

Cloud isn't the answer to every question

**A Pragmatist's View of IT Delivered
as a Service**



Abstract

Cloud is different from earlier shifts in the way that it is notably less about the technology and more about the business

This paper's purpose is to help readers consider cloud services in a pragmatic way that has relevance to today's IT landscape.

It seems like cloud marketing is everywhere and the hype means that even Corporate Boards are asking about 'this cloud thing'. In addition, the economic situation means that budgetary pressures are not easing and, all the while, the accessibility of cloud-based services means that user departments are increasingly procuring functionality directly. The sum of this is that cloud is different from earlier shifts in the way that it is notably less about the technology and more about the business. In this environment, objective engagement is the only way for IT departments to influence these initiatives.

Capgemini does not view cloud services as the answer to every problem. They are, however, sufficiently mature to be considered as part of the answer in most enterprise IT environments. Our 2008 paper 'An Early View of Cloud Computing', suggested that such services are an extension to rather than replacement for current delivery platforms. Two years of experience has validated that while the majority of cloud services still provide extensions to the existing IT landscape, in several areas of operation, cloud services have now reached a level of maturity and robustness where they can be considered as a direct replacement to current platforms. In practical terms, we see a Compute Continuum that runs from owned discrete assets dedicated to a function, through best practice in estate management, virtualization, consolidation, multi-tenancy, and eventually to public cloud in its truest sense.

In many ways, cloud is the culmination of several strands of technological progress that together represent the fulfillment of a much desired feature set. Virtualization is now a mature technology, industrialization has reduced manual intervention to near zero in the life of a virtual server with concomitant improvements in scale and quality, whilst automation and intelligence make systems self-governing within approved boundaries. From this, end clients benefit from:

- operational flexibility;
- quicker return on investment;
- avoidance of capital outlay;
- cost reduction;
- greener IT.

Capgemini believes that cloud services have gone past their tipping point because we are receiving client invitations that may not specify cloud delivery but do require the associated benefits; in particular, avoiding capital expenditure, requiring flexible operations, seeking variable-use and consumption-based pricing. In the long term, even the best operational practices within single estates cannot meet these expectations. So, although big businesses want to build their confidence in the cloud through an evolutionary not revolutionary approach, their economic expectations make that evolution inevitable.

Capgemini strongly believes that cloud-based services offer our clients the opportunity to build powerful custom solutions from vendors' standard elements – in effect systems integration for the cloud model. Backing up our client-first approach and vendor-agnostic positioning called the Collaborative Business Experience™, Capgemini was the first global systems integrator to offer partnerships supporting both Google Apps and Amazon Web Services (AWS). We have since added specialists like Cordys, RightScale, Salesforce.com, Eloqua and Kognitio. These round out a cloud ecosystem bedded on powerful, mature alliances with Microsoft, Oracle, SAP, HP, IBM, Cisco, EMC and VMWare who are all developing their established portfolios to encompass cloud services specifically for enterprises.

The Compute Continuum

NOT ANOTHER DEFINITION

Many papers, including previous ones from Capgemini, have tried to provide a definition of Cloud Computing. Consequently, there are many definitions and this paper will not add to the noise. The US National Institute of Standard and Technology (NIST) has a definition¹ that seems less controversial than most and provides a great starting point for further debate².

In defining Cloud Computing, NIST lists five 'Essential Characteristics':

- on-demand self-service,
- broad network access (for the end purpose – the Internet is not necessarily required),
- resource pooling (i.e. some form of multi-tenancy),
- rapid elasticity, and
- measured service (use-based charges);

three 'Service Models':

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

and four 'Deployment Models':

