delivering compute and/or storage resources on demand, the main aim being to provide an alternative to in-house infrastructure, i.e. servers and storage devices running in your own data centre or computer room. What you do with those resources is then up to you – apart from providing capacity in the form of virtual devices of one kind or another in line with an agreed service level, the IaaS provider is doing little else.

Services in the IaaS domain vary in terms of the degree of physical and commercial flexibility offered (e.g. with regard to mechanics and cost of acquiring or releasing resources as demands change), but some are extremely flexible, allowing resources to be allocated and de-allocated with ease, perhaps even automatically on a continuous real time basis. This is a principle that many refer to as 'elasticity', and it is often accompanied by a 'pay as you go' commercial approach.



As we move up the stack from the underlying infrastructure, we next reach the platform layer, which in traditional on-premise systems is where we find middleware, database management, security management and other facilities that applications require to run effectively and securely. The corresponding 'PaaS' concept simply takes all of that, and provides it as a hosted cloud service, typically complete with the necessary development, deployment and management tools.

PaaS has particular potential for those needing to build and execute applications that fluctuate significantly in terms of demand, e.g. consumer facing Web applications, as providers tend to underpin them with the kind of elastic infrastructure described above.

The top of the traditional systems stack is where we find the application layer, and it is common for people to speak of the equivalent of this in cloud terms being ‘SaaS’. A simple way of thinking about SaaS is therefore as the domain of hosted applications delivered in the same spirit of flexibility that we have been discussing in relation to IaaS and PaaS. Just as in the traditional sphere of application software, however, we have a myriad of options and possibilities in the application services space – from simple business productivity tools at one end to full blown enterprise applications like CRM and ERP at the other.

Given this, the stack or architectural view of the world starts to run out of steam as a vehicle for helping us understand the types of service on offer. It is therefore necessary to switch to our third important perspective on cloud, the taxonomy view, in which services are categorised more from a functionality or capability perspective:

ON-DEMAND SERVICE CATEGORIES	
<b>Business application services</b>	It is in this area that the term ‘Software as a Service’ (SaaS) was originally coined. Services at this level are typically focused on the delivery of complete business functionality, e.g. CRM, ERP, etc.
<b>Hosted productivity tools</b>	Services here are more concerned with horizontal capability ranging from desktop suites for end users, through to modelling, development and project management tools for analysts and developers.
<b>Hosted comms/collaboration</b>	Spearheaded initially by hosted email and web conferencing, the number of services offerings in this area has exploded to include full unified communications and/or social media (directories, blogs, wikis, etc).
<b>Trading community services</b>	As supply chain automation has gathered momentum in some industry sectors, services have emerged aimed at facilitating the way in which customers and suppliers collaborate and transact electronically.
<b>Plug-in services</b>	A myriad of services exist which do not provide complete business functionality but ‘plug into’ existing applications to enhance or extend them. Examples include everything from mapping to credit checking.
<b>Application platform services</b>	As an alternative to consuming pre-built services from external providers, application platform services provide development and runtime environments allowing custom applications to be built and hosted online.
<b>Operational services</b>	This often overlooked but highly important category is where we find services concerned with online backup, archiving, security (e.g. email filtering), etc., and even full blown monitoring and management tools.

While even this taxonomy is not exhaustive, it illustrates how, by moving beyond the more generic ‘as a Service’ definitions, we can make more sense of what is available and where offerings might fit into the enterprise solutions landscape. It also underlines the danger we highlighted previously about too much generalisation, as many of these services are quite different, and need to be evaluated and assessed according to correspondingly different criteria.

## Cloud in context

One of the most common arguments we hear from those promoting cloud is what we affectionately refer to in Freeform Dynamics as the ‘stars are aligning’ story. It starts out by highlighting the great advances that have been made in the areas of hardware efficiency, virtualisation, software architecture, and management capability. It goes on to point out that all of this is changing the economics of hosting, allowing service providers to provide ever more flexible and cost effective alternatives to the traditional on premise approach. Put this together with a greater focus among customers on cost management and the optimisation of business services, and QED, the stars are aligning to make cloud computing the centre of everything looking forward.

