

CIO - Discussion Document: The Private Cloud is Clearly Overhead

Does Cloud Computing hold a silver lining or will it just rain on my parade?

An investigation in what substance truly lies behind this latest IT industry's marketing catchphrase: "Cloud Computing" and what does it mean for the internal IT organization.

In recent industry conferences, analyst briefings & webinars, IT vendors have been at great pains to extol their Cloud Computing credentials to all whom will listen. But what does it mean and from where did this metaphoric terminology arise from?

It would appear that origin of the term was coined originally by Amazon, when they described their initial foray into their web delivered services as the "Elastic Compute Cloud", obviously giving reference to the elasticity in supply or dynamic scalability of these popular services. Back in August 2006 Business Week magazine featured the CEO of Amazon on their cover questioning the wisdom of "Bezos big bet?" Well it would appear that Jeff, once again, this pioneer of finding new markets with the internet, has backed another winner.

Since only it's very recent origins we have come to generally understand the primary purpose of Cloud Computing is to provide companies and consumers alike with a powerful computing platform that is delivered over the web. Both market segments can avail of this welcome powerful and immediately scalable computing resource. It provides access to popular applications and resources that traditionally would have proven cost prohibitive for use in personal or SME environments.

In our Credit Crunch fuelled need for developing new enterprise, this truly represents a groundbreaking development for any web based or technological based start-up firms, where Cloud effectively removes their major establishment cost based, market entry barrier.

However for the majority of established organisations, the cloud concept is creating a very real competitive challenge to the conventional provision of internal IT resources. The rapid adoption of cloud services across the populous has shifted the common perspective in how IT services can be delivered in terms of speed and cost. Especially as organisation leaders themselves are among this rapidly growing band of satisfied web services users; who under the current economic pressure, are viewing internal costs very differently and wondering if they can do so quick and so cheaply – then why can't we?.

It is therefore, this widespread shift in perspective throughout user organisations that will spearhead the pressure to undergo transformational change in the service delivery model.

In responding to this trend The IT marketing guru's conjured up their latest hot buzz word.... "The Private Cloud".

As the naming suggests this is simply the internal delivery of Cloud like computing services, securely confined to the organisational boundaries of purely authorised users of IT services.

However in traditional IT Marketing style, this technological progression, has been heralded as the once in a career, cataclysmic change for the entire IT Profession; or even positioned as - "The industrialization of IT" as one corporate blogger describes it, comparing Cloud computing with the provision of critical civic utility services; electricity, water, gas etc.

It's important to cut through the hype and not get carried away here. Firstly there is a major fundamental difference in relation to the comparative minimum scale efficiencies. Civic utility services require massive infrastructure investments to bring the cost of supply to affordable and viable price levels for households and businesses. The continual cycle of technological development and the economics surrounding the principles of Moore's Law of miniaturization has provided the IT data centre industry with comparatively low minimum scale efficiencies.

This means that the internal data centre can potentially provide a private cloud computing service at comparable costs to public cloud providers, provided they follow new operational disciplines in regard to architecting and managing their IT infrastructure. This shift is already being widely embraced; and the first step on the path after virtualization is the optimization of existing infrastructures which means this transformation can be self funded through returns on immediate utilization gains. Therefore one can safely conclude that the in – house data centre and internal IT function will be around for much longer than the foreseeable future.

Transformational Change

From a technological perspective cloud computing forms the confluence of several technical developments that have been evolving for some time. The pervasive availability of low cost high capacity bandwidth networks prompted the development of thin client applications and the subsequent emergence of service orientated architecture (SOA). However the trigger mechanism for this irreversible swing toward Cloud Computing has been the widespread adoption of server virtualization technologies. Unquestionably Virtualization technologies has enabled this economically fuelled shift in the re-architecting and management of an organisations IT infrastructure to deliver computing resources with much greater agility.