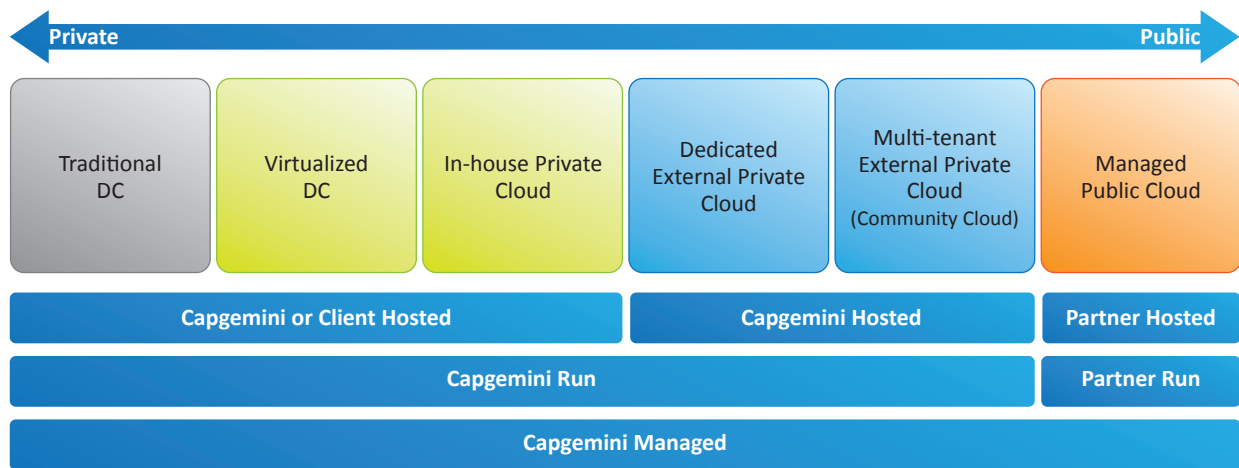
- private cloud (private' refers to use by one organization, it does not have to be owned by or located on that organization's premises),
- community cloud (use by a group with similar requirements, e.g. support or security),
- public cloud, and
- hybrid cloud (the mixing of deployment models).

Clearly, when judged in those terms, many of the products on the market today look somewhat less than cloudy. Practically, it is unhelpful to make any assessment in binary terms as businesses will naturally want to pick and choose the characteristics that make sense to them. In addition, there is the question of granularity when assessing against these criteria. For example, does multi-tenancy have to mean different companies or just different billing departments inside the same company? It is because of this ambiguity that Capgemini proposes the Compute Continuum.

¹ <http://csrc.nist.gov/groups/SNS/cloud-computing/>

² The Jericho Forum's "Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration" is an excellent paper to further explore cloud models.
http://www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf

THE COMPUTE CONTINUUM



SHIFT RIGHT

Presently, our clients have most of their resources concentrated to the left of the Continuum. Moving right, infrastructure becomes increasingly flexible, whether simply to run a more efficient estate or to enable application modernization.

The first step is the now well-established process of server consolidation through virtualization. As clients virtualize their compute, pressure rapidly grows to move further right by increasing the level of industrialization and automation; it simply makes no sense to build virtual machine (VM) estates with any other method than from a library of standardized, centrally maintained, high quality, master VMs. With that increase in tooling, some of the 'essential characteristics' in the NIST definition start to emerge – notably broad network access within the organization and the potential to provide on-demand self-service resources with use-based charging via some form of portal. With these services available, it is easy to see how many IT teams start to brand this as an in-house private cloud.

As noted by NIST, the fact that a service is private does not mean it has to be deployed in-house but only that it has to be used solely by one organization. Capgemini Outsourcing has been providing a virtualized server solution that looks very similar to this level of cloud for the last two years.

Tooling can be very expensive and not every organization can afford to build or buy-in their own fully capable cloud infrastructure. An alternative is to use a service provider who can deliver an enhanced capability through a shared management platform across dedicated compute and storage resources. These dedicated but external private cloud resources may offer self-management (autonomic computing), better billing granularity (down to the server hour or gigabyte of data), more managed machine types, a better reserve capacity and, of course, technical service desks that offer end users access to better support around the clock in different languages.