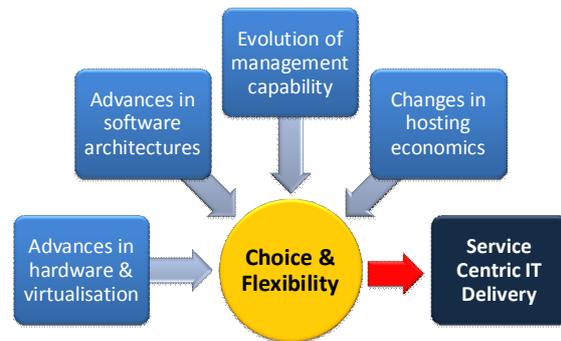
It’s a very convincing and indeed compelling story, and it is hard to argue with its constituent parts. However, there is a more objective way of looking at things that leads to a different conclusion. While many of the things we have been exploring under the cloud umbrella come about as a direct result of the industry advances called out, the real outcome is not a wholesale shift to ‘the cloud’, but the creation of choice and flexibility.

And if we take this further, what is in turn enabled is a more 'service-centric' view of IT service delivery.

So what do we mean by this? Is it just more management consulting gobbledegook?

Well no. In actual fact, it is a concept that is very much in keeping with the way the relationships between IT departments and their customers within the business have been developing over recent years. We are referring here to the increased focus we have observed through our research in many organisations on thinking less in terms of systems, and more in terms of services and service portfolios, against which expectations, requirements and service levels can be more naturally defined in business terms.

Another way of putting this is a shift in the IT-Business conversation from the 'how' of IT to the 'what', on the premise that business users really don't care that much about the componentry that enables the IT facilities they use, provided they have access to the right capability presented in a usable and reliable manner with the right level of performance and support.



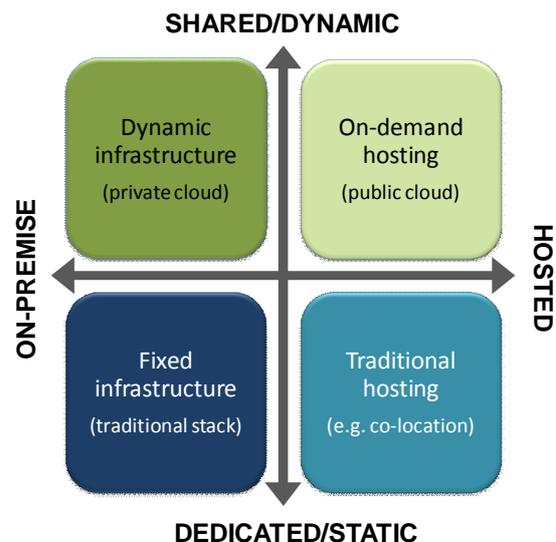
And this is the nub of matter when it comes to the significance of cloud computing. When it is put into its proper perspective, it's not the technology and services associated with cloud that matter the most, but the choices and freedom they create to mix and match approaches and solutions to meet business needs in the most optimum and flexible manner.

## Choice in action

Few organisations, even smaller ones, have the luxury of wiping the slate clean and starting again with IT. While those responsible for running the systems infrastructure might indulge themselves in such a fantasy from time to time, the reality is that new ideas and ways of doing things generally build on what's in place already.

With this in mind, it is less about considering a 'move to the cloud', which is the kind of language we often hear in this space, and more about extending your IT delivery approach to embrace cloud options as you continue to evolve and modernise your systems landscape, and consider the best way of dealing with new business requirements.

A good framework to think through the practicalities of this is the quadrant shown on the right. Most IT requirements today are met through the fixed infrastructure approach that is depicted in the bottom left hand corner of this diagram. Hardware and software here are generally quite tightly coupled, and this remains the case even if virtualisation technologies have been used to achieve a level of consolidation. There is nothing necessarily wrong or bad with this model, and it continues to have relevance as many applications and workloads have relatively fixed and predictable resource requirements, and are therefore well served by this traditional approach.