Comparison Chart: Conventional Computing / Private Cloud Computing

Management Aspects	Traditional Silo Based Data Centre	Next Generation Cloud Data Centre	Future Considerations
Infrastructure Strategy	Physical & Dedicated Silo Fixed Computing Resource	Virtualized Shared Utility Variable Computing Resource	Private / Public Cloud Computing
Operational Perspective	Horizontal View- IT Components	Vertical View- Integrated IT Stack	100% Integration Virtualization Led
Operational Controls	At Each IT Component Level App, Dbase, Server, Storage	Up/ Down through the Stack App, Dbase, Server VM Storage	100% Integration- Virtualization Led
Business Unit Service Model	Dedicated Fixed Platforms	Shared Variable Platform	Service Bureaux
Service Level Chargeback	Fixed Organisation Cost Notional Accountability	Variable Utility Model- Consumption Costing	Competitive Cloud Cost Benchmarking
Consumption of IT Services	Internally Generated Demand Peaks (private user network)	Externally Generated Demand Peaks (collaborative web 2.0)	Availability of Cloud for Peak Demands
Server Resource	Dedicated Servers; Fixed Resource	Virtualized Servers; Variable Resource	Increasing VM's Server Densities
Storage Capacity Provisioning	Over Provisioned Low Utilisation (<50%)	On Demand JIT High Utilisation (>80%)	More Data: Less Disk
Performance Considerations	Low Utilisation : High Performance	High Utilisation : Same Performance	Balanced Controls Doing More With Less
Speed of Change	Slow & laborious	Fast & Easy	Increased Agility
Ease Of Change	Complex Architecture	Utility Architectures	Flexibility
Procurement Process	Point Decisions & Isolated Buying Process	Supply Chain Frameworks Storage & Servers	Dynamic Resourcing; incremental installs
Energy Consumption	Organisational Cost Notional Accountability	Full Utility Model- Consumption Costing	Constant Reduction of Costs & Carbon Footprint
Competitive Outsourcing	Expensive Services Switching Costs	Containerized Applications Reduced switching costs	New Competition Economics Forcing Change

Organizational Change:

Unquestionably a progressive step change in the delivery of internal IT services is underway.

The motivation and driver for change is the considerable economic pressure from both within the organisation needing to pay less for their IT service and the emerging viable economic alternative from external data-centre service providers that are leading the way in gaining cost and energy saving efficiencies.

A growing consideration for IT management today is their organisations concern with internal consumption of carbon producing energy. This is an added competitive pressure for the IT executive. In energy efficient outsourced Cloud Utility Computing environments, customer organisations will be held only accountable for the net power they consume for the precise server and storage resource. This option potentially delivers a significant reduction in carbon footprint when compared with existing in-house operating conventions which encompasses considerable over provisioning of empty storage capacities and spare server resource.

A very real concern in the use of Public Cloud is the issue of 'security' but the reputable outsourced provider's now adhere to exceptionally high standards in the protection of their client's information on their platforms and have largely overcome this restrictive obstacle, to the point that now governments, particularly the US federal government are exploring Cloud as a real option, particularly in infrequent high volume interaction with citizen's – census etc.

Interestingly what is actually happening is that many senior IT executives are building their private cloud strategies, with an element of Public Cloud resource included for contingency purposes i.e. low cost disaster recovery, infrequent peak loads on their systems, new on-line business initiatives - where consumer demand is hard to predict and also organisations are increasingly using popular Software as a Service offerings, like CRM applications for one.

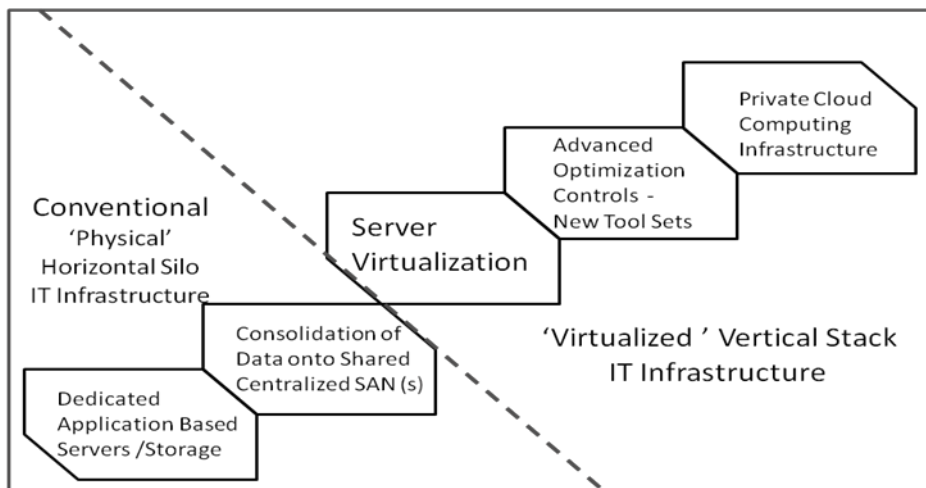
Clearly many IT executives are now adopting the development of an internal Private Cloud computing capability as their prime objective. Organisations wishing to achieve this goal will need to develop the capability within their IT function to deliver computing resources internally at comparative speed, scale and cost as service levels offered by reputable public cloud computing service providers, where a credible external cloud alternative is available.

This should be achievable as the external companies must gain a profitable return on their considerable investments and for many organisations the issue of VAT gives an advantage. Additionally as we have discussed previously, minimum IT scale efficiencies are relatively low and many of the changes become self funding through the fast cost efficiency returns.