There are typically two cost breaks in the process of 'shifting right' along the Continuum. The first is achieving critical mass in virtualization and the second is the mutualization of resources through multi-tenancy in either a private or public cloud. It is this second break that Capgemini is addressing now to meet client needs for a competitive, flexible platform. In NIST terms, the creation of a 'Community Cloud' within Capgemini data centers allows us to address requirements including privacy, location specificity in relation to legal jurisdictions, service and operational level agreements, and collocation with legacy systems. Along the Continuum, a multi-tenant private cloud is currently the best compromise of cost and trust for many clients as functionality builds out and operational enterprise experience of cloud services grows.

Of course, as first-to-market with key public cloud managed services, Capgemini advocates their use too. In our experience, most clients can put certain loads or functions onto these services today without having to turn current delivery models upside down. The trick, obviously, is to do so in a managed way whether the use is for production or the provision of sandboxed environments for business functions to use or test within.

By now, you hopefully have a good sense of what Capgemini means by the Compute Continuum and how the trend for systems is to Shift Right over time. There are three related points of finesse within this:

1. Hybrid solutions are inevitable. The nature of differing application and system requirements means that one size will not fit all. Consequently, complex IT delivery will see elements scattered along the Continuum with potentially some specialist services staying on-premise indefinitely, others to a private cloud, while the core front and back-office functions transfer to a public cloud footprint.
2. Different types of services will Shift Right at different rates. CRM, messaging and people management all have great traction in the SaaS space today and are rarely hosted privately but instead use the massive scale of key cloud vendors. At the base of the NIST Service Model stack, IaaS is well developed and offers incredible flexibility with aggressive commercial terms. In a sense, the limitation on these virtual infrastructures is imagination because of the open, flexible elements available for developers to build stacks on. It is worth noting that IaaS solutions have to lead ahead of cloud applications though; it is impossible to build dynamic applications without an equally flexible infrastructure. Of the three NIST models, it is the possibilities offered by PaaS that remain intriguing. Clients like the idea of porting code and data rather than whole architectures but the rapid evolution of these services tends to push our clients towards evaluation rather than implementation; they need stability to plan against and to be sure that they aren't spending money solving challenges that the vendor will bake into their offer in short order.

3. Not all services are created equal as reflected by their pricing. Cloud definitions do not generally specify the level of granularity for service measurement. The bigger the slices of time, capacity or any other metric, the less complicated the billing tools need to be, reducing the investment that the provider needs to recoup. In contrast, bigger slices equal stability and the more stable the load on a service, the more cheaply it can be run. The greater the demand for elasticity, the more resources have to be held in standby, whose costs have to be built into the service model. Greater security or more stringent auditory requirements will also inevitably drive up costs. The list of complexities could go on but the point has been made: the better a client knows their requirements the easier it will be to select a service that is both suitable and competitive.

It's About the Apps

Capgemini wants to create 'graceful' applications that can scale with demand by provisioning and releasing resources as needed

Interestingly, the easy categorization of infrastructure, platform and software into a three-layer model seems to encourage consideration in discrete terms. However, as with the delivery continuum, they are not discrete in any practical sense.

Nowhere is this more noticeable than in the approach to IaaS. Naturally, the start point for many evaluations is to ask how a new product or service will let a company do what it does today for less money. Lifting and shifting loads onto an IaaS is an apparently simple model; yet, despite the initial numbers often looking attractive, teams quickly come up against challenges:

- How can the applications and services match the flexibility of the underlying infrastructure?
- Whilst very similar to normal Intel-based server architectures, public resources have subtle differences that require the stack to be tuned for best performance.
- If moving production environments how is continuity maintained?
- What are the regulatory, security and privacy issues?

Consequently, there's a rapid reversion to the evaluation criteria applied for test and development, storage, and other less complex services. The limitation of this approach is that the proportion of a complete estate that can be migrated unmodified onto an IaaS is small compared to the whole. As such, the savings, though good on a like-for-like basis, are marginal on total expenditure.

Although this reversion to the lowest common denominator is understandable for clarity in comparison, it cannot change the fact that the move to cloud is about applications. Businesses don't buy resources for fun but to run applications on them. This makes an Application Portfolio Strategy essential. CIOs can prioritize their business' IT evolution on different vectors. Amongst others:

- by functions that will benefit most from the flexibility of cloud enablement;
- by the applications re-platformed most simply; or
- by services already architected for commodity Windows and Linux stacks.