Such fixed and predictable requirements can also be dealt with effectively through the traditional hosting model in which a set amount of resource is 'rented' to host specific applications or

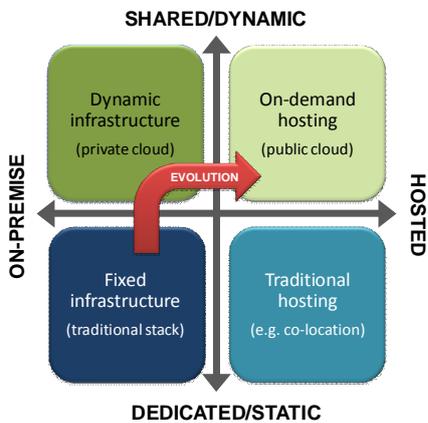
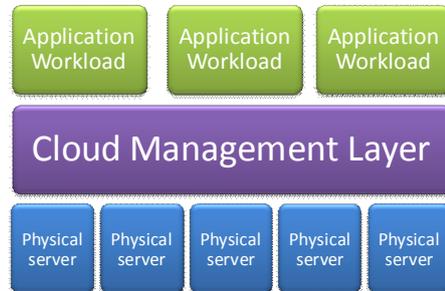
application portfolios. Again, this approach, which has been established for decades, continues to be relevant when you consider the hosting of large bespoke systems, complex and tailored packaged applications (e.g. ERP), mainframe environments, and so on. Hosting in this space is often coupled with full managed services in the spirit of total outsourcing.

It's in the top two quadrants that new options are beginning to emerge, however. If we take the top left to begin with, this is the space in which we see advanced virtualisation and 'private cloud' technologies allowing the creation of a much more flexible internal IT environment. At an infrastructure level, we are specifically talking about the creation of resource pools, in which a subset of the server estate, e.g. a farm of x86 servers, is managed as a single virtual resource, or 'cloud' that can provide compute power on demand. The idea is that applications are totally decoupled from individual servers, so each workload receives just the right amount of resource required as it needs it.

In the conceptual diagram shown on the right, which presents a deliberately simplified view of private cloud, the block we have labelled 'Cloud Management Layer' is actually comprised of a number of different technologies. This includes advanced management tools to allow rapid provisioning (i.e. the preparation of physical servers so they can participate in the running of a given application), as well as load balancing capability to ensure that work is distributed appropriately across the resources allocated. The exact nature of such solutions and the degree to which they support automated or semi-automated resource allocation varies by vendor and by the inclination of the organisation to exploit them.

The private cloud approach as we have described it here is in many ways an evolution of first generation virtualisation technology, and indeed server virtualisation is an important precursor to private cloud computing. Those who have invested in virtualisation already, e.g. spending time on the virtualisation readiness of applications, and the management and monitoring tools that go with this, will find that a lot of that effort will be directly exploitable as the more dynamic virtualisation approach that underpins private cloud is embraced.

### Private Cloud Concept



Furthermore, a move to private cloud provides a natural stepping stone into the on-demand hosting model that many refer to as 'public cloud'. If we look at the evolution of mainstream IT, we therefore expect the progression shown to the left to represent the most common way forward for most organisations.

To be absolutely clear, however, we are talking about an extension of the IT delivery framework, not the migration of everything into to the top right hand quadrant. This clearly wouldn't make sense as the kind of flexibility offered by on-demand hosting will invariably come at a price premium, particularly if the pay as you go approach to payment is part of the equation. For more fixed/predictable requirements,

on-demand hosting is likely to be more expensive than the traditional hosting (bottom right quadrant) approach. This has nothing to do with technology, it is simply the several thousand year old commercial principle that the more you commit to buying in one transaction (or in this case service agreement), the lower the price you pay.

But even the evolutionary path we have outlined from virtualisation through private cloud to public cloud is a simplification. For those interested in taking advantage of cloud services sooner rather than later there are some interesting short term opportunities. Some of these become apparent when we look at the types of applications and workloads early adopters have already found

represent a good fit with the on-demand hosting model. In another recent Freeform Dynamics study<sup>[2]</sup>, for example, the following were called out:

**Lightweight workloads:** e.g. workgroup applications that only occasionally 'tickle' the server they sit on. On-demand cloud hosting is an alternative to traditional virtualisation for consolidating these

**Periodic workloads:** Applications that spin up in a demanding way for discrete periods of time, sitting idle in between. Typically batch oriented (e.g. billing and payroll), but can be online (e.g. periodic administration activity).

**Bursty workloads:** Applications that create a continuous load, but one that fluctuates significantly. The most common examples here are consumer facing Web applications, though solutions ranging from call centre support to telemetry can also fall into this category.