Regardless of this internal IT operations will have to adopt a similar 'Utility' architecture as is deployed in the shared services platforms developed by the leading public cloud services providers. In this regard they will increasingly provide the blue print to follow for making changes, as they progressively transform themselves from being a co-location space and power provider of their clients dedicated environment; to becoming a full utility computing provider for hosting many users on their syndicated common services computing platform.

As we discussed earlier the IT user, particularly the senior executive's perceptions of how computing services and resources can be delivered going forward is irreversibly changing. Under this new private cloud model everything will change in terms of how IT operations will build and manage their IT infrastructure and especially how users will consume resources.

The Transformational Process

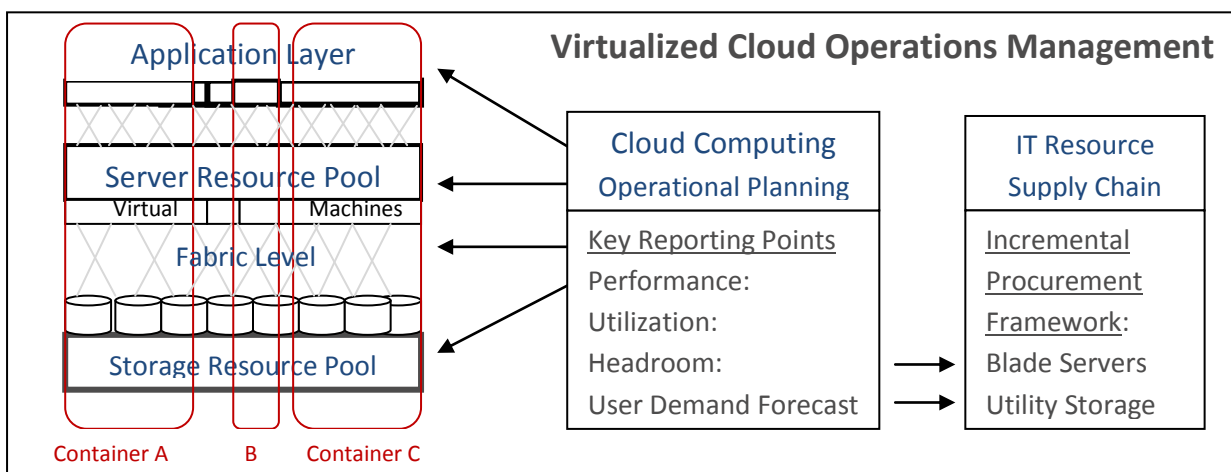


Virtualization – The Key Enabler:

VM Ware the undisputed leader in X86 virtualization have positioned their vSphere™ as being the first operating system for organisations to build private cloud computing platforms.

They equip a set of virtualized computers to provide users with the advanced ability to start and stop servers, or use compute cycles only when needed and to only pay for their usage.

This capability combined with obvious economical savings and energy reduction forms the technological catalyst for the transformational change from traditional fixed resource delivery of computing resource to a modern dynamic provision on a Just-In-Time variable basis. In many ways IT operations need to encompass critical key reporting elements of the long established MRP processes, whereby the goal is to achieve optimal output with the minimum of inputs. These critical key information elements need to be real-time reporting on capacity utilisation and available capacity headroom but in this environment performance reporting at each element in the applications virtualized infrastructure stack becomes absolutely vital.



Advanced Optimization & Controls

This burgeoning requirement for independent Infrastructure reporting has created an entirely new segment in the IT infrastructure industry. Specialist companies have emerged in recent times and are successfully challenging the use of expensive IT vendor's proprietary systems

for monitoring servers and unwieldy Storage Resource Manager (SRM) tools. These new challengers have made huge advances in automated, agent-less infrastructure analysis and are almost a magnitude of cost less expensive, which justifies their use in a matter of weeks.

As also previously discussed the primary server and particularly the storage vendors, are not well motivated to dramatically increase the levels of utilization of existing systems. What they do provide to help with this increasing operational issue, is a solution in the guise of them generating lucrative professional services revenue, probably using these independent tools. It is very clear this critical management capability must be developed and retained in-house.

With virtualization there clearly comes the added operational complexity of breaking the dedicated physical connectivity between servers and associated storage. Despite publicised savings in reduced server count, recent analysis is showing that provisioning of storage related to virtualized servers has actually reduced utilization levels / increased wastage. This compromise of the original justification virtualization proves the need for new analysis tools.

Most industry analysts predict that server density levels and storage utilisation levels will increase dramatically and these new breed of low cost automated tools are the enablers.