There is a risk, however, that an Application Portfolio Strategy will prioritize quick wins at the cost of avoiding the bigger but more profound mid- to long-term issues. This can hide some potential benefits of the Shift Right because the more difficult services are never fully evaluated.

As an example, when working with clients on their application portfolios, Capgemini normally finds considerable ongoing investment in legacy services. Indeed, older applications and their associated hardware have been remarkably sticky. Amongst other reasons, a key stumbling block has concerned capital expenditure needed for new infrastructure. Implicitly (and sometimes explicitly), however, every definition of cloud services excludes capital expenditure on the service in favor of consumption-based pricing. By removing capex as a barrier and making it easy to limit the cost of testing, the business case to refresh more difficult applications gets easier to formulate and more compelling.

Beyond the particular apps that will be good candidates for direct migration, the portfolio as a whole needs a strategy built around a staged migration.

- Re-platforming
- Modernization
- Re-architecting

In simple terms, Capgemini sees re-platforming as migration to commodity infrastructure – most likely Linux on Intel though perhaps Windows too, where significant application rework is not necessary. With efficient employment of capital ever in mind, the standard target today is a virtualized environment. Subsequently, the modernization and re-architecting of the applications is an inextricably linked process; one cannot happen without an awareness of the other.

GRACEFUL APPLICATIONS

Capgemini's target is to create clear portable code or implementations that are not just cloud aware but that are able to be part of an intelligent feedback loop with the underlying infrastructure. We want to create 'graceful' applications that can scale with demand by provisioning and releasing resources as needed.

Graceful applications are the way that Capgemini's clients will maximize their return through these services' tight linkage of demand with costs and revenues. This, quite evidently, makes it easier to manage the old rule of a business bringing in more revenue than it spends.

THE NEW NEW THING

So far, this section has looked at existing portfolios. It is important to also be aware of the persuasive arguments that the most critical client value is to be found in exploiting what might be called the 'new new' thing: to build solutions that were hitherto impossible for technological or monetary reasons. Anecdotes exist about pharmaceutical companies running up large grids temporarily for trials where the capital investment would previously have made the business case very difficult. Amazon Web Services talks about financial service companies using Monte Carlo simulations to better evaluate risk. *The New York Times* is on the record for converting 4TB of back catalogue scans into 11 million PDFs inside 24 hours for around \$240. In identifying what these new applications might solve, looking to issues frustrating the business is usually a rich seam. Spotting that genuinely new killer offer, however, is much more of an art...

Cloud Service Integration

Our Enterprise Architects refine the nuances of building in clouds to account for things like the overheads of virtualization, the latencies involved in different service components, and how new, cloud-specific services can be taken advantage of

Cloud is not a standalone proposition. As a multi-disciplinary company with both Technology and Outsourcing service lines, Capgemini is building extensions to its services portfolio that encompass cloud delivery wherever it makes sense. We have not, as Forrester Research so eloquently put it, 'cloud washed' everything that we do because most of a client's engagement with us is around more established solutions. Likely as not, that is no great surprise and our client demand acts as another proof point for the concept of the Compute Continuum.

Where Capgemini does offer cloud services the practical advantage is that upwards of 70% of the requirement is typically based on tried and tested competencies. Nothing about taking a cloud service obviates the requirements for:

- competent program and project management;
- system-wide risk management and security services;
- experienced deployment skills;
- established product expertise;
- ITIL-based service management;
- change management and help desk services;
- extensive experience with enterprise clients.

Of the remaining 30%, broadly 20% is an extension to existing skills, like the way that our Enterprise Architects refine the nuances of building in clouds to account for things like the overheads of virtualization, the latencies involved in different service components, and how new, cloud-specific services can be taken advantage of. Overall, entirely new skills represent by far the smallest element of a client's service.