**Compute-intensive workloads:** Number crunching applications that benefit from parallel processing, which are often ad hoc or periodic, making on-site special equipment hard to cost justify.

**Development and testing:** Particularly where there are transient requirements for developers and testers to emulate target environments for a period of time as part of the application build/test/release cycle.

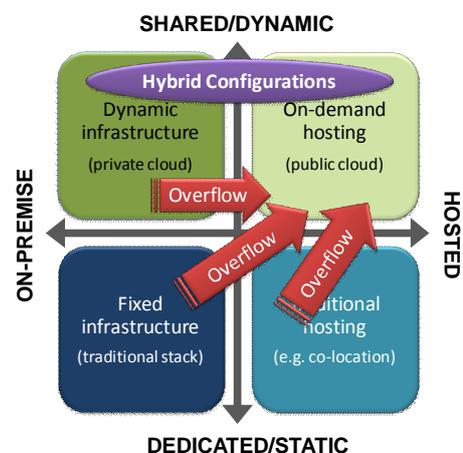
As new requirements emerge in areas such as these, there may be opportunities for you to consider the on-demand cloud hosting option in a very controlled manner, without having to worry about the wholesale movement of a significant part of your systems landscape. The last area on the above list, development and testing, for example, represents a particularly low risk place to start if there are cultural and/or practical concerns about going down the hosted services route (so long as no sensitive data is involved).

On a specific point, it is also worth considering emerging PaaS services, that naturally live in the top right hand quadrant of our diagram, for either lightweight or bursty workloads in particular. Not only does this overcome the problem of having hardware sitting around idling because of over-provisioning to cater for peak demand, PaaS environments are specifically designed from a software and data architecture perspective to cater for requirements from the very small to the very large, without the need to code for scale. They are therefore very well suited to the rapid prototyping approach that often helps when working in a fast moving Web application context.

Lastly on the topic of choice, it is important to acknowledge that not everything needs to fit into a single box. While some applications and workloads will naturally fit into one quadrant or another, a significant amount of activity is likely to cross boundaries in most mainstream environments. With dynamic virtualisation as the common denominator between private and public clouds, it becomes possible to spread activity between the on-premise and hosted environment. This is useful if some parts of an application are dynamic and others are relatively static, e.g. the web facing or business intelligence elements of an ERP landscape versus the traditional back office functions. It also provides scope to cater for overflow scenarios, in which applications for the most part run on premise, or in a fixed hosting environment, but can 'grab' more resource from the cloud when demand exceeds normal capacity.

The possible scenarios are actually wide and varied, and go beyond the ones we have mentioned here, which further underlines the principle of choice and flexibility.

But before we get carried away with the possibilities, we are talking about technology and not magic here, so there are some very practical implications to consider.



## Practical considerations

One of the problems with cloud computing at the moment is that most of the commentary and advice that is published tends to focus on the benefits, only paying lip service to the practicalities and challenges. However, part of evaluating if, where and how the different forms of cloud computing we have been discussing might fit into your plans and activities, is having a clear understanding of the real world implications.

At the highest level, given the dynamic hybrid vision we have described, with cloud technology and services coexisting with traditional infrastructure and hosting, one of the most obvious questions is what should run where as things evolve. The discussion so far in this paper will hopefully have provided a high level steer on this, but when it gets to the next level, there is a need to model workloads and requirements a bit more precisely. The best advice we can give when doing this is to start with the business service definition (in line with the service-centric IT delivery approach) and work backwards to the systems level requirements.

Apart from looking at the resource demand profile of applications, it will also be necessary to consider potential constraints and issues. It is beyond the scope of this paper to be exhaustive here, but some of the main things to think about include:

- Any relevant constraints with regard to information and storage, e.g. there may be compliance or security considerations associated with the maintenance of data in a service provider environment. Regulation may impose certain restrictions on where some kinds of data may reside, or at least mandate that you know where it is stored, which can be a problem with providers whose infrastructure crosses national boundaries.
- Practical and commercial implications of data transport and storage. Where high volumes of data are concerned, for example, it is generally more performant and economical to bring the processing to the data rather than vice versa. Routinely shipping huge volumes of information into a hosted compute cloud for number crunching can actually work out more expensive and time consuming than building a private cloud infrastructure to do the same job.
- Appropriate due diligence before signing up to cloud services. With a myriad of players out there ranging from niche start-ups, through consumer focused providers dabbling in the business space, to big global hosting companies, figuring out what's on offer and the constraints under which it is being provided can be a minefield. The advice here is to make no assumptions about what's included in a service (e.g. backup, archiving, customisation tools, level of support, etc) and always think ahead. One of the most common mistakes is committing to a service that fulfils an immediate tactical need, without considering what happens when your requirements escalate and something different or more sophisticated is required. It is all too easy to find yourself locked into a proprietary service, with no easy way to switch providers or bring the application (or your data) back in-house.
- Integration across domains, either with hybrid configurations or simply when applications need to communicate across boundaries. In one Freeform Dynamics study<sup>[3]</sup>, for example, IT professionals were generally more comfortable with the SaaS proposition in relation to discrete applications such as sales force automation, and were much less certain about solutions with many integration touch points, such as ERP. At the root of this were concerns about physical integration (development, maintenance and support of interfaces) and the associated question of accountability (i.e. who is responsible for what when problems arise and/or interfaces need changing), but issues with visibility and policy management across domains were also highlighted.

This last point highlights a critical enabler of effective cloud computing, namely an appropriate IT management environment in terms of processes and tools. Those that have moved along the route from traditional IT into more virtualised infrastructures are already likely to have discovered that the management game changes<sup>[4]</sup>. Tracking, maintaining and monitoring a mix of physical and virtual assets in a coherent way requires tools that support a holistic view of the environment, complete with all of the dependencies. And the challenge here escalates when we extend into the more

dynamic private cloud approach, where everything is more fluid and mappings and dependencies change on an almost continuous basis.

Add to this the increasing need for end to end monitoring and management as a result of the service-centric IT delivery approach, including root cause problem analysis to troubleshoot the more complex environment, and the hard reality is that traditional management tools and processes are unlikely to be up to the job. As you move forward with a more dynamic cloud style infrastructure, it is therefore critical to allocate time, resource and funding to the management side of the equation, as neglecting this will not only make it difficult to realise the potential we have been discussing, but will also create some very real risks.

The degree to which you might want to or be able to monitor and manage assets in a service provider environment is debatable. Arguably, with the right service level agreements in place, that level of activity should not be required (though we realise there is a big due diligence question here, and there will always be exceptions). One element that will need managing across the on-premise and service provider domains, however, is provisioning and resource allocation. This is particularly the case if you are interested in taking advantage of hybrid configurations and overflow capabilities – e.g. if a Web application runs out of resources in your data centre because of an unanticipated peak in demand, you are going to need the capability to provision virtual servers in the hosted domain and move workloads onto them with ease.

Fortunately, some of the larger players in the market have been working on such capability for a while now, so provided certain compatibility criteria are met, the hybrid and overflow models can be implemented today (though again, you need to do your supplier due diligence).

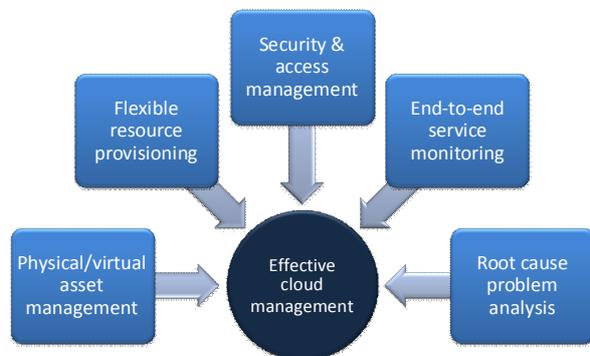
On a specific point, one of the areas it is necessary to pay particular attention to with cloud computing is security and access. Beyond the obvious question of whether a cloud service is secure per se, there is the question of policy management. The last thing you need, for example, is the overhead and risk of maintaining one set of policies and rules in-house, then another for each provider you make use of. Again, this is a management issue, and while standards and best practices mature, there is a need to be careful in this area.

These are just some of the areas to consider as we have said, and the main take-away from this is that cloud computing does not magically make all problems disappear (as some seem to claim). Furthermore, everything we have been discussing can significantly affect the way things are done in IT from an operational perspective in particular. But contrary to the line we often hear, cloud does not 'change everything', so let's take a more balanced look at the most likely impact on the IT department.