Private Cloud Computing Infrastructure:

The Cloud concept has been enthusiastically seized on by the primary infrastructure vendors which means they see an opportunity for extraordinary profitability, by developing their own integrated stack or cloud computing components; which primarily comprise of the following;

Virtualization Operating Environments	<i>Key Strategic Future Battleground (Emerging players)</i>
Physical Servers	<i>Efficient Blade Technologies (Competitive Market)</i>
Storage Systems	<i>Utility Designed Systems (Agile Provisioning)</i>
Security Systems	<i>Concentration of Leaders (R&D Spend is key)</i>
Management Software	<i>Proprietary IT Vendors Vs. Independent Open Vendors</i>
Professional Services	<i>Core Competency vs. External Control = Outsource?</i>

The big strategy debate at present; is whether organisations should build their private cloud with a single IT vendor or establish a policy of using the best in class for each component? In a recent analyst briefing a prominent infrastructure vendor stated their motive to enter a joint venture with another major player to fill out their stack was to uncover new profit pools. The lead must be taken by the Data centre organisations, that solely exist by delivering cloud computing services on their infrastructure and they are taking the best of breed option.

Operations Management academics consider operations best contribute to organizational strategy by excelling on each aspect of the Five Production Performance Objectives Model. Therefore Infrastructure stack decisions should be based on the highest aggregate rating for each of the cloud computing components; to decide best of breed based on strategic value.

The Five Production Performance Objectives Model

(N. Slack et al Warwick Business School 1995)

Operational Aspect	Strategic Value	Component Rating
Ability to Change	Flexibility Advantage	1/2/3/4/5/6/7/8/9/10
Speed of Change	Agility Advantage	1/2/3/4/5/6/7/8/9/10
Doing Things On Time	Dependability Advantage	1/2/3/4/5/6/7/8/9/10
Doing Things Right 0 Errors	Quality Advantage	1/2/3/4/5/6/7/8/9/10
Doing Things Less Expensively	Cost / Price Advantage	1/2/3/4/5/6/7/8/9/10

Manufacturing / Operations Management & Academic Perspective;

I recently presented the Private Cloud Computing concept to a non-IT industry audience of seasoned manufacturing management professionals and business school 'operations management' (OM) academics. They considered my portrayal of this private cloud concept and the IT conventions of which it seeks to replace. They concluded this to be nothing other than IT professionals embracing the proven economics and principals of modern operational management disciplines. This is where meeting user demand is synchronised with sufficient supply of computing resource and performance on a variable MRP type supply chain basis.

They continued to expound on this consideration by making the observation that for the first time, in the now lengthy (50+ years) existence of the IT industry, the focus of the IT function appears to be shifting from them making their organisation more productive, to them actually having to be much more cost efficient and a highly productive function in themselves. They considered that these operational trends are clearly being forced by the prevailing economic climate, and is characteristic of the IT industry's maturation and its former aura of mystique and resulting in 'super profits' gathered by proprietary IT vendors being confined to history.

They were greatly amused as we discussed the widespread conventions of over provisioning expensive storage capacity resources. Where recent studies have shown capacity utilisation on primary database systems were as low as 15% and rarely moving above 50% for the entire lifespan of the power hungry storage array systems. They voiced the poor economics of such practice, based on both environmental costs and the constant price erosion of this technology. They rightly considered this aspect to be the priority for senior management intervention with the introduction of closer reporting capabilities to eliminate such wastage.

I did however explain that this inefficient convention arose due to the combined function of complexities in changing capacity provisions on conventionally architected storage arrays and until recently a distinct lack of automated and affordable real time control monitors. Both of which greatly restricted IT operation professionals from being able to critically balance an increase in capacity utilisation, whilst maintaining overall system performance levels. The lack of tools also made capacity allocation changes laborious and error prone. Consequently best practice was considered to over provision on a just-in-case basis. This is now ceasing.

We were all in agreement that this gap in critical operational reporting was probably likely to have persisted for so long, as it was not in the interest of IT vendors to change this profitable practice of selling approx 50% more hardware than was required, at a much more favourable point on the technology price erosion curve. Complexity generated lucrative spare capacity.

After brief systems demonstrations; there was general acceptance that these new reporting tools and their widespread adoption will drive capacity efficiencies across existing IT infrastructures, certainly denting hardware sales in the short term for established vendors.

There was no surprise that these innovations to deliver the balanced optimization controls have been developed by emerging independent specialist players rather than the large established infrastructure providers. It was also considered 'Independence' a critical aspect.

Discussion Paper Written by Mike Higgins,

Founder of Sanalyse Technologies Ltd. and a regular industry commentator on developing new operational efficiencies to gain competitive advantage within Data Centres.