This practical approach was drawn together in the launch of a Capgemini initiative called Infostructure Transformation Services (ITS) at the start of 2010. The initiative aligns expertise embedded in our service lines to enable clients to work on various levels as they evolve their services – particularly to define how to take immediate tactical advantage of emerging technologies whilst formulating an overall cloud strategy. In this, ITS envisages a three stage evolution relevant to this paper:

- data center optimization;
- virtualization; and
- cloud computing and services.

In March 2010 at the London Cloud Computing Congress, two enterprise adopters of private and public cloud services made it clear that cloud should not be the top of a CIO's hit list if existing services had not been optimized onto a virtual server estate. Not a very thick slice of industry opinion of course but a useful, objective comment nevertheless that synchronizes with the ITS approach.

Although Capgemini's client engagements cover a mix of project and lifecycle work, the common core value is the way that our experience across many clients is being mutualized so that any one client can benefit from the sum of our knowledge. Moreover, that experience is not within a discrete section of an enterprise estate but part of an overall library of established and maturing technologies and practices.

An example of that evolving best practice is the guidance of our architectural staff on the segmentation of our clients' back-end systems. Currently, management systems and SLAs are applied as a blanket across many of our clients' estates. We contend that as systems become distributed along the Continuum this is not the best use of resources. Going forward, there should be greater granularity; not every system needs to be managed as tightly or be as available as every other. Built the right way, the failure of nodes should be irrelevant because of the tolerance in the architecture – meaning that SLAs also have to evolve. The template would be desktop and workplace services where industry best practice already differentiates between a machine operator and the CEO, whether in terms of equipment, support hours, or responsiveness. This approach lends itself well to the Shift Right vision, where it is easy to envisage loosely coupled, non-critical systems moving fast to the right with considerably lower management system complexity in the destination environments. Capgemini already has a client who would have traditionally followed this single management model but, after conversation with us, has moved their SAP development environment onto a public cloud infrastructure which is managed solely through exception.

Summary

- Client buying criteria based on current economic expectations have driven providers past the tipping point for adoption of cloud in their service portfolios.
- Cloud services are not a binary decision, clients can pick and choose the delivery model and service location on private or public resources that meets their needs.
- Many services will inevitably shift right from private, dedicated to public, mutualized solutions.
- Migrating inflexible applications onto flexible infrastructure only saves money from virtualization and maybe mutualization. The real benefits are in graceful apps that can autonomically scale based on load.
- There are many providers of discrete services. There is value in considering a partner with deep enterprise experience and familiarity across mature systems as well as emerging technologies.

Recommendations

- Do not plan discrete cloud migrations at the expense of estate-wide consolidation and/or virtualization.
- Make a point of considering cloud services as an answer in all decommission and rebuild situations when infrastructure, systems or applications reach end of life.
- Make re-platforming proprietary or legacy applications a priority.
- Assume that messaging, CRM and HR are bound for the cloud.
- Do examine requirements that would apply to mutualized services such as jurisdiction, privacy and SLAs.
- Our client base would suggest that the best services to consider building or moving to a cloud base are:
 - email, calendaring and collaboration;
 - CRM and sales force automation;
 - other front-office or consumer-facing applications;
 - web portals and web consumed applications;
 - static storage, archival, information lifecycle management;
 - analytics / search services;
 - massive, dynamic / volatile processing tasks;
 - compliance services.
- Build an Application Portfolio Strategy with a target of making key, variable load applications more graceful.
- Engage the business to help bring common minimum standards to directly procured services.



About Capgemini and the Collaborative Business Experience

Capgemini, one of the world's foremost providers of consulting, technology and outsourcing services, enables its clients to transform and perform through technologies. Capgemini provides its clients with insights and capabilities that boost their freedom to achieve superior results through a unique way of working, the Collaborative Business Experience™. The Group relies on its global delivery model called Rightshore®,

which aims to get the right balance of the best talent from multiple locations, working as one team to create and deliver the optimum solution for clients. Present in more than 30 countries, Capgemini reported 2009 global revenues of EUR 8.4 billion and employs 90,000 people worldwide.

More information is available at www.capgemini.com

For more information on Capgemini's Cloud Computing
please contact:

Richard Payling
VP, Global Outsourcing
richard.payling@capgemini.com