## Impact on the IT department

In order to consider the likely impact on the IT department, we shall make the assumption that the IT services landscape is headed towards the blended delivery approach, with a mix of activity across all four of the quadrants we have highlighted. In terms of changes, this basically means two things – the in-house infrastructure will become more virtualised and dynamic, and more use will be made of external hosted services. Then, of course, there is an implied need for everything to work together effectively.

Given that IT departments are multi-faceted in nature, with many disciplines and processes involved in the IT delivery equation, it is most useful to look at how cloud computing will affect the activities of specific functions. In the following table, we therefore walk through each of these and look at the implications that are likely to be most significant, bearing in mind some of the practicalities we were discussing above.



<b>Summary of impact on core IT functions</b>	
Architecture, design and development	The importance of a well thought out architectural approach is accentuated as cloud options are embraced. Structure and consistency (with regard to server, networking and storage infrastructure) is key to building effective private clouds, and is critical to achieving interoperability with IaaS services. A more architectural approach to software design and development is also necessary to allow ease of integration between on-premise and hosted service domains (no application is an island), and to permit migration of application workloads both within and between resource pools. Sound architectural and development expertise is therefore an important pre-requisite for success with cloud computing, indeed driving down the cloud route without a good understanding of how everything will hang together is a recipe for nasty surprises down the line.
Security, identity and access management	The fundamental principles in this area remain largely unchanged, but the scope and complexity of their application is extended. In the short to medium term, it is unrealistic to expect all but the most mature cloud service providers to implement mechanisms for automated policy integration, e.g. through directory synchronisation. When we consider the ease with which end users can unilaterally sign up to cloud services, it is also unrealistic to expect IT to have control of everything that's used. Skills and experience will therefore be a requirement to assess hosted service related risks, figure out what's necessary to mitigate them, then define the appropriate policies and implement the necessary integration.
Storage and information management	Storage management is an often neglected area that cannot remain so with cloud. As the infrastructure becomes more virtualised and dynamic, including the storage elements, the skills and expertise to design, implement and manage effective tiered storage architectures (which may possibly extend into the IaaS domain) becomes doubly important. At the next level up, content management and information governance skills are critical to ensure compliance and integrity as data moves into and around the various internal and external clouds. There is an obvious loop that needs to be closed here with security and access, and another back to storage as things like backup, recovery, archival and discovery are considered in relation to data maintained in a hosted environment.
Systems monitoring and management	Beyond the need for enhanced monitoring and management processes to cater for the dynamic virtualised server, storage and networking environment, the increasing presence of hosted services creates additional challenges. Those responsible for IT service delivery will need to become adept at defining SLAs with providers and monitoring performance against them, often with limited visibility into the provider environment. Understanding and mapping dependencies between hosted services and on-premise systems, and between hosted services themselves, adds an additional level of complexity to the game of end-to-end monitoring and management, and, of course, troubleshooting.
End user support	Some aspects of end user support, e.g. in relation to the desktop and on-premise business application functionality, remain largely unchanged. However, troubleshooting issues in a dynamic environment when dependencies are transient probably means more hand-offs to other groups (e.g. in operations). There is then the question of integration of support between the in-house help desk and hosted service providers, which is important to resolve to prevent user confusion and dissatisfaction.

Our treatment of the impact on IT presented here is by no means exhaustive; we merely aim to illustrate that cloud mostly changes the scope and emphasis of activity rather than completely restructuring it.

## Discussion and Conclusion

When we consider the ground we have covered in this paper, it is clear that far from being just marketing hype, cloud computing represents a tangible set of developments and opportunities that are likely to be relevant to most organisations looking forward. As we have also seen, however, cloud doesn't revolutionise everything, and is best thought of, and indeed embraced, as an evolution of capability that will sit alongside traditional IT infrastructures moving forward.

As part of this, any notion of cloud computing devaluing the role IT professionals is a bit of a myth. Perhaps there will be less need for some routine manual activities around server and storage provisioning and maintenance, but as we have seen from our discussion about the impact on core IT activities, new requirements emerge elsewhere which can be higher value from a business contribution perspective.

Beyond the technical disciplines, we do, however, expect a significant impact from an overall IT management perspective. It is a fair bet, for example, that for most organisations, especially larger enterprises, their needs will not be met by a single service provider, or even a handful of providers. As the number of service contracts and supplier relationships escalates, this will change the emphasis of the IT leader role. CIOs will need to act more as orchestrators of services rather than commanders of empires.

The need for this orchestration approach becomes even clearer when we consider the accessibility of cloud services. Local departments and workgroups, even individual users, can now sign up to take advantage of online applications completely independently of IT. Whether it's personal communications, social media, cloud storage, business functionality or any other of the endless options now a credit card transaction away from adoption, this is already a fact of life in many organisations.

This kind of ad hoc adoption of cloud services has obvious risks and potential overhead associated with it, though trying to prevent it is both practically difficult and politically risky. Come down too hard, and there could easily be a backlash from users and business stakeholders, with the same activity simply going underground. In response to this, some of the more forward thinking CIOs we have spoken with have highlighted visibility as one of the biggest priorities – i.e. as much as possible, tracking what employees are using.

The idea here is not so much to police or act as 'big brother', but more to monitor where and how to step in. When a particular service is observed to be gaining popularity, then it may be time to embrace it and make sure it is properly supported and integrated, and that appropriate policies are defined for its use. When similar but fragmented activity is observed, that may be a prompt to put an enterprise class solution in place to deliver the same capability. Of course if particularly risky behaviour is seen, then that may have to be blocked, but the idea is to acknowledge that IT no longer has absolute control so a different approach may be appropriate.

Our discussion of cloud would not be complete if we didn't touch on the area of IT governance and financing. While the service centric approach makes conversations around business requirements and service levels much more natural, it also disrupts the way most organisations have accounted for IT in the past. No longer do we have an application running on dedicated hardware with the cost of the whole stack allocated to a given department or cost centre. The fluidity of dynamic cloud infrastructure destroys the traditional notion of physical IT asset ownership by specific parts of the business. In these early days of private cloud, IT departments are working around this by playing the old game of tucking horizontal infrastructure spend under individual project or system budgets, but such 'frigs' are not sustainable over the longer term.

The unilateral adoption of SaaS at departmental level may actually be of help here, as it engenders a mindset of paying for what you use. This may pave the way for more of a chargeback accounting approach to be adopted, though this has implications in terms of systems and process overhead. A preferable outcome would be organisations investing in IT infrastructure and services as a genuine corporate asset, looking at overall business contribution and return on investment, and not getting too obsessed with granular cost-centric accounting. More mature organisations already work in this way, and it is probably no coincidence that in these same entities, there is a high degree of trust and harmony between IT and the business.

As a final thought, while there is a great deal of emphasis in the industry at the moment on technology and services, cloud computing is as much a philosophy as anything else. While there are some clear opportunities to take advantage of specific offerings, the most profound impact of cloud comes about as a result of the choice and flexibility it introduces, and the service-centric approach to IT delivery it encourages.

We hope that this paper has helped you to understand both the promise and practicality in this highly important area of IT.

## References and Further Reading

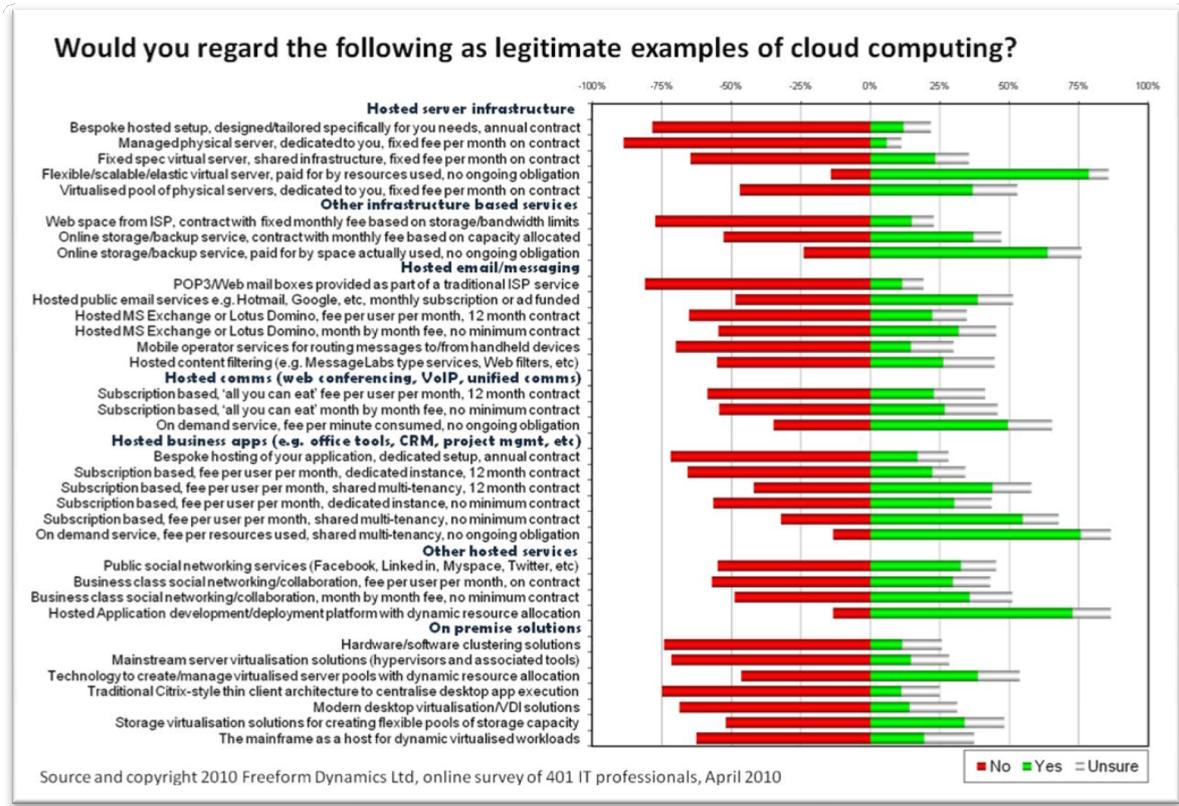
The following reports and papers are referred to in this paper:

1. **But is that really cloud computing?**  
*The problem of ill-defined terminology*  
<http://www.freeformdynamics.com/fullarticle.asp?aid=1068>
2. **Understanding the market for elastic cloud services?**  
*Will Amazon really inherit the earth?*  
<http://www.freeformdynamics.com/fullarticle.asp?aid=1105>
3. **IT on the front foot**  
*Sourcing, architecture and the progressive IT organisation*  
<http://www.freeformdynamics.com/fullarticle.asp?aid=318>
4. **Evolution of hosted server computing**  
*The emerging cloud alternative*  
<http://www.freeformdynamics.com/fullarticle.asp?aid=1104>

All of these, plus a range of other research based reports and papers, are freely downloadable from the Freeform Dynamics website at <http://www.freeformdynamics.com>.

## Appendix A

The following chart summarises the results of an online survey of 401 IT professionals conducted in April 2010 to investigate the clarity of cloud computing related messaging and terminology.



For more details of this study, please see the report entitled "But is that really cloud computing?", which is downloadable from [www.freeformdynamics.com](http://www.freeformdynamics.com).

## About Freeform Dynamics



Freeform Dynamics is a research and analysis firm. We track and report on the business impact of developments in the IT and communications sectors.

As part of this, we use an innovative research methodology to gather feedback directly from those involved in IT strategy, planning, procurement and implementation. Our output is therefore grounded in real-world practicality for use by mainstream business and IT professionals.

For further information or to subscribe to the Freeform Dynamics free research service, please visit [www.freeformdynamics.com](http://www.freeformdynamics.com) or contact us via [info@freeformdynamics.com](mailto:info@freeformdynamics.com).

## About Microsoft



Founded in 1975, Microsoft (NASDAQ "MSFT") is the worldwide leader in software, services and solutions that help people and businesses realise their full potential.

### Terms of Use

This document is Copyright 2010 Freeform Dynamics Ltd. It may be freely duplicated and distributed in its entirety on an individual one to one basis, either electronically or in hard copy form. It may not, however, be disassembled or modified in any way as part of the duplication process. Hosting of the document for download on the Web and/or mass distribution of the document by any means is prohibited unless express permission is obtained from Freeform Dynamics Ltd.

This document is provided for your general information and use only. Neither Freeform Dynamics Ltd nor any third parties provide any warranty or guarantee as to the suitability of the information provided within it for any particular purpose.